NSP
Network Services Platform
Network Functions Manager - Packet (NFM-P)
Release 24.11

Classic Management User Guide

3HE-20021-AAAC-TQZZA
Issue 1
December 2024

Legal notice

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

Contents

About this document

Purpose

The *NSP NFM-P Classic Management User Guide* provides information about using the NFM-P to manage service-aware IP/MPLS networks, including GUI operations, device and network management, and policy and service management.

Scope

The scope of this document is limited to the NFM-P, which can accomplish many configuration, monitoring, and assurance functions.

Document organization

This guide contains the following volumes:

- Getting started—contains general NFM-P information including the following:
  - a system overview
  - basic GUI operation instructions
  - workspace customization and GUI/topology map management
- Device management—contains information about device functions that are not directly related to networking, including the following:
  - device support
  - device commissioning and management, and device discovery
  - working with network objects and device object configuration on logical groups, shelf and cards objects, and port and channel objects
  - NFM-P device and equipment management functions such as inventory management, TCA, and bulk operations
- NE maintenance—contains information about the functions that facilitate the maintenance of managed NEs, including the following:
  - when to deploy NE configuration changes and how to configure them
  - backing up device configurations on demand, or using a schedule, and restoring device configurations
  - upgrade device software on demand, or using a schedule, and how perform device configuration rollback
  - monitor and troubleshoot operations in progress and how to browse NE file systems
- Network management—contains information about network functions including the following:
  - general routing and forwarding object configuration
  - protocol-specific configuration
  - traffic management using MPLS, MPLS TP, and service tunnels
  - NE redundancy
- Policy management—contains information about configuring and applying NFM-P policies that define rules that govern how network traffic is handled and prioritized
- Service management—contains information about configuring and managing customer services.

- Service assurance—contains information about the proactive detection of service degradation and SLA verification including the following:
  - service verification using the STM and specific OAM diagnostic tests
  - networking troubleshooting activities including fault detection/verificatio/isolation and performance monitoring
  - root-cause analysis audits of services and physical links
  - service throughput configuration

Safety information

For your safety, this document contains safety statements. Safety statements are given at points where risks of damage to personnel, equipment, and operation may exist. Failure to follow the directions in a safety statement may result in serious consequences.

Document support

Customer documentation and product support URLs:

- Documentation Center
- Technical support

How to comment

Please send your feedback to documentation.feedback@nokia.com.

Part I: Getting started

Overview

Purpose

This part provides information about the NFM-P user interface, workspaces, navigation tree, map, and schedules.

Contents

1 NFM-P GUI

GUI overview

1.1 GUI workspace elements

1.1.1 Overview

The GUI on an NFM-P client station allows you to use a standard keyboard and point-and-click operations to perform device, network, policy and service management functions. The GUI simplifies operations such as service configuration, OAM, security, statistics, and other functions by allowing point-and-click selection and intuitive data entry, rather than using a CLI. Multiple NFM-P GUI clients can connect to an NFM-P server.

The main GUI window displays standard workspace elements such as a menu bar, toolbar, taskbar, and status bar.

1.1.2 Navigation tree

The navigation tree lists the managed equipment, services, protocols, and other objects that are configured on the network. By expanding the tree, you can navigate to and select network objects for configuration or to view information. You can open up to five navigation tree windows at one time.

See Chapter 3, “NFM-P navigation tree” for more information about the navigation tree.

1.1.3 Topology map

The topology map provides a view of network elements and connections. You can arrange elements on the map and create equipment groups to suit your requirements. You can select map objects and use contextual menus for configuration or to view information. You can open multiple topology map windows at one time.

See Chapter 4, “Topology map management” for more information about the topology map.

1.1.4 Windows and forms

The NFM-P displays forms in response to menu options or other on-screen selections. Forms are typically used to configure parameters and to view information. Forms allow configurations, display lists and searches, and provide other options; see “NFM-P forms” (p. 82) in this chapter.

All currently open windows and forms, except external windows, are represented by icons on the NFM-P taskbar. All currently open windows and forms, including external windows, are listed in the Window menu of the NFM-P main menu.

The NFM-P GUI provides floating windows and forms that you can move, resize, hide, and bring to the foreground to optimize the workspace. For information about managing the display of windows and forms, see 1.13 “To manage the display of windows and forms” (p. 105). A window or form in the GUI can be moved out of the main workspace and managed as an external window on a second monitor; see 1.15 “To manage a window or form as an external window” (p. 107).

Messages, warnings, dialogs, and other information appear in pop-up windows in the GUI, but are not managed as floating windows.

1.1.5 Main menu bar

The main menu bar contains menus and submenus that allow you to navigate to the forms required to perform many NFM-P functions. The available menu items vary depending on licensing and the scope of command role of the current user.

1.1.6 Toolbar

The toolbar contains icons that allow you to quickly access windows and forms. The icons represent menus or menu items in the NFM-P main menu. When a button represents a menu, the button displays the name of the menu. Clicking the button opens a drop-down of the associated menu items. When an icon represents a menu item with no submenu items, clicking the icon opens the associated form or window. You can set user preferences to show or hide the toolbar; see 1.23 “To configure NFM-P user preferences” (p. 116). You can customize the toolbar as part of a custom workspace; see 2.11 “To customize toolbars” (p. 150).

1.1.7 NFM-P objects

The term “object” in the NFM-P typically refers to an entity on the network, such as an NE, card, port, routing instance, or any other physical or logical object that is displayed in the navigation tree or map. The term “object” can also refer to entities in the NFM-P, such as policies, tests, or other objects that are displayed in lists or on configuration forms.

1.1.8 Contextual menus

Contextual menus appear when you right-click certain objects or areas in the GUI; for example, in the navigation tree and topology map. Contextual menus allow you to perform actions that are specific to the object or area.

1.1.9 Tooltips

The NFM-P provides tooltips to identify many icons, buttons, fields and other GUI elements. When you mouse over these elements, information is displayed.

1.2 GUI customization

1.2.1 Custom workspaces

The NFM-P supports custom workspaces that allow you to configure and arrange workspace elements to suit your operational requirements. You can save and select custom workspaces and update saved workspaces after an NFM-P upgrade. You must have the required scope of command permissions to customize workspaces. Permissions for workspaces are set by a system administrator. See Chapter 2, “NFM-P custom workspaces” for more information about custom workspaces.

1.2.2 Span of control

The GUI displays objects and allows object configuration based on the span of control of the user and user preferences. By default, the GUI displays the objects that are included in the View Access and Edit Access spans of the current user. To reduce the number of objects displayed, you can configure user preferences to show only the objects in the Edit Access span of control of the user; see 1.23 “To configure NFM-P user preferences” (p. 116) . The user preferences setting affects the display of objects in the GUI. You can temporarily override the user preferences setting when you configure filters to search lists; see 1.39 “To filter using span of control” (p. 136) .

For more information about span of control, see the section on creating a span of control in the *NSP System Administrator Guide*.

1.2.3 User preferences

The NFM-P allows you to configure workspace elements and functions to suit your operational requirements. You can set many preferences using the User Preferences form accessed from the main menu. User preferences are specific to a particular user; the settings are associated with the user ID and are applied to the system when that user is logged in. See 1.23 “To configure NFM-P user preferences” (p. 116) for more information about user preferences.

You can set tab preferences to show or hide tabs on forms; see 1.24 “To set local tab preferences for configuration forms” (p. 118) . You can also configure custom workspaces; see Chapter 2, “NFM-P custom workspaces”.

1.2.4 Tab preferences

Some NFM-P property forms contain numerous tabs, which you can display, hide, and arrange in a preferred sequence according to your operational requirements. On some forms, tabs are hidden by default. You can temporarily display hidden tabs, and you can configure and save your tab display preferences; see 1.24 “To set local tab preferences for configuration forms” (p. 118) and 1.25 “To temporarily display hidden tabs on property forms” (p. 119) .

You can configure a workspace to use local tab preferences or custom tab preferences. Local tab preferences are configured using the tab selector and are saved for a specific user. When your current workspace uses local tab preferences, saved local tab preferences become the default that is displayed when a form opens. Custom tab preferences are saved as part of the workspace, and are not user-specific. See 2.9 “To configure tab preferences” (p. 147) .

Local tab preferences can be exported and shared; see 1.26 “To export local tab preferences” (p. 120) and 1.27 “To import local tab preferences” (p. 121) .

An administrator can configure system preferences to enable or disable tab preferences and to define the default local behavior for hidden tabs. See the section on system preferences configuration procedures in the *NSP System Administrator Guide*.

1.2.5 Localized language support

The GUI supports localized language display. Localized language display, also known as internationalization, displays GUI text in a specified language. The localized language setting applies to most GUI objects, except system components and database objects. Contact Nokia technical support for more information about localized language support.

Note: The NFM-P supports localized language settings using predefined strings, and does not translate data to different languages.

The XML API does not support language localization.

1.3 Additional GUI operations

1.3.1 Shortcut keys and accessibility

The NFM-P supports shortcut keys to enable operation by keyboard alone, without a mouse or pointer. The following keyboard shortcuts are available:

- Standard Java-based shortcuts. The NFM-P supports many conventional keyboard shortcuts for accessibility and convenience.
- Mnemonics. You can use mnemonics to open menu items and access forms. To open a menu item, simultaneously press Alt-*key*, where *key* is the underlined letter of the menu item. For example, when the main menu is displayed, pressing Alt-P opens the Policies menu. The menu item must be displayed for the mnemonic to work.
- Keyboard shortcuts for specific NFM-P functions.

The following table describes some of the keyboard shortcuts for specific functions in the NFM-P. For multiple keystrokes, a hyphen means that the keys are pressed simultaneously.

Keyboard shortcut	Description
Ctrl-A	Selects all GUI client sessions on the Select Sessions form Selects all objects on the topology map
Ctrl-C	Stops a command, when using the command window
Ctrl-F	Opens the “Find Attribute in” form when a property form is the active GUI element Opens or closes the search panel when the navigation tree is the active GUI element
Ctrl-O	Opens a directory and lists its contents, when a directory is selected in an NE file system
Ctrl-T	Opens the tab selector dropdown on a property form
Ctrl-Shift-A	Temporarily displays all hidden tabs on a property form
F3 and Shift-F3	When search results for the navigation tree produce multiple matching objects, the F3 key selects the next object. Shift-F3 returns to the previous matching object.

1.3.2 NFM-P clipboard

As well as the standard clipboard functions for copying text, the NFM-P provides a clipboard function that copies property form identifiers. A property form identifier is a unique internal address that the NFM-P assigns to a property form.

Identifiers are copied to the clipboard using the Clipboard icon or the Copy to Clipboard option in a contextual menu. You can use the clipboard to open forms, send identifiers to other users, and configure search filters.

For more information about the NFM-P clipboard, see 1.20 “To use the NFM-P clipboard” (p. 112).

1.3.3 Broadcast messages

You can send messages to other selected NFM-P users or to all active users logged into the NFM-P. This is useful for sending maintenance and similar notifications to other users.

When you receive a message, you can click Reply to respond to the sender.

Messages are uniquely identified with the Client ID number of the sender.

See the procedures on sending a text message to other NFM-P users and sending a broadcast message to GUI clients in the *NSP System Administrator Guide* for more information.

1.3.4 Task manager

The task manager allows you to monitor the progress of operational tasks. The task manager monitors the following operations:

- all write operations that are performed from the GUI; for example, when you click Apply or OK
- all write operations that are performed using the XML API
- some read operations; for example, when you click Resync or Collect All

See 1.21 “To monitor the NFM-P Task Manager” (p. 114) for more information.

An NFM-P administrator can change the task manager settings. See the procedure to change the NFM-P Task Manager settings in the *NSP System Administrator Guide*.

You can also send the tasks displayed in the Task Manager to a file using the findToFile method. See “Inventory retrieval methods” in the *NSP NFM-P XML API Developer Guide* for more information.

NFM-P forms

1.4 Forms overview

1.4.1 Types of forms

The NFM-P displays forms in response to menu options or other on-screen selections. Forms display lists or allow you to view and configure properties. Some forms display other information.

The GUI displays the following form types:

- list forms
- configuration forms
- step forms
- property forms
- problems encountered forms
- other forms such as messages, warnings, or dialog boxes. These forms are also sometimes referred to as windows.

See Figure 1-1, “List form” (p. 83) and Figure 1-4, “Service configuration form with built-in tree” (p. 88) for examples of NFM-P forms.

Forms can be sized, positioned, and managed using standard OS functions and NFM-P display features; see 1.13 “To manage the display of windows and forms” (p. 105) . A form can be displayed anywhere in the GUI. A newly opened form is displayed in the foreground. You can do the following:

- organize forms according to your preferences
- compare information on multiple open forms
- navigate quickly to another open form
- save a set of open forms and reopen them later with one operation, to provide quick access to forms that you use often; see 1.14 “To save or open a set of forms for quick access” (p. 106) .

Forms are identified by a titlebar at the top. The displayed form name is the name specified during object creation. If the object is not named, a default name is used. When a form is minimized to the taskbar, a tool tip on the taskbar icon of the form displays the title bar information. The Window menu on the main menu lists all open forms.

Forms typically contain tabs, buttons, fields, and other elements that you can click to perform actions, depending on the form type. When these elements are not available or not applicable, they are dimmed.

1.4.2 More Actions button

Buttons that perform various functions appear at the bottom of configuration forms and on the right side of list forms. When a form is resized to a smaller size, or when there are too many buttons to fit in the available space, some of the buttons are consolidated into a More Actions button. The consolidated buttons are available as menu items when you click More Actions; see Figure 1-1, “List form” (p. 83) and Figure 1-4, “Service configuration form with built-in tree” (p. 88) .

1.5 List forms

1.5.1 General information

A list form displays a list of objects in response to a menu option, tab, or other on-screen selection. Figure 1-1, “List form” (p. 82) shows the features of a list form.

Some list forms populate automatically; for others, you must click Search to see results.

Figure 1-1 List form

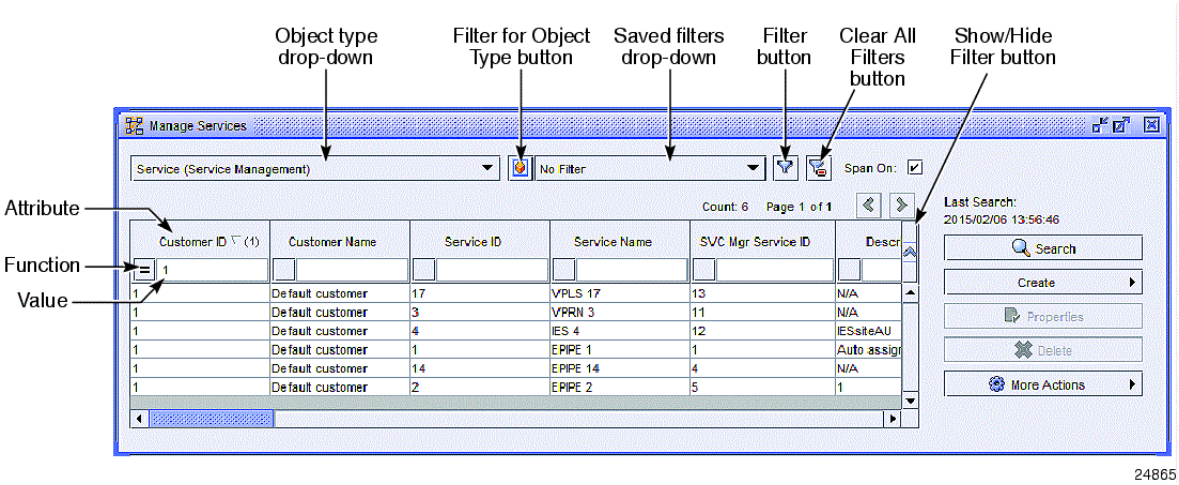

1.5.2 Using list forms

Some list forms contain an object type drop-down that allows you to select the type of objects in the list. When an object type drop-down menu contains more than ten items, you can apply a filter to reduce the number of drop-down items; see 1.38 “To filter object types” (p. 136) . You can customize the object drop-down and select the default object type in a custom workspace; see 2.13 “To customize list forms” (p. 153) .

Objects on a list form are displayed in table format. The column headings correspond to the attributes (parameters) for the object, for example, the Customer ID for a service. The column headings vary depending on the object type. Each object in the list is displayed in a row, showing the values for the attributes of the object. On some forms, the first column on the left displays a deployment icon for an item, when required; see 11.1.2 “Object deployment status” (p. 334) in Chapter 11, “Working with network objects” .

Objects are displayed based on the span of control of a user and user preferences. See 1.2.2 “Span of control” (p. 79) in this chapter.

List forms can contain many objects, in some cases hundreds or thousands. You can reduce the number of items in a list using search filters; see “NFM-P searches” (p. 92) in this chapter. Most list forms display the time at which the last search was performed on the form.

List forms typically display a column of buttons along the right side. On some list forms, not all buttons are displayed, but you can access them using the More Actions button. See 1.4.2 “More Actions button” (p. 82) in this section.

Many list forms contain a Customize button that allows you to quickly make changes to the list form default settings in a custom workspace. See 2.13 “To customize list forms” (p. 153) for more information.

The NFM-P provides features that allow you to manage the display of listed results and to set and save your display configuration. See 1.8.4 “Managing search results and lists” (p. 97) in this chapter, and 1.29 “To manage the display of listed information” (p. 122) . You can also save the results to a file; see 1.22 “To save listed information to a file” (p. 115) .

You can set the following user preferences for list forms:

- number of items per page
- object display based on user span of control
- default file extension for saving listed information

See 1.23 “To configure NFM-P user preferences” (p. 116) for more information about user preferences.

1.6 Configuration forms

1.6.1 General information

Configuration forms allow you to create and modify NFM-P objects and to view information about objects; see 1.17 “To use configuration forms to configure or view parameters” (p. 109) . Figure 1-4, “Service configuration form with built-in tree” (p. 88) describes the main features of a configuration form in the NFM-P.

When you first create an object in the NFM-P, you use a configuration form to enter values for the parameters that define the object. For example, when you create a service, you use a service configuration form; when you create a policy, you use a policy configuration form. After object creation, the form associated with the object is typically called a property form. See 1.6.12 “Property forms” (p. 89) in this section.

You can set user preferences to suppress warnings and messages for configuration forms; see 1.23 “To configure NFM-P user preferences” (p. 116) .

1.6.2 Step forms

Some configuration activities lead the operator through a series of forms, each of which represents a step in the configuration process. Such a form is called a step form. You must click Next to proceed to the next step. The following figure shows the first step in a step form sequence. When the configuration sequence is complete, you must click Finish to commit the changes.

Figure 1-2 Step form

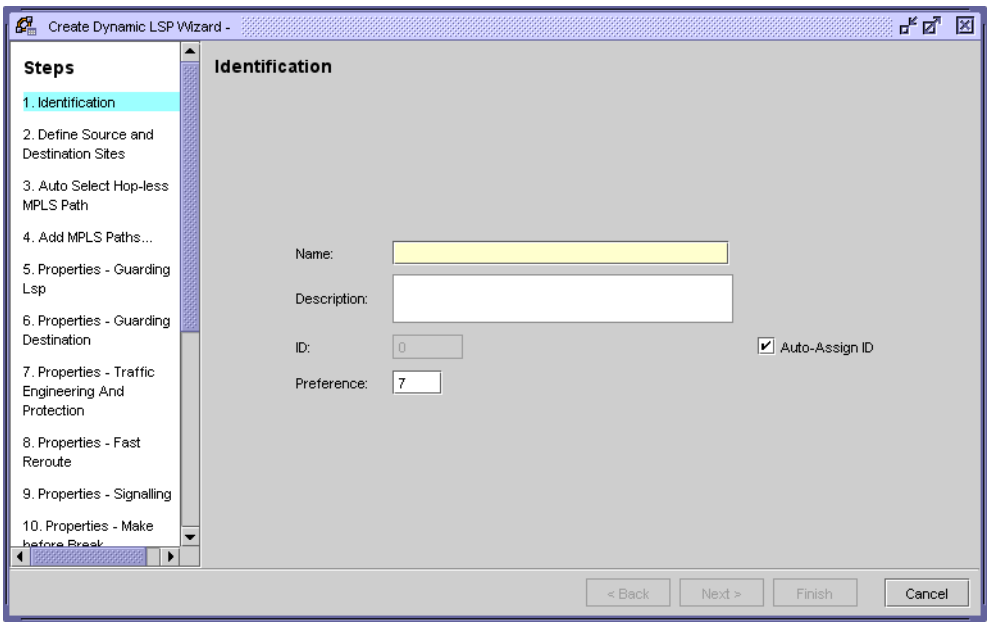

Some steps open a new step form. You must complete the steps in the new form before you can return to the previous form. After you click Finish on the new form, the previous form reappears.

1.6.3 Tabs

Configuration forms typically display tabs. You can use the tabs to open related forms for additional configurations or to open lists or view other information, such as deployment or fault status. Property forms allow you to show or hide tabs to suit your requirements, and to set preferences for tab display; see 1.2.4 “Tab preferences” (p. 79).

1.6.4 Collapsible panels

Configuration forms are often subdivided into specific information areas using panels. Panels are framed by border lines, with a title bar at the top that identifies the panel. Panels can be collapsed to hide areas that are not of interest, or expanded to display all information on a tab or form; see 1.17 “To use configuration forms to configure or view parameters” (p. 109).

When a configuration form closes, the current state of the panel is saved for that object type. When a user opens a configuration form, the last saved state of the panel is displayed.

1.6.5 Parameters

When you create or modify an object, you typically configure parameters, sometimes referred to as attributes. Parameter settings define the properties of an object. Some parameters are read-only, and cannot be configured. Configurable parameters typically display a white field, check box, or drop down arrow. Mandatory fields are yellow. Read-only parameter values display in a grey field.

Non-applicable parameters are dimmed. See 1.17 “To use configuration forms to configure or view parameters” (p. 109) for more information about parameters.

1.6.6 Saving configurations

When you complete a configuration, you can click OK, Apply, or other buttons to save the configured values. When you click the OK button, the information is saved and the form closes. When you click the Apply button, the information is saved and the form remains open to allow you to perform additional actions. See 1.17 “To use configuration forms to configure or view parameters” (p. 109) for more information.

1.6.7 Warning and confirmation messages

When you save configuration changes or close a form, the NFM-P typically displays a warning or confirmation message. You must acknowledge the message. Not all configurations require confirmation.

When you save changes for a child object configuration form that was launched from a parent object form, the NFM-P displays a warning that the changes are not committed until the parent object form (called the containing window) is also saved. You must acknowledge the warning. You can set user preferences to suppress containing window warnings; see 1.23 “To configure NFM-P user preferences” (p. 116). If you then attempt to close the parent form without saving it first, the NFM-P displays a warning, regardless of the user preferences setting.

1.6.8 Indicators and icons

Indicators and icons inform you of activity that is in progress or requires attention, and can appear and disappear depending on the activity that is occurring in the NFM-P, as shown in the following figure.

Figure 1-3 Indicators and icons

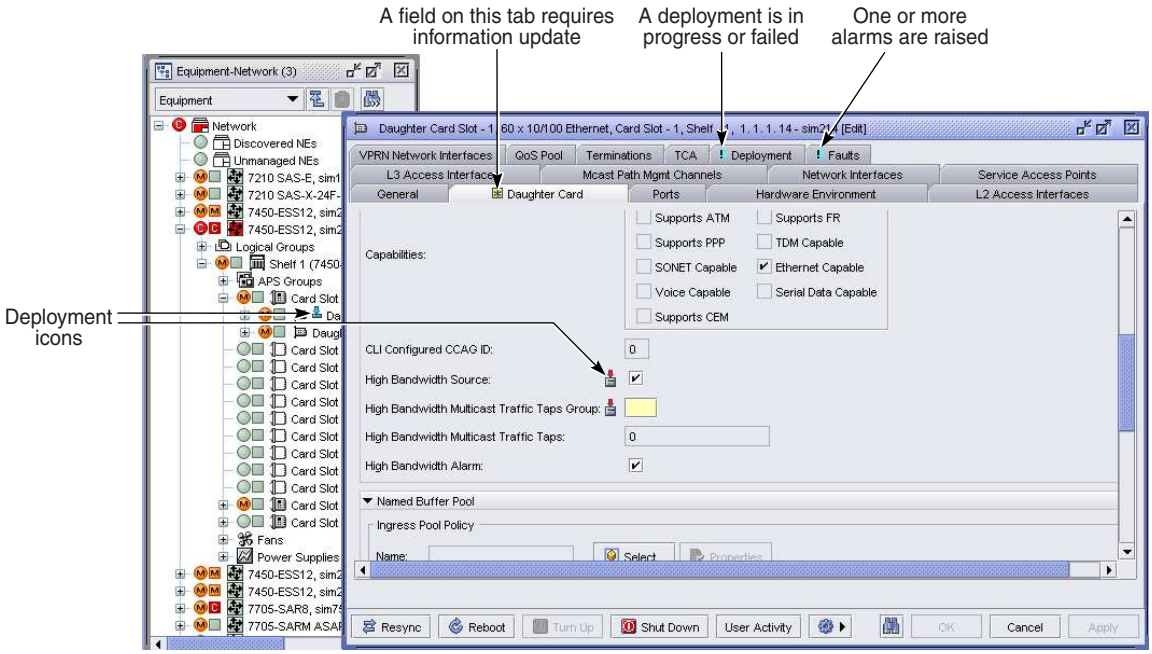

22870

A yellow asterisk icon on a tab or panel title bar indicates that a field contains incorrect data, or that a mandatory field requires data.

A warning indicator appears on the Deployment tab when a configuration change is not fully deployed to an NE.

When a deployment is in progress or has failed, a deployment icon appears beside parameters on an object properties form, and beside affected objects in the navigation tree, and in list forms. See 11.1.2 “Object deployment status” (p. 334) in Chapter 11, “Working with network objects” for more information.

1.6.9 Action buttons

Configuration forms typically have a row of action buttons along the bottom. The available buttons vary depending on the form. The Find icon is available on all configuration forms, and allows you to quickly locate a specific tab, panel, or parameter; see 1.37 “To locate an attribute on a configuration form” (p. 135). On some configuration forms, not all buttons are displayed, but you can access them using the More Actions button. See 1.4.2 “More Actions button” (p. 82) in this section.

1.6.10 Configuration forms with built-in navigation trees

Services and NEs typically require configuration at several levels of the service or equipment object tree. For example, when you create a service, you configure parameters at the service, site, and interface levels. NEs require configuration at the device, shelf, card, and port levels; see Chapter 11, “Working with network objects”.

NFM-P service configuration and NE property forms display a navigation tree on the left side of the form, which provides quick access to service and NE objects. Other object types also display navigation trees in configuration or property forms. See “NFM-P navigation tree” (p. 161) for information about using navigation trees. Objects in the tree provide access to configuration forms using contextual menus.

Service navigation trees for VPLS, I-VPLS, B-VPLS, MVPLS, IES, and VPRNs display a maximum of 50 access interfaces, SAPs, MSAPs, subscriber SAPs, and spoke and mesh SDP bindings at a time. When this limit is reached, you can click on the message beside the containing object under which the SAPs or access interfaces are listed to open a tab with the full list of objects.

Figure 1-4 Service configuration form with built-in tree

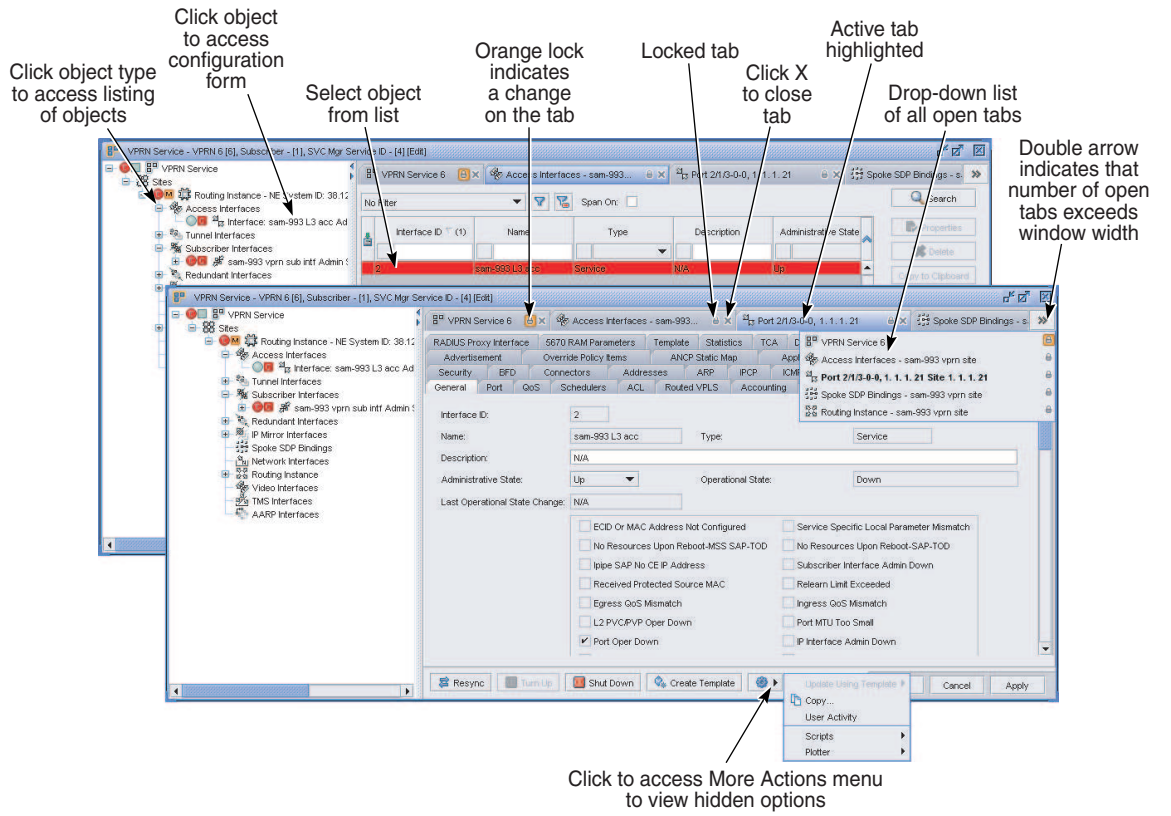

21805

When you first open a configuration form with a built-in navigation tree, the form is unlocked. When you select another object in the tree, a form for the selected object replaces the unlocked form of the previous object. You can lock forms to keep them open when you select additional objects. The form for each object is displayed as a tab in the main configuration form.

You can undock tabs to detach them from the main configuration form. Undocked tabs are displayed as separate forms and are managed like other standard form in the NFM-P. You cannot undock Create forms, tabs for object type lists, or tabs with unsaved changes.

See 1.18 “To manage configuration forms with built-in navigation trees” (p. 110) for more information about managing configuration forms with built-in navigation trees.

1.6.11 Modifying configuration forms with built-in navigation trees

When you modify a docked configuration form, it becomes locked. The lock icon appears in orange until the changes are saved. An orange lock icon indicates unsaved changes on the form. You cannot undock tabs with orange lock icons. When you save the changes, the form remains locked but the icon is no longer orange.

If you attempt to close a modified configuration form that contains unsaved changes, a message prompts you to specify whether to discard the changes or return to the display. If you attempt to close a modified configuration form that has related forms with unsaved modifications, the message lists the forms that have unsaved changes.

For service configuration forms, if an object is added to the service configuration, the new object appears in the built-in navigation tree and the configuration form is displayed with the orange lock icon. The icon remains orange until the modifications are saved or discarded.

For NE property forms, if an object is added from the navigation tree, the change will be made immediately. However, if a change, such as adding or removing an object, is made to an equipment configuration form, the change will be reflected on the form immediately, but it will not be reflected on the navigation tree until the OK, Apply, or Apply Tab button is clicked.

Note: Changes implemented by pressing the OK or Apply buttons affect all tabs in the form. For some configuration forms, you can click the Apply Tab button to save changes for the current object-level tab and related objects. A message lists the affected objects.

If you save a set of forms for quick access when child forms are displayed as tabs in a parent form, all of the child tabs are saved as separate forms. When you reopen the forms, the child tabs are displayed as separate forms. See 1.14 “To save or open a set of forms for quick access” (p. 106) .

1.6.12 Property forms

Configuration forms for objects that already exist in the NFM-P (NEs for example) are typically called property forms. You can access property forms by using contextual menus, the Properties button on list forms, or by double-clicking on objects in the navigation tree that have no child objects.

A property form has a property form identifier that you can copy to the clipboard and paste for various functions; see 1.20 “To use the NFM-P clipboard” (p. 112).

Most property forms contain tabs that provide specific information. The Deployment tab allows you to monitor the deployment status of an object; see 11.1.2 “Object deployment status” (p. 334) in Chapter 11, “Working with network objects”.

1.6.13 NE sessions

You can open a CLI window using the Telnet Session or SSH Session button on NE property forms. You can access the NE file system using the File Browser button. These options are also available using the contextual menu for NE objects in the navigation tree.

1.6.14 User Activity button

The User Activity button opens a form that lists the recent NFM-P user actions performed on the object. See the section on user activity logging in the *NSP System Administrator Guide* for more information.

1.6.15 Resync button

The Resync button on a property form resynchronizes the data in the NFM-P with the current state of the corresponding object. The NFM-P requests the configuration from the object and updates the NFM-P network model accordingly. Resynchronization does not affect the contents of the historical statistics database.

1.6.16 Turn Up and Shut Down buttons

The Turn Up and Shut Down buttons on some property forms provide a convenient method for changing the administrative state of an object. When you click the Turn Up or Shut Down buttons, the change is effected immediately. When you modify the Administrative State parameter, the change is not effected until you click the OK or Apply button for the form.

1.6.17 Multi-edit property forms

The NFM-P provides a multi-edit function that allows you to select multiple objects at one time for editing of specific properties. The objects must be of the same type (for example, services must all be VPRN, not a mixture of VPRN and other service types).

Changes for the selected objects are entered on a Multiple Instances (Edit) property form. When the objects selected for multi-edit are configured with differing parameter values, the NFM-P typically alerts you by displaying multi-value fields or drop-downs for the parameters whose values differ. Multi-value fields and drop-downs are identified by triple borders along the top and right side; see Figure 1-5, “Multi-value property form” (p. 91). Both configurable and read-only parameters may display as multi-value properties.

You can modify configurable multi-value fields. When a value is entered in a multi-value field, the value is propagated to all selected objects in the multi-edit group, and the parameter no longer displays as a multi-value property. See 1.19 “To modify multiple objects at one time (multi-edit)” (p. 111) for more information.

The NFM-P also provides a bulk operations function; see Chapter 20, “Bulk operations”.

Figure 1-5 Multi-value property form

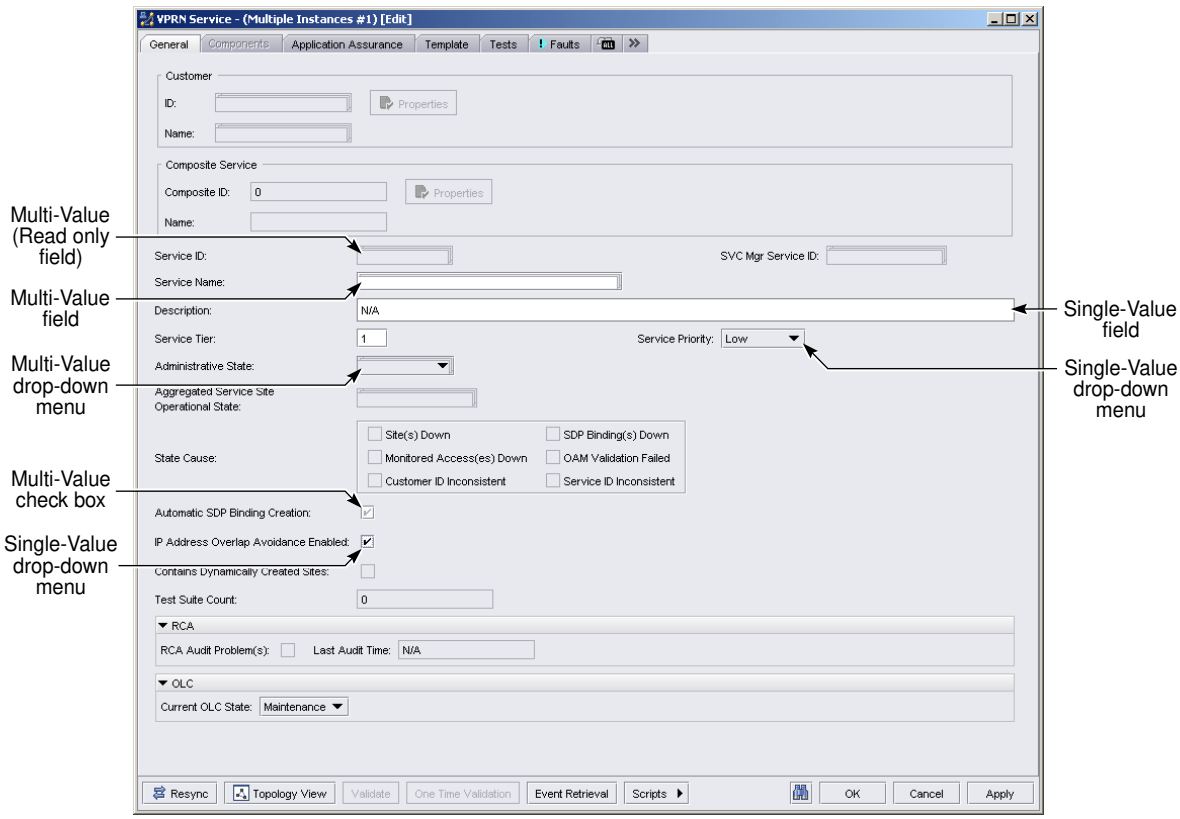

23740

1.6.18 Problems encountered notifications

When you enter a faulty configuration or instruction in the NFM-P, a Problems Encountered form may open. The Problems Encountered form displays problems in a list. You can double-click on a problem entry, or select and click Properties to display a property form with more information about the problem. The Problems Encountered form contains a Deployment button that you can click to access the Deployment form for additional information and actions; see 25.5 “To view and manage failed deployments” (p. 769) for more information. You can copy the property form identifier for the Problems Encountered form to the NFM-P clipboard; see 1.20 “To use the NFM-P clipboard” (p. 112).

NFM-P searches

1.7 Search overview

1.7.1 Search types

The NFM-P typically manages large numbers of network objects, policies, and configurations. Searches allow you to find the objects and information required for your operations. Searches are available for list forms, navigation tree objects, configuration parameters, and online documentation.

This section provides an overview of NFM-P search features. See "Procedures for searching" (p. 127) for search procedures.

Table 1-1 NFM-P search types

Search type	Description	See
Find attribute	Locate a tab, panel, or parameter on a form using the Find icon	1.37 "To locate an attribute on a configuration form" (p. 135)
Simple search	Search lists using column headings to create filters	1.8.1 "Simple searches" (p. 95) in this section 1.30 "To perform a simple search from an object list form" (p. 127)
Advanced search	Search lists using the filter configuration form	1.8.2 "Advanced searches" (p. 95) in this section 1.31 "To perform an advanced search from an object list form" (p. 128)
Endpoint search	Search for LSPs, tunnels, and paths by specifying endpoints	1.32 "To perform a search by specifying endpoints" (p. 131)
Equipment Group filters	Search for equipment objects, and service objects, based on equipment groups	1.40 "To configure and save equipment group filters" (p. 138) and 1.41 "To use a saved equipment group filter" (p. 139)
Navigation tree	Search for network objects in the navigation tree using the Find icon or Find field	3.7 "To locate objects in the navigation tree" (p. 165)
NSP Help Center	Search the NFM-P user documentation	Help→Open Help Center in the NFM-P main menu.

1.7.2 Search filters

Some search types use configurable search filters. Filter configuration varies depending on the type of search; see Table 1-1, "NFM-P search types" (p. 92).

Search filters for lists in the NFM-P typically consist of one or more filter properties that are based on the type of objects in the list.

A filter property is a combination of an attribute, a function, and a value:

- Attribute: defines the information type, or field, to be searched. Attributes typically correspond to

the properties and parameters for objects in the list. For example, you could filter a list of services by using the Customer attribute to search for services associated with a particular customer. Attributes are specific to the object type being searched.

- Function: defines the operation that the search performs. For example, the CONTAINS function will search in the list for instances where the attribute contains a specified value, such as a number or text string. The available functions vary, depending on the selected attribute. See Table 1-2, “Search functions” (p. 93) for information about search functions.
- Value: specifies the character string or parameter setting used by the function. The available values depend on the selected attribute. For example, if the Customer ID attribute is selected, you can configure numerals in the Value field. If the Current OLC State attribute is selected, you can choose In Service or Maintenance from the drop-down menu for the Value field.

You can configure filter properties using the column headings of a list; see 1.8.1 “Simple searches” (p. 95) in this section. You can open the filter configuration form to display, create, and modify filter properties. The filter configuration form is used for advanced searches; see 1.8.2 “Advanced searches” (p. 95) in this section. Figure 1-6, “Filter configuration form” (p. 96) shows the filter configuration form.

You can save search filters for reuse or modification; see 1.33 “To save search filters” (p. 132) . Saved filters appear in the filter list drop-down menu for the list form; see Figure 1-1, “List form” (p. 83) .

When a saved search filter contains attributes or values that are not available for the object type you are searching, the NFM-P alerts you by displaying red text. See 1.8.3 “Invalid attributes or values” (p. 96) in this section for more information.

If the number of search results exceeds the allowed limit, the NFM-P displays a dialog box that indicates the number of results returned. You can modify a filter to refine the search criteria and reduce the search results.

On some list forms, you can enable a span of control filter to display only the objects that are within your Edit span of control; see 1.39 “To filter using span of control” (p. 136) .

Table 1-2 Search functions

Operator	Description
APPROXIMATELY EQUAL	Filters by object timestamp. The function allows you to specify a timestamp value using 1 min resolution rather than the 1 ms default resolution. The search returns items that match the specified timestamp value. Select a time unit in the display (yyyy/mm/dd/hh:mm) and click on the arrows in the column heading to specify a time value.
BETWEEN	Filters by object timestamp. The function allows you to specify a start time and end time. The search returns items that have a timestamp between the configured values. Click the Select icon, select a time unit in the display (yyyy/mm/dd/hh:mm), and click on the arrows in the column heading to specify a time value for Start Time and End Time. If the start and end time are the same, the OK button is dimmed.

Table 1-2 Search functions (continued)

Operator	Description
CONTAINS	Filters by comparing character strings. The search returns an item when the search term matches any part of the character string of the item. For example, the filter property “Description CONTAINS (ab)” returns all items with the string “ab” in the Description attribute, such as “ABC Industries” and “Calgary, AB”. An empty CONTAINS search returns all items. The function is not case-sensitive.
DOES NOT CONTAIN	Filters by comparing character strings. The search returns an item when the search term does not match any part of the character string of the item. For example, the filter property “Description DOES NOT CONTAIN (ab)” returns items that do not contain the string “ab” in the Description attribute. Items with Descriptions such as “ABC Industries” and “Calgary, AB” are not returned in the search results. The function is not case-sensitive.
EQUALS	The search returns an item when the search term exactly matches the attribute value of the item. The function is case-sensitive.
GREATER OR EQUAL	Filters by comparing numerical values.The search returns items that have a value greater than or equal to the search term.
GREATER THAN	Filters by comparing numerical values.The search returns items that have a value greater than the search term.
IS BLANK	The search returns items that do not have a value. For example, the filter property “Associated Template IS BLANK” returns services that have no configured value in the Associated Template field. When the IS BLANK function is selected, you cannot configure a value for the filter property.
IN THE PAST	Filters by object timestamp. The function allows you to enter a period of time. The search returns items that are within the specified period of time before the present; for example, in the past two hours. Click the Select icon in the column heading and enter a period of time.
LESS OR EQUAL	Filters by comparing numerical values.The search returns items that have a value less than or equal to the search term.
LESS THAN	Filters by comparing numerical values.The search returns items that have a value less than the search term.
MATCHES	Applies to simple searches using column headings, for a single property of a group of checkboxes. The MATCHES function filters by comparing checkbox settings. The search returns an item when the checkbox setting (selected or deselected) of the search term matches the checkbox setting of the item.
MATCHES ALL	Applies to advanced searches using the filter configuration form, for attributes that show settings for a group of checkboxes. When you choose the MATCHES ALL function, you must click Properties in the Value field and select or deselect the checkboxes associated with the attribute. The MATCHES ALL function filters by comparing checkbox settings. The search returns an item when all of the selected checkbox settings of the search term are selected for the item.

Table 1-2 Search functions (continued)

Operator	Description
MATCHES ANY	Applies to advanced searches using the filter configuration form, for attributes that show settings for a group of checkboxes. When you choose the MATCHES ANY function, you must click Properties in the Value field and select or deselect the checkboxes associated with the attribute. The MATCHES ANY function filters by comparing checkbox settings. The search returns an item when one or more of the selected checkboxes of the search term matches the corresponding checkbox of the item.
NOT EQUAL	The search returns an item when the search term does not exactly match the attribute value of the item. The function is case-sensitive.
WILDCARD	For filters based on character strings (letters, numbers, or special characters). Wildcards substitute for characters in the search term. Use a question mark (?) to substitute for a single character. Use a single asterisk (*) to substitute for multiple characters. The function is case-sensitive.

Equipment group filters

Equipment group filters allow you to search for objects that are organized by equipment group, such as managed equipment objects, and managed service objects. An equipment group is sometimes called a topology group.

An equipment group filter is created by selecting from existing equipment groups. Filters can be saved for reuse or modification. See 1.40 “To configure and save equipment group filters” (p. 138) and 1.41 “To use a saved equipment group filter” (p. 139) for more information about equipment group filters.

1.8 Searching tips

1.8.1 Simple searches

You can perform a simple search on a list form using the column headings of the list table. Each column heading displays an attribute. You can configure a filter property for an attribute by selecting a function and a value from the column heading fields.

Filter properties for simple searches are displayed in the column heading, and show the attribute, the function (as a symbol, such as = or >), and the configured value. See Figure 1-1, “List form” (p. 83) . You can configure multiple filter properties at one time.

See 1.30 “To perform a simple search from an object list form” (p. 127) for more information about simple searches.

1.8.2 Advanced searches

You can perform an advanced search on a list using the filter configuration form.

Figure 1-6 Filter configuration form

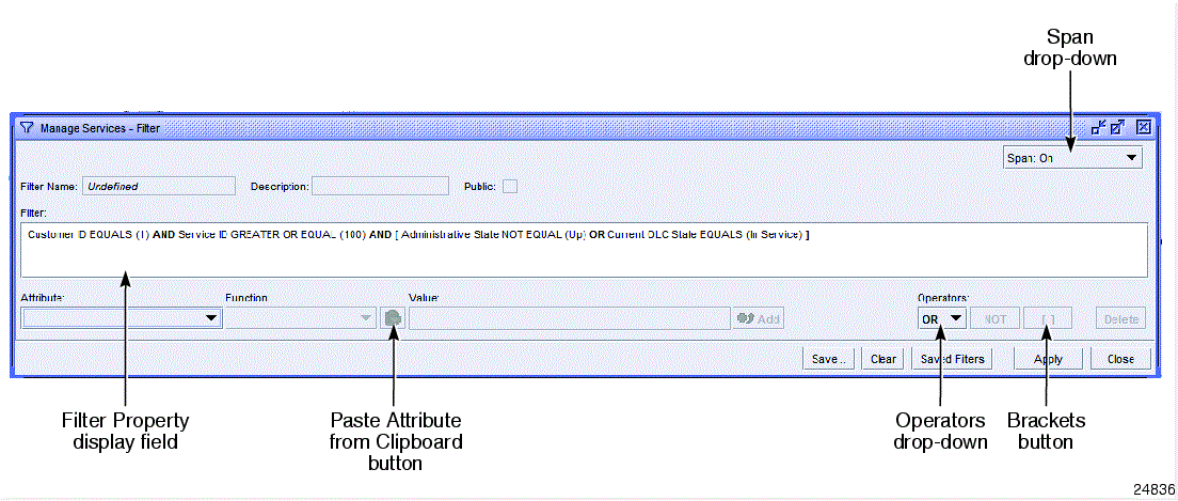

With an advanced search, you can apply more than one function and value to a single attribute. For example, you could configure a filter property Customer ID GREATER THAN (1) and also, in the same filter, the property Customer ID LESS THAN (100). This is not possible with a simple search.

An advanced search filter combines filter properties using Boolean operators. You can group properties together using brackets, and apply Boolean operators to the groups. See Table 1-5, “Boolean search operators” (p. 128) for information about Boolean operators.

You can use the filter configuration form to modify simple searches. See 1.31 “To perform an advanced search from an object list form” (p. 128) for more information about advanced searches and the filter configuration form.

1.8.3 Invalid attributes or values

When a saved search filter contains attributes or values that are not available for the object type you are searching, the NFM-P alerts you by displaying red text. The filter name is displayed in red text in the saved filters drop-down of a list form. The invalid attributes or values are displayed in red text on the filter configuration form.

You can select filters with invalid attributes or values and use them to perform a search. However, the invalid attributes are not evaluated. You can also display and modify such filters using the filter configuration form. For information about using the filter configuration form, see 1.31 “To perform an advanced search from an object list form” (p. 128) .

When a filter property that contains an invalid attribute is displayed on the filter configuration form, the attribute, function, and value for that property are all shown in red text. The attribute is surrounded by asterisks (**). Both the attribute and value are displayed as a best-effort string and may not support language localization.

When you select an invalid attribute in the filter display field, it is displayed in the Attribute drop-down as a best-effort string that may not support language localization. If the attribute is not recognized, the Function drop-down and the Value field display only the configurations that are

shown in the Filter display field. They cannot be changed until a different attribute is chosen from the Attribute drop-down menu.

As long as invalid attributes or values exist, the Save button on the filter configuration form is dimmed.

When only the value in a filter property is invalid, the value is displayed in red text. The value is surrounded by asterisks (**) and may not support language localization. The Value drop-down does not display a value, but is configurable.

1.8.4 Managing search results and lists

Most list forms allow you to manage the display of search results. You can:

- specify the number of items per page and navigate through pages
- remove, resize, and rearrange columns
- sort lists based on a selected column in ascending or descending order
- save your display preferences

See 1.29 “To manage the display of listed information” (p. 122) for more information about managing the display of lists.

You can save listed information to a file; see 1.22 “To save listed information to a file” (p. 115) .

1.8.5 Parameter search

See the XML API Reference for information about parameters and configuration forms in the NFM-P.

Information about NFM-P parameters that are not found in the XML API Reference is available in Appendix A, “Parameters”.

Procedures for opening and closing the GUI

1.9 To open a single-user GUI client configured for one NFM-P system

1.9.1 Purpose

Perform this procedure to open a single-user GUI client that is configured to connect to only one NFM-P system.

By default, the NFM-P GUI client login form lists the main servers in one NFM-P system. If you manage multiple NFM-P systems at the same release, you can configure the login form to list the main servers in each system as login options; see the *NSP Installation and Upgrade Guide*.

Note: To perform the procedure, you must be logged in to the single-user client station as one of the following:

- the user who installed the client software
- a user with sufficient permissions on the client files and directories, such as a local administrator

1.9.2 Steps

1

Double-click on the NSP NFM-P Client (*server*) icon on the RHEL or Windows desktop, where *server* is a main server IP address or hostname.

If you are currently logged in to the NSP and have suppressed security messages, the NFM-P GUI opens.

2

If the form in Figure 1-7, “Do you want to run this application?” (p. 98) is displayed, perform the following steps.

1. To suppress the form when opening subsequent GUI sessions, select the check box.
2. Click Run.

If you are currently logged in to the NSP, the NFM-P GUI opens.

Figure 1-7 Do you want to run this application?

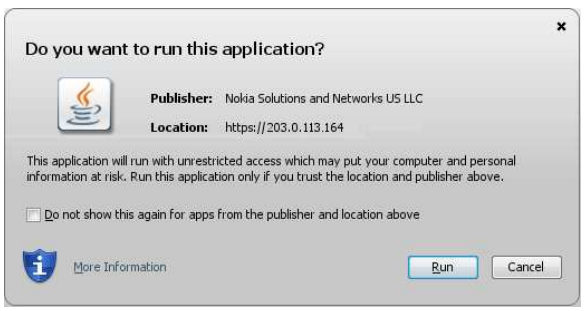

3

If you are not currently logged in to the NSP:

- The splash screen in Figure 1-8, “Network Functions Manager - Packet client” (p. 98) is displayed.
- The default browser opens to the NSP sign-in page.

Enter your NSP user credentials, acknowledge the security statement, if present, and click SIGN IN.

Note: If you do not enter valid credentials within the allowed authentication period, the splash screen indicates that the client has timed out; you must close the splash screen and return to Step 1.

The client splash screen indicates that the client is loading, and then closes as the NFM-P GUI opens.

Figure 1-8 Network Functions Manager - Packet client

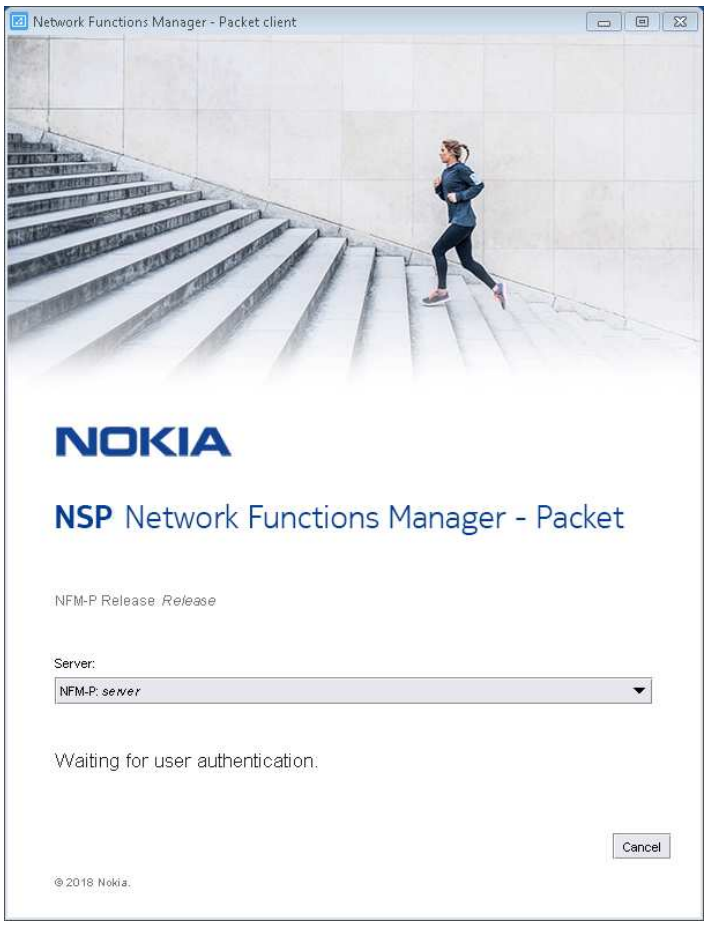

END OF STEPS

1.10 To open a single-user GUI client configured for multiple NFM-P systems

1.10.1 Purpose

An NFM-P GUI client can be configured to connect to multiple NFM-P systems at the same release. After you configure a GUI client login form to list multiple NFM-P systems, as described in the *NSP Installation and Upgrade Guide*, perform this procedure to open a GUI client session on a specific NFM-P system.

Note: To perform this procedure, you must be logged in to the single-user client station as one of the following:

- the user who installed the client software
- a user with sufficient permissions on the client files and directories, such as a local administrator

Note: In deployments in which one NFM-P GUI client is installed on a client station and configured to point to multiple independent standalone NFM-P servers, if one of the servers is upgraded, duplicate NFM-P entries appear in the GUI for client upgrade. To work around this issue, remove the extraneous entries from the nms-client.xml file. See "To configure a GUI client login form to list multiple NFM-P systems" in the *NSP Installation and Upgrade Guide* for information about configuring this file.

1.10.2 Steps

1

Double-click on the NSP NFM-P Client (*server*) icon on the RHEL or Windows desktop, where *server* is a main server IP address or hostname.

2

The form in Figure 1-9, “Network Functions Manager - Packet client” (p. 101) is displayed.

Choose a server from the drop-down list and click Sign In.

If you are currently logged in to the NSP and have suppressed security messages, the NFM-P GUI opens.

Figure 1-9 Network Functions Manager - Packet client

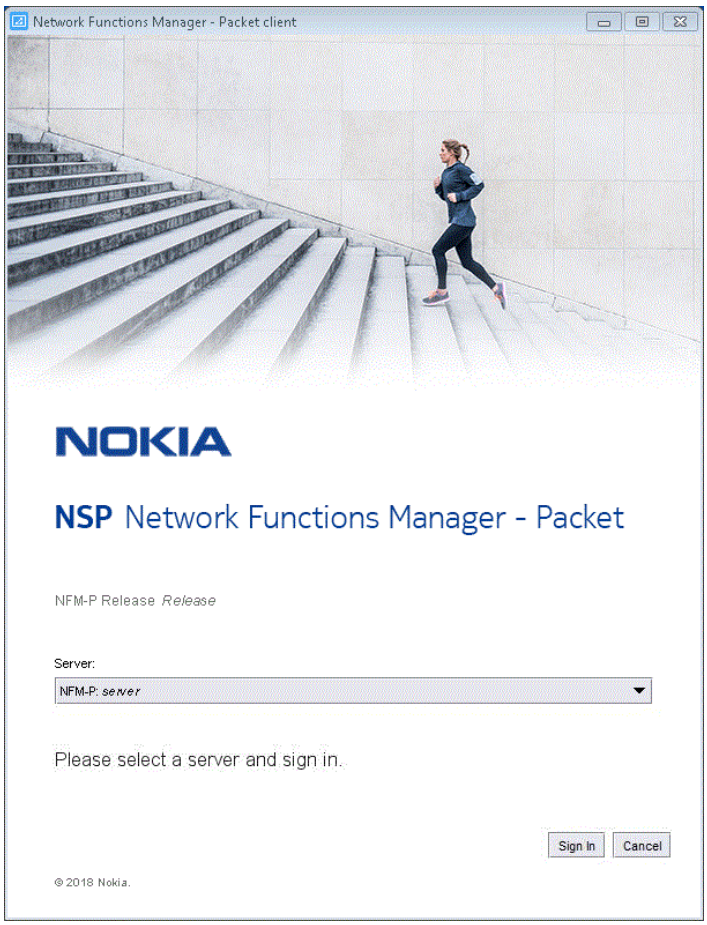

3

If the form in Figure 1-10, “Do you want to run this application?” (p. 102) is displayed, perform the following steps.

1. To suppress the form when opening subsequent GUI sessions, select the check box.
2. Click Run.

If you are currently logged in to the NSP, the NFM-P GUI opens.

Figure 1-10 Do you want to run this application?

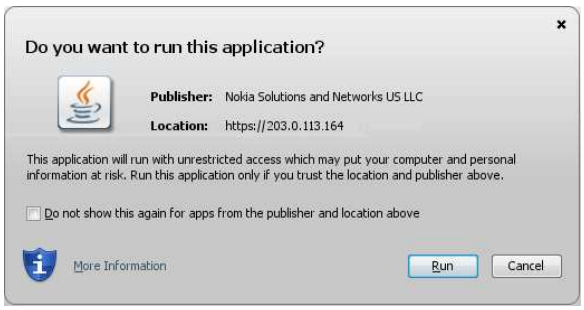

4

If you are not currently logged in to the NSP:

- The splash screen in Figure 1-11, “Network Functions Manager - Packet client” (p. 103) is displayed.
- The default browser opens to the NSP sign-in page.

Enter your NSP user credentials, acknowledge the security statement, if present, and click SIGN IN.

Note: If you do not enter valid credentials within the allowed authentication period, the splash screen indicates that the client has timed out; you must close the splash screen and return to Step 1.

The client splash screen indicates that the client is loading, and then closes as the NFM-P GUI opens.

Figure 1-11 Network Functions Manager - Packet client

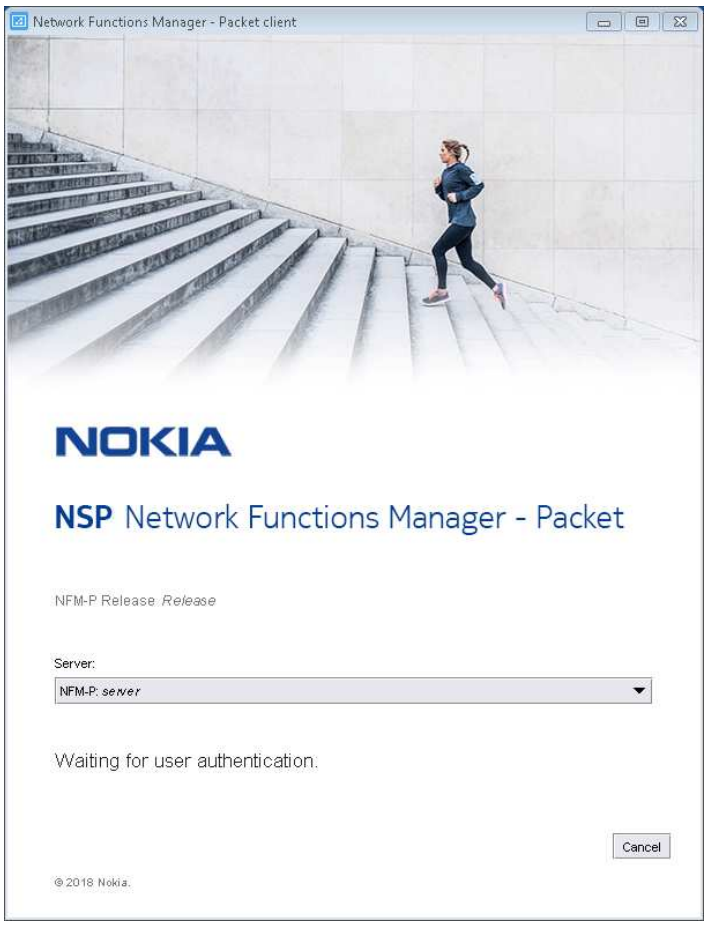

END OF STEPS

1.11 To open the NFM-P GUI through a client delegate server

1.11.1 Steps

1

If you are not in the same physical facility as the client delegate server, log in to the facility using the appropriate access tool, for example, a Citrix gateway.

2

Log in to the client delegate server station using the appropriate method:

- RHEL client delegate server—X window session

- Windows client delegate server—Windows Remote Desktop session

3

If the client delegate server is installed on a RHEL station, redirect the display to the station that you are using. See the RHEL documentation for information about display redirection.

4

Open the GUI, as described in 1.9 “To open a single-user GUI client configured for one NFM-P system” (p. 98) or 1.10 “To open a single-user GUI client configured for multiple NFM-P systems” (p. 100).

END OF STEPS

1.12 To close the NFM-P GUI

1.12.1 Steps

1

Choose Application→Exit from the NFM-P main menu.

2

Select the Save opened configuration forms check box, if required. See 1.14 “To save or open a set of forms for quick access” (p. 106) for information.

3

Click Yes. The NFM-P GUI closes, and the client session ends.

END OF STEPS

Procedures for using the GUI

1.13 To manage the display of windows and forms

1.13.1 Purpose

You can rearrange the display of windows and forms to optimize your workspace. You can also save window layouts in custom workspaces; see 2.8 “To customize window layouts” (p. 147).

Open windows and forms display an icon on the NFM-P taskbar, and are listed in the Window menu of the NFM-P main menu. You can resize, minimize, activate, and close windows and forms using standard OS methods. The following procedure describes additional methods for managing internal windows and forms.

1.13.2 Steps

1

To open the navigation tree, perform one of the following:

a. Click the toolbar icon for the required window.

   If the toolbar is hidden, you can modify user preferences to show the toolbar; see 1.23 “To configure NFM-P user preferences” (p. 116) . If the toolbar does not contain the required icon, see information about custom workspaces in Chapter 2, “NFM-P custom workspaces” .

b. Choose Application, then the name of the required window from the NFM-P main menu.

c. Use the shortcut keys; see 1.3.1 “Shortcut keys and accessibility” (p. 80) in “GUI overview” (p. 77) .

You can open multiple navigation tree windows.

2

To make a window or form the active window, choose Window from the NFM-P main menu, then choose the name of the required window.

3

To arrange multiple open forms so that they overlap, perform one of the following:

a. Choose Window→Cascade & Reset to Preferred Sizes from the NFM-P main menu. The forms are displayed overlapping in their existing sizes, showing the titlebar for each form.

b. Choose Window→Cascade from the NFM-P main menu. The forms are displayed overlapping and equally sized, showing the titlebar for each form.

4

To arrange multiple open forms without overlapping, perform one of the following:

a. Choose Window→Tile Vertical from the NFM-P main menu. The forms are displayed side by side.

b. Choose Window→Tile Horizontal from the NFM-P main menu. The forms are displayed one above the other.

c. Choose Window→Tile Square from the NFM-P main menu. The forms are displayed in a grid.

5

To minimize all open windows and forms to the taskbar, choose Window→Minimize All from the NFM-P main menu.

6

To restore windows and forms to their preferred settings, choose Window→Reset to Preferred Sizes from the NFM-P main menu.

7

To close all open windows and forms, choose Window→Close All Internal Windows from the NFM-P main menu. External windows are not affected.

When you close all internal windows, the navigation tree is closed. To reopen it, see Step 1 .

END OF STEPS

1.14 To save or open a set of forms for quick access

1.14.1 Purpose

You can save a set of open forms and reopen them all later with one operation, to provide quick access to forms that you use often or need to return to.

CAUTION

Service Disruption

Network activity can cause changes to quick-access forms when they are not open. When reopened, the saved forms may collectively exceed the memory threshold for the NFM-P GUI.

Nokia recommends that you limit the number of quick-access forms in a set. Close any unnecessary forms before you save the set of open forms. When you save a parent form, consider whether you need to save the open child forms (tabs).

If you exceed the memory threshold, restart the NFM-P client, then save a reduced number of quick access forms.

1.14.2 Before you begin

Consider the following when you save a set of forms for quick access.

- Only configuration forms that support a property form identifier are saved in the set. See 1.3.2 “NFM-P clipboard” (p. 81) in “GUI overview” (p. 77).
- When quick-access forms are reopened, they may not appear exactly as they did when they

were saved. The size and location of configuration forms is not preserved. The current values and configurations may be different from those in effect at the time the set of forms was saved.

- If an object is deleted from NFM-P management, the configuration form for that object is removed from the saved set of forms.
- Open configuration forms that are displayed in external windows are saved with the set of quick-access forms, but they are displayed as internal windows when the forms are reopened.
- All quick-access forms display an icon on the taskbar when reopened.
- If you save a set of forms when child forms are displayed as tabs in a parent form, all of the child tabs are saved as separate forms. When you reopen the forms, the child tabs are displayed as separate forms. See 1.18 “To manage configuration forms with built-in navigation trees” (p. 110) for more information.

1.14.3 Steps

1

To save a set of open forms:

1. Open the configuration forms you need to save for quick access and close any unnecessary forms.
2. Choose Window→Save Opened Configuration Forms from the NFM-P main menu. The open forms are saved as a set of quick-access forms.

   Alternatively, when you close the NFM-P client, you can select the Save opened configuration forms check box in the confirmation dialog.

2

To reopen the set of quick-access forms, choose Window→Open Saved Configuration Forms from the NFM-P main menu. The saved set of quick-access forms opens.

Alternatively, you can click the Open Saved Configuration Forms icon on the toolbar, if available.

END OF STEPS

1.15 To manage a window or form as an external window

1.15.1 Purpose

A window or form in the NFM-P GUI can be moved from the main workspace and managed as an external window, for example, on a second monitor.

External windows are managed by the client computer OS. External windows maintain the window icon that is used in the NFM-P GUI, but the icon is placed on the OS taskbar. The NFM-P GUI taskbar does not display external windows, but they are listed in the NFM-P Window menu.

Any windows or forms that are opened from an external window appear as separate external windows.

An external window or form cannot be moved back to the NFM-P GUI. You must reopen the window in the NFM-P GUI.

1.15.2 Steps

1

To move a window or form to an external window:

1. Click on the required internal window or form to make it the active window.
2. Choose Window→Move to External Window from the NFM-P main menu.

   Alternatively, you can right-click on the title bar of the selected window and choose Move To External Window.

   External windows are listed in the NFM-P Window menu. You can activate an external window by clicking on the window or form name in the Window menu list.

2

To close all external windows, choose Window→Close All External NFM-P Windows from the NFM-P main menu.

END OF STEPS

1.16 To send a text message to other NFM-P users

1.16.1 Steps

1

Choose Application→Text Message from the NFM-P main menu. The Text Message window opens.

2

Enter the text message in the text area.

You can use the NFM-P clipboard to send a property form identifier with the message. Copy the identifier for an object and click Paste to paste it into the message area. You can paste multiple objects, up to a limit. If multiple objects do not paste, reduce the size of the clipboard contents. See 1.20 “To use the NFM-P clipboard” (p. 112) for more information about the clipboard.

The most recent information copied to the system clipboard is pasted, even if it is not a property form identifier. If you use the system clipboard to copy text or other types of information after you copy a property form identifier, the identifier is no longer available.

3

Click Send To. The Select Sessions form opens with a list of active client sessions displayed.

4

Select one or more users in the list. To select all GUI client sessions, click in the list and press Ctrl-A.

5

Click OK. The form closes and the text message is sent to the selected clients.

END OF STEPS

1.17 To use configuration forms to configure or view parameters

1.17.1 Steps

1

To open a configuration form, use one of the following methods:

a. Choose the required NFM-P main menu or submenu option.

b. Right-click on an object in the navigation tree and choose the required option.

c. Double-click on the lowest-level navigation tree object.

When you double-click on an object that has a child object, the tree expands to show the child object. When you double-click on an object that has no child objects, the object properties form opens.

d. Double-click on an object in a list.

e. Select an object in a list and click Properties.

f. Choose Application→Go To Property Form from the NFM-P main menu and paste a property form identifier from the NFM-P clipboard; see 1.20 “To use the NFM-P clipboard” (p. 112) .

2

To expand or collapse panels, click on the titlebar of the panel.

You can collapse or expand all of the panels on a tabbed form. Right-click on the tab header and choose Collapse All Panels or Expand All Panels.

When a configuration form closes, the current state of the panels is saved for that object type. When a user opens a configuration form, the last saved state of the panels is displayed.

3

Configure or view the parameters on the form.

If the configuration form is a step form, follow the form prompts.

Note: A parameter field that has a yellow background is mandatory and must be configured before you can go to the next form or apply the changes on the current form. Avoid colons in the name field because the NFM-P uses colons as separators for the object full name.

4

Save the configuration changes. Perform one of the following:

a. Click Apply to save the changes without closing the form.

b. Click OK to save the changes and close the form.

c. Click Finish to save the changes in a step form.

d. Click Cancel or Close to close the form without saving the changes.

Note: Using the Ctrl-F4 keys to close a form may cause problems if the parent form is open. Nokia recommends that you do not use Ctrl-F4 to close a form that is contained by another form or that has a parent form.

Some forms display warnings or require confirmation before changes are committed. You can configure user preferences to suppress some confirmation messages; see 1.23 “To configure NFM-P user preferences” (p. 116) .

END OF STEPS

1.18 To manage configuration forms with built-in navigation trees

1.18.1 Steps

1

Open a configuration form that displays a built-in navigation tree.

When the form opens, it is unlocked, as indicated by the open lock icon on the tab header for the form.

2

To keep the form open when additional tree objects are selected, click the lock icon. Alternatively, you can right-click on the tab header and choose Lock. The icon indicates that the form is locked.

An unlocked form closes and is replaced when another object is selected in the tree.

3

Select additional objects in the built-in navigation tree. A configuration form for each object opens as a tab on the main configuration form. Lock each tab to keep the corresponding form open when a new object is selected.

If the tabs of the open forms exceed the width of the window, double arrows are displayed. Click on the double arrows to display a drop-down menu of the open configuration forms.

For service configuration forms, you can select objects in two ways:

a. Expand the tree and select an object.

b. Open an object-type list.

An object type is an item in the tree that can be expanded to display child objects.

This method is convenient when an object type contains many child objects that require a search filter.

Note: Service navigation trees for VPLS, I-VPLS, B-VPLS, MVPLS, IES, and VPRNs display a maximum of 50 access interfaces, SAPs, MSAPs, subscriber SAPs, and spoke and mesh SDP bindings at a time. When this limit is reached, you can click on the message beside the containing object under which the SAPs or access interfaces are listed to open a tab with the full list of objects.

1. Click on an object type in the tree. A list form opens as a tab and displays the objects available for the object type.
2. Select an object in the list and click Properties. If the list form was not locked, it is automatically locked. Depending on the object type, a form for the selected object will open as a tab in the main configuration form, or a separate form will open for the object.

4

To undock a form from the main configuration form, click on the tab header and choose Undock. The tab disappears and the form is displayed as a standard NFM-P form, separate from other tabs and the main configuration form.

END OF STEPS

1.19 To modify multiple objects at one time (multi-edit)

Note: Selecting large numbers of objects for multi-edit may affect system performance.

1.19.1 Steps

1

Select multiple objects in the navigation tree or a list, then right-click and choose Properties. The object type - Multiple Instances (Edit) form opens.

2

Configure the required parameters.

When a multi-value field is changed to empty and the change is applied, and that field was previously populated with a value for one of the selected multi-edit objects, a warning message is displayed. The warning message indicates the tab and parameter affected.

3

Save your changes and close the form.

END OF STEPS

1.20 To use the NFM-P clipboard

1.20.1 When to use

You can use the NFM-P clipboard to copy property form identifiers for the following:

- property forms
- statistics forms
- objects in list forms
- objects in the navigation tree
- problems encountered notifications

A property form identifier is a unique internal address that the NFM-P assigns to a form. Property form identifiers are copied to the clipboard using the Clipboard icon or by selecting the Copy to Clipboard option from a contextual menu.

You can use the NFM-P clipboard to open forms, send identifiers to other users, configure search filters, and search for objects or create physical links on the topology map.

When you copy items from a list or from the navigation tree, you can copy multiple items at one time. For other NFM-P clipboard functions, only one item at a time is copied.

Each time a property form identifier is copied to the clipboard, the previous clipboard contents are replaced.

1.20.2 Steps

1

To copy property form identifiers to the NFM-P clipboard, perform one of the following:

a. Copy objects from a list. Select one or more items in the list and click Copy to Clipboard.

b. Copy objects from the navigation tree.
   1. Expand the tree as required.
   2. Select one or more objects from the tree. The Copy to Clipboard icon is activated.

      Objects with no properties do not activate the Copy to Clipboard icon (for example, the Fans parent object). If multiple selected objects include an object without properties, the Copy to Clipboard icon is not activated.
   3. Click the Copy to Clipboard icon.

      Alternatively, you can use the contextual menu for single objects or multiple objects that are at the same level in the tree.

c. Copy a property form identifier from a form.
   1. Open the required form.
   2. Right-click on the title bar for the form and choose Copy Property Form Identifier.

      Alternatively, you can choose Application→Copy Property Form Identifier from the NFM-P main menu.

d. Copy a property form identifier from a tab.

1. Open the required tab.
2. Right-click on the tab and choose Copy Property Form Identifier. The property form identifier link is copied to the clipboard.

   Alternatively, you can choose Application→Copy Property Form Identifier from the NFM-P main menu.

2

To open a property form from the clipboard, perform one of the following:

a. Use the Go To Property Form menu item.
   1. Choose Application→Go To Property Form from the NFM-P main menu. The Go To Property Form window opens.
   2. Click Paste. The clipboard contents appear in the Identifier field.

      The most recent information copied to the system clipboard is pasted, even if it is not a property form identifier. If you use the system clipboard to copy text or other types of information after you copy a property form identifier, the identifier is no longer available.

      If the clipboard contains multiple objects, only the first object is pasted.

      You can use this method to open a specific tab, if a tab was used to save the property form identifier.
   3. Click OK or Apply. The properties form for the object opens.

b. Use the Clipboard window.
   1. Choose Application→Clipboard from the NFM-P main menu. The Clipboard window opens.
   2. Select an object and click View Object. The properties form for the object opens.

      Use this method to open the property form for an NE. To open a specific tab, use the method described in Step 2 a .

      The clipboard window shows the most recent property form identifiers copied to the NFM-P clipboard. Property form identifiers are preserved in the clipboard window even when the system clipboard is subsequently used for other types of information. However, they are replaced when new property form identifiers are copied to the NFM-P clipboard.

3

To send the clipboard contents to other users on the NFM-P server, see 1.16 “To send a text message to other NFM-P users” (p. 108) . You can also paste the clipboard contents into another application.

4

To use the clipboard for configuring advanced search filters, see 1.31 “To perform an advanced search from an object list form” (p. 128) .

5

To create physical links, see 4.10 “To create a physical link” (p. 183) .

END OF STEPS

1.21 To monitor the NFM-P Task Manager

1.21.1 Steps

1

Choose Application→Task Manager from the NFM-P main menu. The Task Manager form opens.

The All Users check box appears for NFM-P administrators. If you are an NFM-P administrator or a user with an assigned administrator scope of command role, go to Step 2 . Otherwise, go to Step 3 .

2

Select the All Users check box to monitor operations that are performed by all users in the NFM-P server.

3

Double-click on a task. The property form for the task opens.

Note: The Task Manager automatically refreshes when the value of the autoRefreshInterval parameter is reached. The default is 20. You can also click Search to refresh the list of tasks.

4

Click on the tabs to view information about the task.

5

To view the NFM-P GUI or OSS user activity associated with a task, click User Activity. The Activity form opens and displays additional information about the task. See the section on user activity logging in the *NSP System Administrator Guide* for more information.

6

Close the forms.

END OF STEPS

1.22 To save listed information to a file

1.22.1 Purpose

You can save listed information to a file, for purposes such as:

- record keeping
- inventory management

Perform this procedure to save listed information from a list form.

1.22.2 Steps

1

Open a list form and apply a filter or click Search to produce a list.

2

Right-click on any column heading and choose Save To File. The Save As form opens.

Note: The NFM-P uses the user home directory as the default location for saved files.

3

Specify the name and location of a file for the listed information. You can save the information in the following file types:

- All Files
- HTML
- CSV

You can specify the default file type on the User Preferences form; see 1.23 “To configure NFM-P user preferences” (p. 116) .

4

Click Save. The information is saved and the Save As form closes.

END OF STEPS

Procedures for configuring user preferences

1.23 To configure NFM-P user preferences

1.23.1 Purpose

The NFM-P GUI allows you to configure workspace elements and functions to suit your operational requirements. You can set many preferences using the User Preferences form accessed from the NFM-P main menu. You can also configure custom workspaces; see Chapter 2, “NFM-P custom workspaces” . You can set tab preferences to show or hide tabs on forms; see 1.24 “To set local tab preferences for configuration forms” (p. 118) .

The User Preferences form allows you to set preferences for the following GUI functions:

- show or hide the toolbar
- set the default time zone
- suppress or allow warnings and messages
- set defaults for list forms
- configure the workspace selector

You can also use the User Preferences form to set preferences for:

- STM forms
- statistics plotters; see the *NSP NFM-P Statistics Management Guide*
- scripts
- service encapsulation dot1q values
- browser path

You can import local tab preferences from a file using the User Preferences form; see 1.27 “To import local tab preferences” (p. 121) .

User preferences are specific to a particular user; the settings are associated with the user ID and are applied to the system when that user is logged in.

1.23.2 Steps

1

Choose Application→User Preferences from the NFM-P main menu. The User Preferences form opens.

2

Click on the tabs and configure the required parameters. Table 1-3, “User preferences parameters” (p. 117) describes user preferences parameters.

The Workspaces tab allows you to configure the workspace selector; see 2.14 “To configure the workspace selector” (p. 155) .

Table 1-3 User preferences parameters

Parameter	Description	See
Access Interface Encap Value (Dot1q only)	Specifies whether the Auto-Assign ID parameter is enabled by default for dot1q encapsulation	—
Apply User Span Of Control	Specifies when enabled that the navigation tree, maps, and list forms display only the objects in the Edit Access span of control of the user	1.2.2 “Span of control” (p. 79) in this chapter 1.39 “To filter using span of control” (p. 136)
Browser Path	Specifies the file path to the system browser which must then be used for the NSP. You can configure the parameter manually or by clicking the Browse button and navigating to the file path.	—
Command Helper Key	Allows you to designate the key that activates the command helper function in the Script Editor	—
Debug STM Mode	Specifies whether the NFM-P provides additional options in the object drop-down menu for the Service Test Manager. The options allow you to configure some limits and view additional information for tests.	89.14 “To configure OAM diagnostic test limits on the STM and view additional test configuration information” (p. 2957)
Default Client Time Zone	Specifies the default time zone used by the NFM-P client. Choose a time zone from the drop-down menu. The NFM-P server provides a menu default. You can modify the current client time zone without changing the default client time zone.	1.28 “To configure the current client time zone” (p. 122)
Default Polling Interval (seconds)	Specifies the default polling interval for real-time statistics in the Statistics Plotter form. The range is 10 to 3600.	*NSP NFM-P Statistics Management Guide*
Enable Command Helper	Specifies whether the Command Helper function in the Script Editor is enabled.The Command Helper is a one-key function that completes commonly occurring script commands. You can designate the command helper key.	—
Enable Confirmation for Bulk Change Actions	Specifies whether confirmation messages are displayed before the NFM-P carries out a bulk change operation	“Bulk operations using the NFM-P” (p. 713)
GUI Builder in the Editor	Specifies whether the Velocity GUI Builder in the Script Editor is enabled	—
Maximum Data Retention Time (seconds)	Specifies the number of seconds to keep statistics data in the Statistics Plotter. The range is 3600 to 86400. The default is 43200.	*NSP NFM-P Statistics Management Guide*
Save To File Default Extension	Specifies the default file type when list information is saved to a file	1.22 “To save listed information to a file” (p. 115)
Show Toolbar	Specifies whether the NFM-P toolbar is displayed in the main GUI window	1.1.6 “Toolbar” (p. 78) in this chapter

Table 1-3 User preferences parameters (continued)

Parameter	Description	See
Specify # of Items Per Page	Specifies the maximum number of items displayed per page on a list form. The range is 1 to 9999.	1.29 “To manage the display of listed information” (p. 122)
Suppress Containing Window Warning	Specifies whether warnings are displayed when saving changes to child object forms that are launched from parent object forms	1.6.7 “Warning and confirmation messages” (p. 86) in this chapter
Suppress Template Execution Warning	Specifies whether warning messages are suppressed when you execute a template	—
Suppress Template Generation Message	Specifies whether template generation messages are suppressed	—
Suppress Service and Composite Service Map Load Warning	Specifies whether load time warnings for service and composite service maps are suppressed	—

3

Save your changes and close the form.

END OF STEPS

1.24 To set local tab preferences for configuration forms

1.24.1 When to use

Custom workspaces are configured to use either local tab preferences or custom tab preferences. Local tab preferences are configured for the current user. Custom tab preferences are set as defaults in the custom workspace. See 2.9 “To configure tab preferences” (p. 147) .

When the current workspace uses local tab preferences, you can set and save tab preferences using the tab selector. When the current workspace uses custom tab preferences, you cannot set tab preferences using the tab selector.

You can export local tab preferences by saving them to a file; see 1.26 “To export local tab preferences” (p. 120) . You can import local tab preferences from a file; see 1.27 “To import local tab preferences” (p. 121) .

1.24.2 Steps

1

Open the property form for an object.

2

Click on the tab selector button at the right of the row of tabs. The tab selector drop-down opens.

Note: The keyboard shortcut to display the tab selector drop-down is Ctrl-T. The keyboard focus must be on one of the main tabs of the form. Press the Tab key to transfer the keyboard focus to the drop-down, then use the down-arrow key.
If the Tab Preferences parameter for the current workspace is set to Custom, the buttons on the tab selector do not appear.

3

Perform any of the following:

a. To display a hidden tab, select a tab name in the Hidden Tabs list and click the Move to Displayed Tabs button.

b. To hide a displayed tab, select a tab name in the Displayed Tabs Order list and click the Move to Hidden Tabs button.

c. To reorder the sequence of tabs, select a tab name in the Displayed Tabs Order list and click the Move Up or Move Down button. You can select and move multiple tabs.

Note: Tab settings are specific to the selected object type. For example, if you open the Network Element (Edit) form for a 7705 SAR-H and hide all tabs except the General tab, then open the same form for a Wavence SM, the form displays the General tab and may also display some Wavence SM specific tabs.
At least one tab must be displayed on a form. If all tabs are moved to the Hidden Tabs list, the Save Tab Settings button is dimmed.
Mandatory tabs (such as the TCA, Faults, and Deployment tabs) do not appear in the tab selector.

4

To save the changes, click on the Save Tab Settings button. The form updates with the configured tab display.

Note: If you close the tab selector without saving, all unsaved tab preferences are lost. The Find Attribute form lists all parameters on the form including parameters located on hidden tabs. When you click on the Show on Form button, the tab is temporarily displayed.

END OF STEPS

1.25 To temporarily display hidden tabs on property forms

1.25.1 Purpose

You can temporarily display the hidden tabs on a property form. Hidden tabs that are temporarily displayed are hidden again the next time the form is opened.

1.25.2 Steps

1

Open the property form for an object.

2

Perform one of the following:

a. To temporarily display all hidden tabs on a form, click Show All Tabs at the right side of the row of tabs. The form displays all available tabs.

   If no available tabs are currently hidden, the Show All Tabs button is dimmed.

   Note: The keyboard shortcut to temporarily display all hidden tabs is Ctrl-Shift-A. The keyboard focus must be on one of the main tabs for the form.

b. To temporarily display one hidden tab, perform the following:

   1. Click the tab selector button at the right of the row of tabs. The tab selector drop-down opens.

      Note: The keyboard shortcut to display the tab selector drop-down is Ctrl-T. The keyboard focus must be on one of the main tabs for the form. Press the Tab key to transfer the keyboard focus to the drop-down, then use the down-arrow key.

   2. Double-click on a tab in the Hidden Tabs list. The form displays the chosen tab.

Hidden tabs that are displayed temporarily are shown in italics in the Displayed Tabs Order list in the tab selector drop-down.

If the Tab Preferences parameter for the current workspace is set to Local, you can click on the Save Tab Settings button to set all temporarily displayed tabs as saved tab preferences. If the Tab Preferences parameter for the current workspace is set to Custom, the buttons on the tab selector do not appear.

END OF STEPS

1.26 To export local tab preferences

1.26.1 Purpose

You can export local tab preferences and import them into an NFM-P client of a different user.

1.26.2 Steps

1

Choose Application→User Preferences from the NFM-P main menu. The User Preferences form opens.

2

Click on the Tab Preferences drop-down and choose Export. The Export Directory window opens.

3

Specify the export directory in the Save In drop-down menu, or create a directory or folder.

4

Click Save. The local tab preferences are exported and saved to the specified directory.

5

Close the User Preferences form.

END OF STEPS

1.27 To import local tab preferences

1.27.1 Steps

1

Choose Application→User Preferences from the NFM-P main menu. The User Preferences form opens.

2

Click on the Tab Preferences drop-down and choose Import. The Import Directory window opens.

3

Use the Look In drop-down menu to navigate to the required directory.

Note: The folder that contains the tab preferences must contain tab preferences of only one user. If there are multiple tab preferences of different users in the specified directory, an error message appears.

4

Click Open. A warning message is displayed.

5

Click Yes to overwrite your local tab preferences with the tab preferences from the selected directory.

All affected users that currently have a client session open, other than the client session where the import has been initiated, will receive a system-generated text message informing them that their tab preferences have been changed and they should restart their NFM-P client or risk losing the changes.

The user can click Reply to reply to the message.

6

Close the User Preferences form.

END OF STEPS

1.28 To configure the current client time zone

1.28.1 Purpose

The default time zone used by the NFM-P client is configured in user preferences; see 1.23 “To configure NFM-P user preferences” (p. 116) . Perform this procedure to choose a time zone other than the default, or to return to the default client time zone.

1.28.2 Steps

1

Click the Current Client Time Zone icon in the bottom right corner of the GUI. The Current Client Time Zone form opens.

2

Perform one of the following:

a. Select Use Default Client Time Zone.

b. Select Use Specific Client Time Zone and choose a time zone from the list.

3

Save your changes and close the form.

END OF STEPS

1.29 To manage the display of listed information

1.29.1 Purpose

Perform this procedure to configure and save display preferences for list forms. You can resize and rearrange the sequence of columns, and remove columns that are not required. You can sort lists based on a selected column (attribute), and choose to display objects in ascending or descending order.

You can save listed information to a file; see 1.22 “To save listed information to a file” (p. 115).

The following table describes the order in which different types of data are sorted from the top of the column to the bottom, based on the direction of the arrow in the column header.

Table 1-4 Sorting listed information

Entry type	Down arrow (Descending)	Up arrow (Ascending)
Numbers only	Sorted from the lowest number to the highest number	Sorted from the highest number to the lowest number
Letters	Sorted left to right, character by character, in alphabetical order	Sorted left to right, character by character, in reverse alphabetical order

Table 1-4 Sorting listed information (continued)

Entry type	Down arrow (Descending)	Up arrow (Ascending)
Alphanumeric	Sorted left to right, character by character, in the following order: from 0 to 9, then in alphabetical order	Sorted left to right, character by character, in the following order: in reverse alphabetical order, then from 9 to 0
Alphanumeric for IP Address fields[1]	Addresses are sorted first by type, then alphanumerically from the lowest to the highest number. For example, IPv4, IPv6 and DNS, in the following order: • 2.1.1.1 • 10.2.2.2 • 21.2.1.1 • 21.10.1.1 • 14:C8AA:0:0:0:0:0:EA • 24AA:C8:0:0:0:0:0:EA	Addresses are sorted first by type, then alphanumerically from the highest to the lowest number. For example, DNS address, IPv6, IPv4 in the following order: • 24AA:C8:0:0:0:0:0:EA • 14:C8AA:0:0:0:0:0:EA • 21.10.1.1 • 21.2.1.1 • 10.2.2.2 • 2.1.1.1
Special characters	Sorted left to right, character by character. The following characters are sorted in the following order before numbers, and uppercase and lowercase letters in a list: (space) ! # $ % & ‘ ( ) * + , . / The following characters are sorted in the following order after numbers in a list, but before uppercase and lowercase letters: : ; < = > ? @ The following characters are sorted in the following order after numbers, and uppercase and lowercase letters in a list: { | } ~	Sorted left to right, character by character. The following characters are sorted in the following order before numbers, and uppercase and lowercase letters in a list: ~ } | { The following characters are sorted in the following order before numbers in a list, but after uppercase and lowercase letters: @ ? > = < ; : The following characters are sorted in the following order after numbers, and uppercase and lowercase letters in a list: / . , + *) (‘ & % $ # ! (space)
Blanks	Blank entries are placed first at the top of the list.	Blank entries are placed last at the bottom of the list.

Notes:

1. Sorting on an IP Address field may cause performance issues when the NFM-P database is under heavy load.

The following figure describes the features of a list form.

Figure 1-12 List elements

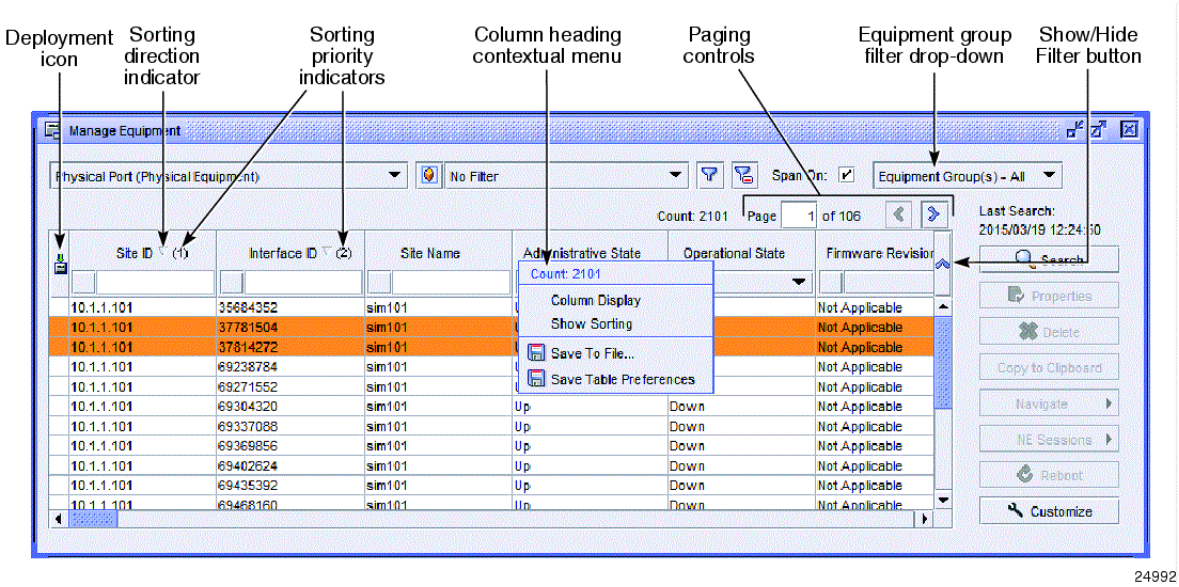

1.29.2 Steps

1

Open a list form and choose an object type. The columns specific to the object type are displayed.

2

Configure column display preferences. Perform any of the following:

a. Remove columns.

1. Right-click on any column heading and choose Column Display. The Column Display form opens.
2. In the Displayed on Table panel, select one or more columns to remove.
3. Click on the left arrow to move the selected items to the Available for Table panel.

   Dimmed items cannot be moved.

   There must be at least one item in the Displayed on Table panel.
4. Save your changes and close the form.

b. Add columns.

1. Right-click on any column heading and choose Column Display. The Column Display form opens.
2. In the Available for Table panel, select one or more columns to add.
3. Click on the right arrow to move the selected items to the Displayed on Table panel.

4. Save your changes and close the form.

c. Rearrange the left-to-right sequence of columns.

Use this feature to place important information in the first few columns of the table, or to place related columns side by side.

1. Right-click on any column heading and choose Column Display. The Column Display form opens.
2. In the Displayed on Table panel, select the name of one or more columns to rearrange.
3. Click on the up or down arrows to move the selected items. The top-to-bottom sequence of the items reflects the left-to-right sequence of the columns in the list table.
4. Save your changes and close the form.

Alternatively, you can change the sequence of columns by clicking and dragging on a column heading.

d. Resize columns. Hover over the vertical line between two columns, then click and drag when the arrows appear.

3

Configure a filter if required and click Search to generate a list.

The total number of listed items is indicated by the Count value displayed on the form and in the contextual menu for column headings.

4

Hide the simple search fields in the column headings. Click on the Show/Hide Filter button.

5

Navigate multiple pages.

When a list uses multiple pages, the paging controls are available. Click on the paging arrows to move forward or backward one page at a time, or enter a page number to jump to a specific page. The page number and total page count are shown on the form.

You can specify a user preference for the number of items to display on a page; see 1.23 “To configure NFM-P user preferences” (p. 116) .

6

Configure sorting preferences.

Lists are sorted top-to-bottom based on selected column attributes. You can display sorted results in ascending or descending order; see Table 1-4, “Sorting listed information” (p. 122) . When no column is explicitly selected for sorting, the first column at the left of the table is used.

Perform one of the following:

a. Sort a list using a single column.

1. Click on the column heading. The list is sorted based on the values for the selected attribute.
2. Click on the column heading to reverse the sorting direction (ascending or descending).

b. Sort using multiple columns.

When a sort using a single column produces many items that all have the same value for the selected attribute, you can refine the sort by choosing additional columns. The priority order for the columns is indicated by a number in parentheses in the column heading.

1. Right-click on any column heading and choose Show Sorting. The Show Sorting form opens.
2. In the Used for Sorting panel, select any column names that are not required for sorting and click on the left arrow to move the selected items.
3. In the Available for Sorting panel, select the column names to be used for sorting and click on the right arrow. The selected items appear in the Used for Sorting panel.
4. In the Used for Sorting panel, click on the items and use the up and down arrows to arrange them in the required priority order.

   The sorting priority decreases from top to bottom; the item at the top of the panel has priority (1).
5. Click to select Sort Ascending or Sort Descending; see Table 1-4, “Sorting listed information” (p. 122) . The selected column headings display numbers in parentheses to show the sorting priority.
6. Close the form.

Alternatively, you can click on a column heading to select the priority (1) attribute, then Ctrl-click on additional column headings to further refine the sort. Priority is based on the order in which the columns are clicked. Press Ctrl-click on an already selected (numbered) column heading to reverse the sorting direction (ascending or descending).

You can navigate quickly through a long sorted list using the first character of the column that has sorting priority (1). Click in the list, then type a letter or number. The first item that begins with that letter or number in the priority (1) column is selected in the list.

7

Save your column display and sorting preferences.

1. Right-click on any column heading and choose Save Table Preferences.
2. Confirm the action. The configured display preferences are applied whenever you open a list table for the same object type.

END OF STEPS

Procedures for searching

1.30 To perform a simple search from an object list form

1.30.1 Purpose

You can perform a simple search using the column headings of a list table. Each column heading displays an attribute. You can configure a filter property for each attribute by selecting a function and a value from the column heading fields. See Figure 1-1, “List form” (p. 83). The available functions and values vary depending on the selected attribute and other settings. You can configure one function and value for each attribute.

You can show or hide the column heading search fields. You can also display a simple search filter on the filter configuration form, then save or modify it using the advanced search functions; see 1.31 “To perform an advanced search from an object list form” (p. 128) . When multiple filter properties are configured using column headings, the properties are combined using the AND Boolean operator.

Simple searches are disabled when:

- the filter configuration form is open
- a saved filter is selected in the saved filter drop-down
- a saved filter is loaded using the Saved Filters form

See “NFM-P searches” (p. 92) in this chapter for more information.

1.30.2 Steps

1

Open a list form and choose an object type, if required.

You can reduce the number of items in the object drop-down menu. See 1.38 “To filter object types” (p. 136) and 2.13 “To customize list forms” (p. 153) .

2

Choose No Filter from the saved filter drop-down menu, if required.

3

Click Show/Hide Filter to show the column heading search fields, if required.

4

Configure one or more filter properties.

1. Decide on an attribute for your filter and locate the corresponding column heading in the list.

   When a column heading has no search fields, the attribute in that column is not searchable.

2. Click on the function field in the column heading and choose a function. The symbol for the function is displayed. See Table 1-2, “Search functions” (p. 93) for a description of search functions.

3. Click on the value field and enter a value or choose one from the list. The value is displayed.

   If you enter a value without choosing a function first, a default function is used.

5

Configure the Span On parameter to override the global span of control preference, if required. See 1.39 “To filter using span of control” (p. 136) .

6

Click Search. A list of filtered results is displayed.

The Search button applies the filter even if the column heading search fields are hidden.

7

Revise the filter properties, if required.

To clear a filter property, choose (NONE) from the function menu for that attribute. The filter property is cleared from the column heading.

To clear all filter properties, click the Clear All Filters icon. All column headings return to an unconfigured state.

END OF STEPS

1.31 To perform an advanced search from an object list form

1.31.1 Purpose

You can use the filter configuration form to perform an advanced search on a list. Figure 1-6, “Filter configuration form” (p. 96) shows the filter configuration form. Advanced search filters are more detailed than simple search filters.

Multiple filter properties in an advanced search filter are combined using Boolean operators. You can apply brackets to properties or groups of properties, and apply Boolean operators to the groups. Table 1-5, “Boolean search operators” (p. 128) describes Boolean search operators.

Table 1-5 Boolean search operators

Boolean operator	Description
AND	When you combine properties using the AND operator, the search returns objects that meet both or all of the specified criteria.
OR	When you combine properties using the OR operator, the search returns objects that meet at least one of the specified criteria.
NOT	When a filter property is preceded by the NOT operator, items that meet the criteria in that filter property are excluded from the results.

You can save search filters for reuse or modification; see 1.33 “To save search filters” (p. 132) .

1.31.2 Steps

1

Open a list form and choose an object type.

You can reduce the number of items in the object menu; see 1.38 “To filter object types” (p. 136) and 2.13 “To customize list forms” (p. 153) .

2

Click the Filter icon. The filter configuration form opens.

3

Configure a filter property.

You can load a previously saved filter and modify it; see 1.34 “To use a saved search filter” (p. 133) .

1. Click on the Attribute field and choose an attribute.
2. Click on the Function field and choose a function. See Table 1-2, “Search functions” (p. 93) for a description of search functions.
3. Click on the Value field and choose a value, or enter a value manually.

   Alternatively, you can use the NFM-P clipboard to enter a value. Copy the property form identifier for a form that contains the same attribute and value that you are searching, then click on the Paste Attribute from Clipboard icon. The value field is populated with the required value. This method is useful when the value you are entering is a long text string, or when your search uses a value (such as a customer name) that you know is on a particular form. See 1.20 “To use the NFM-P clipboard” (p. 112) for more information about using the clipboard.

   The Paste Attribute from Clipboard icon is available when a search attribute is selected in 1 . A value is pasted only when the clipboard contains an attribute of the same name and type as the search attribute. If the clipboard contains multiple objects, the value for the first object is pasted.
4. Click Add. The filter property appears in the Filter display field and the Operators drop-down is available.

4

If required, change the Boolean operator. Click on the Operators drop-down and select an operator. See Table 1-5, “Boolean search operators” (p. 128) for information about Boolean operators.

The Boolean operator shown in the Operators field is inserted between the filter properties when the next filter property is added.

The Operators drop-down is available only when a filter property is displayed in the Filter display field.

5

Repeat Step 3 and Step 4 to configure additional filter properties.

6

Choose a Span setting; see 1.39 “To filter using span of control” (p. 136) and 1.2.2 “Span of control” (p. 79) in this chapter.

7

To modify filter properties, perform any of the following:

a. Add or remove brackets.
   1. Click and drag to select a filter property or properties in the Filter display field.
   2. Click Brackets. The selected properties are wrapped in brackets. When properties are already wrapped in brackets, the Brackets button unwraps the selected properties.

      You can wrap or unwrap multiple properties only when all Boolean operators within the brackets are the same.

b. Replace a filter property.
   1. Click and drag to select a filter property in the Filter display field.
   2. Configure an attribute, function, and value; see Step 3.
   3. Click Replace. The configured filter property replaces the selected filter property.

c. Exclude filter properties.
   1. Click and drag to select a filter property in the Filter display field.
   2. Click NOT. The property is wrapped in brackets and preceded by the NOT operator. Objects that match the selected property are excluded from the results.

      You can exclude multiple properties at one time.

      You can select excluded properties and click NOT to include the properties again.

d. Replace a Boolean operator.
   1. Click to select a Boolean operator in the Filter display field.
   2. Choose a Boolean operator from the Operators menu. The new selection replaces the previous operator.

      When you replace a Boolean operator, the properties associated with that operator are wrapped in brackets.

      You can replace all of the Boolean operators in a bracketed group by selecting the group and choosing a Boolean operator from the Operators menu.

e. Delete filter properties.
   1. Click and drag to select one or more filter properties in the Filter display field.
   2. Click Delete. The selected properties are deleted from the filter.

      To delete all properties from the Filter display field, click Clear.

8

Click Apply to perform the search. The filtered results are displayed.

9

To save the configured filter, see 1.33 “To save search filters” (p. 132). Otherwise, close the filter configuration form.

10

Select the required object in the list.

END OF STEPS

1.32 To perform a search by specifying endpoints

1.32.1 Purpose

Perform this procedure to find objects such as service tunnels, MPLS paths, and LSP paths by specifying endpoints. You can search based on a source site endpoint, a destination site endpoint, or both. You can also search based on multiple source or destination endpoints at the same time.

1.32.2 Steps

1

Choose Manage→MPLS→*<object submenu>* from the NFM-P main menu.

The Manage filter form for the selected object opens.

2

If required, specify the object endpoint type in the Type column, otherwise go to Step 3 .

3

Click the Filter icon. The filter configuration form opens.

4

Choose Endpoints from the Select Filter Type menu.

Note: If a saved filter is selected, or a simple search filter is configured in the column headings, the Select Filter Type displays Advanced Filter by default.

To configure an advanced search, see 1.31 “To perform an advanced search from an object list form” (p. 128) .

5

If you know the source endpoint for the tunnel or path, deselect the Any Source parameter and choose the required source site. Otherwise, go to Step 6 .

You can choose multiple source sites in one filter.

6

If you know the destination endpoint for the tunnel or path, deselect the Any Destination parameter and choose the required destination site.

You can choose multiple destination sites in one filter.

7

Click Apply to perform the search. The filtered results are displayed.

8

To save the configured filter, see 1.33 “To save search filters” (p. 131) . Otherwise, close the filter configuration form.

9

Select the required object in the list.

END OF STEPS

1.33 To save search filters

1.33.1 Purpose

You can save search filters for reuse or modification. Saved filters appear in the filter menu for the list form; see Figure 1-1, “List form” (p. 83) . Saved filters apply only to the specific object type for which they were created, and descendant objects of that object type. You can view saved filters by clicking on the Saved Filters button of the filter configuration form.

1.33.2 Steps

1

On a list form, click the Filter icon. The Filter configuration form opens.

See 1.31 “To perform an advanced search from an object list form” (p. 128) for information about advanced search filters.

See 1.32 “To perform a search by specifying endpoints” (p. 131) for information about endpoint search filters.

2

Click Save. The Save Filter form opens.

3

Configure the required parameters.

The Public parameter specifies whether a filter can be accessed by other users. When the Public parameter is disabled, the filter is private and cannot be accessed by other users.

4

Click Save. The Save Filter form closes, and the saved filter is available for future searches on the same object type.

5

Close the forms.

END OF STEPS

1.34 To use a saved search filter

1.34.1 Purpose

Perform this procedure to load and reuse saved search filters from a list form or from the filter configuration form.

Saved filters can be selected and applied directly from list forms, or from the filter configuration form. You can modify filters that are selected from the filter configuration form. Only filters that are applicable to the current object type and its descendant types are available.

Saved filters may contain attributes or values that are not valid for the current object. When filters contain invalid attributes and values, the NFM-P alerts you by displaying red text. You can select such filters and use them to perform a search. However, the invalid attributes are not evaluated. See 1.8.3 “Invalid attributes or values” (p. 96).

1.34.2 Steps

1

Open a List Form.

2

Perform one of the following:

a. Reuse a saved search filter directly from the list form. Click on the saved filters drop-down and select a filter in the list. The filter is immediately applied and the search results are displayed.

   When a saved filter is selected, simple searches using the column headings are disabled.

b. Load and reuse a saved filter using the filter configuration form.

   When you use this method, you can search for or modify a saved filter.

   1. Click the Filter icon. The filter configuration form opens.
   2. Click Saved Filters. The Saved Filters form opens.
   3. Select a filter in the list and click Load. The Saved Filters form closes and the filter properties are displayed on the filter configuration form.
   4. Modify the filter properties as required; see 1.31 “To perform an advanced search from an object list form” (p. 128) .

5. Save the modified filter, if required; see 1.33 “To save search filters” (p. 132) .

   The filter is saved for the current object type.

6. Click Apply to perform the search, then close the filter configuration form. The search results are displayed.

End of steps

1.35 To delete a saved search filter

1.35.1 Steps

1

Open a list form and click the Filter icon. The filter configuration form opens.

2

Click the Saved Filters button. The Saved Filters form opens.

3

Select a saved filter and click Delete. The filter is deleted.

4

Close the forms.

End of steps

1.36 To copy an advanced search filter

1.36.1 Purpose

You can copy saved advanced search filters for reuse or modification. You can paste a copied filter into the advanced filter configuration form; see Figure 1-1, “List form” (p. 83).

Consider the following when you copy and paste an advanced search filter:

- all of the attributes in the copied filter must have a matching property in the Object Type of the Advanced Filter window into which the filter expression is pasted
- the matching attributes must have the same displayed name, and the displayed name is case sensitive
- the matching attributes must be of the same type, for example, string, Boolean, etc.
- the corresponding properties must have a matching Function

1.36.2 Steps

1

Open a list form.

2

Click the Filter icon to open the Advanced Filter form.

3

Click Saved Filters. The Saved Filters form opens.

4

On the Saved Filters window, choose the filter you want to copy and click Copy Filter.

5

On a new filter configuration form, click Paste Filter. You can paste the filter expression only into a filter for the same object type as the filter you copied.

Note: If the filter is not compatible, the paste is not successful. An error is generated in the server.log file that indicates the attribute that failed in the paste process.

END OF STEPS

1.37 To locate an attribute on a configuration form

1.37.1 Purpose

The Find icon on configuration forms allows you to quickly locate tabs, panels, and parameters on configuration and property forms.

For service configuration and NE property forms with built-in navigation trees and tabs for multiple objects, the search is performed only on the active object.

1.37.2 Steps

1

Open a configuration form.

2

Click the Find icon or press Ctrl-F. The Find Attribute in form opens.

3

Select an attribute in the list and click Show on Form. On the configuration form, the selected name of the tab, panel, or parameter is highlighted and the associated tab is displayed.

Alternatively, you can double-click on the attribute row, or select an attribute and press Enter.

Note: If you select a parameter or panel that resides on a hidden tab, the hidden tab is temporarily displayed.

4

Close the Find Attribute in form.

END OF STEPS

1.38 To filter object types

1.38.1 Purpose

When the object type drop-down menu in a list form contains more than 10 items, you can use a filter to reduce the number of items.

You can also set defaults for list forms in custom workspaces, to reduce the number of items in the object type drop-down menu and to set the default object type; see 2.13 “To customize list forms” (p. 153) .

1.38.2 Steps

1

Open a list form and click the Filter for Object Type icon; see Figure 1-1, “List form” (p. 83) . The Select Object Type form opens.

2

Configure a simple search using column headings; see 1.30 “To perform a simple search from an object list form” (p. 127) .

Alternatively, you can configure an advanced search; see 1.31 “To perform an advanced search from an object list form” (p. 128) .

3

Click Search. A filtered list of object types is displayed.

4

Select an item from the list and click OK. The Select Object Type form closes and the selected item is displayed in the Object Type field.

END OF STEPS

1.39 To filter using span of control

1.39.1 Purpose

The NFM-P GUI displays objects and allows object configuration based on the span of control of the user and user preferences. You can configure user preferences to specify that only objects within the Edit Access span of control of the current user are displayed in list forms and on the navigation tree. See 1.23 “To configure NFM-P user preferences” (p. 116) for more information about user preferences.

For more information about span of control, see the section on creating a span of control in the *NSP System Administrator Guide*.

1.39.2 Steps

1

Open a List Form and choose an entry from the object drop-down menu.

2

Perform one of the following.

a. Configure span settings for a simple search.

1. Configure filter properties, if required; see 1.30 “To perform a simple search from an object list form” (p. 127) .
2. Configure the Span On check box to override the user preferences default span setting.
3. Click Search. A filtered list of objects is displayed.

   When Span On is enabled, only the objects that are within the Edit Access span of control of the current user are displayed.

b. Configure span settings for an advanced search.

1. Click the Filter icon to open the filter configuration form.
2. Configure filter properties, if required; see 1.31 “To perform an advanced search from an object list form” (p. 128) .
3. Configure the Span parameter. The following table describes the Span parameter options.

Option	Description
Span Off	Span of control filtering is disabled; objects in the View Access and Edit Access spans of the current user are displayed.
Span On	Span of control filtering is enabled; only objects in the Edit Access spans of the current user are displayed.
User Preference	Span of control filtering is enabled or disabled, as configured on the User Preferences form.

4. Click Apply and close the filter configuration form. The search results are displayed.

3

Select the required object in the list.

END OF STEPS

1.40 To configure and save equipment group filters

1.40.1 Purpose

Some list forms allow filtering based on equipment groups. An equipment group filter is created by selecting from existing equipment groups. You can save equipment group filters for reuse or modification.

Equipment group filters can return up to 12 000 items. If this limit is exceeded, you are prompted to modify the filter to reduce the number of items.

If you choose an equipment group filter that includes a group that has been modified since the filter was applied, an icon appears beside the equipment group filter drop-down prompting you to click the Search button to update the filter.

1.40.2 Steps

1

Open a list form and choose an entry from the object type drop-down menu. For object types based on equipment, the Equipment Group filter drop-down is displayed.

2

Choose Select Equipment Groups from the Equipment Group filter drop-down menu. The Select Equipment Groups form opens with equipment groups listed in the Available panel.

3

Perform one of the following:

a. Create an equipment group filter.
   1. In the Available panel, select the required equipment groups and click on the right arrow. The selected groups are moved to the Filtered by panel.

      To empty the Filtered by panel, click Clear. All groups in the Filtered by panel move into the Available panel.
   2. To include data from equipment in equipment subgroups, select the Include contained Equipment Groups parameter.

b. Modify an equipment group filter.
   1. Click Saved Filters. The Saved Filters form opens.
   2. Select a filter and click Load. The Saved Filters form closes and the selected filter is displayed in the Select Equipment Groups form.
   3. To remove groups from the Filtered by panel, select the groups and click the left arrow.

      To empty the Filtered by panel, click Clear. All groups in the Filtered by panel move into the Available panel.
   4. To add groups to the Filtered by panel, select the groups in the Available panel and click the right arrow.

4

Save the equipment group filter, if required.

1. Click Save. The Save Equipment Group Filter form opens.
2. Configure the required parameters.

   The Public parameter specifies whether a filter can be accessed by other users. When the parameter is disabled, the filter is private and cannot be accessed by other users.
3. Click Save. The Save Equipment Group Filter form closes, the filter is saved, and the Select Equipment Groups form displays the saved equipment group filter.

5

Click OK to apply the filter. The Select Equipment Groups form closes and the results list is displayed based on the equipment group filter and any additional filters that are applied to the list. The Last Search time is updated, and the equipment group filter name is displayed in the equipment group filter drop-down.

END OF STEPS

1.41 To use a saved equipment group filter

1.41.1 Steps

1

Open a list form and choose an entry from the object type drop-down. For object types based on equipment, the Equipment Group filter drop-down is displayed.

2

Click on the Equipment Group filter drop-down and perform one of the following:

a. Reuse a saved equipment group filter directly from the list form. Choose a saved filter from the drop-down menu. The filter is immediately applied and the search results are displayed.

b. Load and reuse a saved filter using the Select Equipment Groups form.

   When you use this method you can search for or modify a saved filter.

   1. Choose Select Equipment Groups from the Equipment Groups filter drop-down menu. The Select Equipment Groups form opens.
   2. Click Saved Filters. The Saved Filters form opens.
   3. Select a filter in the list and click Load. The filter is displayed on the Select Equipment Groups form.
   4. Click OK to apply the filter. The Select Equipment Groups form closes and the search results are displayed.

END OF STEPS

2 NFM-P custom workspaces

NFM-P custom workspaces overview

2.1 Workspace customization

2.1.1 General information

A workspace is a configuration of the main NFM-P GUI elements, such as the window layout or menu options. System-defined workspaces are provided, but you can also create and save your own customized workspaces to simplify navigation and operation according to your requirements.

There are three types of workspaces:

- System-defined—these workspaces are read-only and are provided as part of the NFM-P product. System-defined workspaces cannot be modified.
- Private—a private workspace contains your customized settings and is visible only to you and an administrator.
- Public—a public workspace is usable by all users on the same server. Only users with the required permissions can modify a public workspace.

You can customize the following workspace elements:

- Window layouts—save your window location and size preferences for the navigation tree and main NFM-P GUI window.
- Tab preferences—save tab preferences as part of a workspace.
- Menus—create custom menus with user-defined names, remove menu items that you don’t need, re-order the sequence in which menu items appear, rename menu items, and add grouping separators and custom mnemonics.
- Toolbar—add, remove, re-order and rename menu icon buttons on the toolbar. You can also add custom toolbar menus, which contain selected menu items and optional grouping separators.
- Tree labels—on the routing and equipment navigation trees, choose the information to display in the label text fields, and the order in which it appears.
- List forms—organize the content of the object type drop-down menus and select which buttons are displayed on some specific list forms.

You can create new custom workspaces based on a system-defined workspace or on another custom workspace, and you can modify existing custom workspaces. You must have the required scope of command permissions to create or modify public workspaces. Scope of command roles and user group workspace permissions are defined by a system administrator.

You can save custom workspaces by exporting them to a specified directory. You can share these saved workspaces with another user or reuse these saved workspaces by importing them into another NFM-P server.

Later releases of the NFM-P may contain newly-introduced menu items that are not in the custom workspaces of an earlier release. You can add new menu items to custom workspaces that were exported in an earlier release and then imported in a later release. For convenience, the NFM-P

separately lists the newly-introduced menu items; see 2.19 “To add new menu items to a custom workspace of an earlier NFM-P release” (p. 159) .

The workspace selector located in the upper right-hand corner of the screen contains a list of system-defined and custom workspaces, and allows you to replace the current workspace with another from this list. You can add or remove custom workspaces in the workspace selector, and change the order in which workspaces appear in the list. See 2.14 “To configure the workspace selector” (p. 155) and 2.15 “To apply a different workspace using the workspace selector” (p. 156) .

Note: Custom settings for tabs on configuration forms are saved as user preferences. They are not configured or saved as part of custom workspaces. See 1.24 “To set local tab preferences for configuration forms” (p. 118) and 1.25 “To temporarily display hidden tabs on property forms” (p. 119) .

2.2 Workflow to administer NFM-P custom workspaces

2.2.1 Purpose

A system administrator can perform the following tasks to configure users and permissions for workspaces, and to export custom workspaces.

2.2.2 Stages

1

Configure the scope of command roles for custom workspaces; see “To create a scope of command role” in the *NSP System Administrator Guide*.

2

Add or remove workspaces for a User Group and set mandatory workspaces settings; see “To add or remove workspaces for a user group” in the *NSP System Administrator Guide*.

3

Save all custom workspaces to a directory using the Export All function; see “To export all workspaces and local tab preferences” in the *NSP System Administrator Guide*.

4

Import workspaces that were previously saved as exported workspaces; see “To import workspaces and local tab preferences” in the *NSP System Administrator Guide*.

2.3 Workflow to customize NFM-P workspaces

2.3.1 Stages

1

Assess your operational requirements to determine your workspace needs.

2

Create a custom workspace; see 2.6 “To create a new custom workspace” (p. 145) .

3

Configure window layouts; see 2.8 “To customize window layouts” (p. 147) .

4

Configure tab preferences; see 2.9 “To configure tab preferences” (p. 147) .

5

Customize menus; see 2.10 “To customize menus” (p. 148) .

6

Customize toolbars; see 2.11 “To customize toolbars” (p. 150) .

7

Customize tree labels; see 2.12 “To customize tree labels” (p. 152) .

8

Customize list forms; see 2.13 “To customize list forms” (p. 153) .

9

Configure the workspace selector; see 2.14 “To configure the workspace selector” (p. 155) .

10

Modify an existing workspace, as required. See 2.7 “To modify an existing workspace” (p. 146) .

11

Update a custom workspace of an earlier NFM-P release with new menu items from a later release; see 2.19 “To add new menu items to a custom workspace of an earlier NFM-P release” (p. 159) .

12

Delete a workspace when it is no longer required; see 2.16 “To delete a custom workspace” (p. 157) .

2.4 Workflow to share workspaces

2.4.1 Purpose

You can share custom workspaces with other users of the NFM-P. The following workflow lists the sequence of steps to follow for sharing workspaces.

2.4.2 Stages

1

To share a private workspace with other users, change the scope of the workspace from Private to Public. You must have the required scope of command permissions for public workspaces. See 2.6 “To create a new custom workspace” (p. 145) .

2

To share a custom workspace by saving it to a directory:

1. Export custom workspaces to a specified directory; see 2.17 “To export custom workspaces” (p. 157) .
2. Import workspaces from a directory; see 2.18 “To import a workspace” (p. 158) .

NFM-P GUI custom workspace procedures

2.5 Overview

2.5.1 Purpose

The following procedures describe how to create and configure custom workspaces in the NFM-P.

2.6 To create a new custom workspace

2.6.1 Steps

1

Choose Application→Manage Workspaces from the NFM-P main menu. The Manage Workspaces form opens.

2

Perform one of the following:

a. Create a new workspace based on the system default workspace. Click Create. The Workspace (Create) form opens.

b. Create a new workspace based on an existing custom workspace:

1. Choose a custom workspace in the list and click Properties. The Workspace (Edit) form opens.
2. Click Copy.

3

Configure the required parameters.

Note: You must have the required scope of command permissions to create a public workspace.

4

Configure window layouts; see 2.8 “To customize window layouts” (p. 147) .

5

Configure tab preferences; see 2.9 “To configure tab preferences” (p. 147) .

6

Customize menus; see 2.10 “To customize menus” (p. 148) .

7

Customize toolbars, see 2.11 “To customize toolbars” (p. 150) .

8

Customize tree labels; see 2.12 “To customize tree labels” (p. 152) .

9

Customize list forms; see 2.13 “To customize list forms” (p. 153) .

END OF STEPS

2.7 To modify an existing workspace

Note: You cannot modify a system-defined workspace.

2.7.1 Steps

1

Perform one of the following:

a. To modify the current custom workspace, choose Application→Customize Current Workspace from the NFM-P main menu. The Workspace (Edit) form opens.

b. To modify a workspace other than the current workspace, perform the following:

1. Choose Application→Manage Workspaces from the NFM-P main menu. The Manage Workspaces form opens.
2. Choose a custom workspace in the list and click Properties. The Workspace (Edit) form opens.

2

Configure window layouts; see 2.8 “To customize window layouts” (p. 147) .

3

Configure tab preferences; see 2.9 “To configure tab preferences” (p. 147) .

4

Customize menus as required; see 2.10 “To customize menus” (p. 148) .

5

Customize toolbars as required; see 2.11 “To customize toolbars” (p. 150) .

6

Customize tree labels as required; see 2.12 “To customize tree labels” (p. 152) .

7

Customize list forms; see 2.13 “To customize list forms” (p. 153) .

END OF STEPS

2.8 To customize window layouts

2.8.1 Purpose

You can customize the sizes and locations for the navigation tree and the main NFM-P GUI window. See 1.13 “To manage the display of windows and forms” (p. 105) for information about how to arrange windows in the NFM-P GUI.

Perform this procedure to specify the window layout that appears when the custom workspace opens. You can specify the system default or the current window layout.

2.8.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

2

Configure the Window Layout parameter.

You can only select the Set to Default option when Custom is displayed in the Window Layout field.

3

Save your changes and close the form.

END OF STEPS

2.9 To configure tab preferences

2.9.1 Purpose

You can configure a workspace to use either local tab preferences or custom tab preferences. Local tab preferences can be set manually using the tab selector, and are saved for a specific user. Custom tab preferences are saved as part of the workspace, with settings copied from either the currently open GUI client or another workspace. When custom tab preferences are used, the tab selector for configuration forms is disabled.

For more information about tab preferences, see 1.24 “To set local tab preferences for configuration forms” (p. 118) .

2.9.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

2

Configure the Tab Preferences parameter.

If you choose Set from Workspace, the Select Workspace form opens; go to Step 3 . Otherwise, go to Step 4 .

3

Select a custom workspace and click OK. The tab preferences of the selected workspace are copied to the workspace you are configuring.

4

Save your changes and close the form.

Note: Use the Copy to Local button to copy the custom tab preferences saved in the workspace to the local tab preferences saved for the user. The Copy to Local button is not shown during initial workspace creation, and is active only when the Tab Preferences parameter for the workspace is set to Custom.

END OF STEPS

2.10 To customize menus

2.10.1 Purpose

Perform this procedure to customize menus in a new workspace or to modify menus in an existing workspace.

2.10.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

Note: If the workspace was created in an earlier release of the NFM-P, see 2.19 “To add new menu items to a custom workspace of an earlier NFM-P release” (p. 159) .

2

Click on the Menus tab.

The All Menu Items panel contains a list of all available menu items in the NFM-P. The Custom Menu panel contains a list of all menu items in the custom workspace.

3

Expand items as required in the All Menu Items panel and Custom Menu panel.

4

Perform any of the following:

a. In the Custom Menu panel, remove any menu items that are not required in the workspace. Select and delete the items using the right-click menu or the keyboard.

Note: You cannot delete the User Preferences or Manage Workspaces menu items.

b. Add menu items to the Custom Menu panel. Select items in the All Menu Items panel and drag and drop or copy the selected items to the Custom Menu panel.

The usage indicator icon beside a menu in the All Menu Items panel indicates whether the menu items for that menu are all used, partially used, or not used in the Custom Menu. Usage indicator icons are described in the Legend at the bottom of the Workspace form. The usage indicators reduce the need to compare menu content between the All Menu Items panel and the Custom Menu panel.

c. Add a new custom menu in the Custom Menu panel.

1. Right-click in the location where you need to add a custom menu and choose Add Menu. A new menu appears under the selected item.

   If you right-click on empty space in the panel, the new menu appears as the last item in the list.

2. Type a name for the new menu and press Enter.

d. To locate a menu item in the All Menu Items panel, display the menu path for the item.

1. In the Custom Menu panel, select an item, right-click, and choose Show in All Menu Items Tree.

   A dialog box displays the path for the item in the All Menu Items panel.

2. Click OK to close the dialog box.

e. To locate a menu item in the Custom Menu panel, display the menu path for the item.

1. In the All Menu Items panel, select an item, right-click, and choose Show in Custom Menu Tree.

   A dialog box displays the path for the item in the Custom Menu. If the item appears more than once in the Custom Menu tree, multiple paths are displayed.

2. Click OK to close the dialog box.

f. Arrange the order of menus and menu items. Drag and drop or cut and paste selected menu items in the Custom Menu panel to other locations in the tree.

If you hold the Ctrl key while dragging, the drop action changes from move to copy.

g. Rename menus and items in the Custom Menu panel.

1. Select the item, right-click, and choose Rename.
2. Type a name and press Enter.

Renamed menu items are displayed in blue.

h. Add grouping separators in sub-menus. Right-click in the location where you need to add a separator and choose Add Separator.

i. Edit the mnemonic for an item in the Custom Menu panel.

1. Select the item, right-click, and choose Edit Mnemonic. A drop-down menu opens.
2. Choose a letter and click OK.

   The mnemonic letter is underlined in the menu item name and is used as a keyboard shortcut.

5

Save your changes and close the form.

Note: When you modify a workspace that is currently in use, the workspace selector displays Workspace Out of Sync. Select the current workspace from the workspace selector to apply the modified settings.

END OF STEPS

2.11 To customize toolbars

2.11.1 Purpose

The NFM-P toolbar contains buttons that allow quick access to menus and menu items; see 1.1.6 “Toolbar” (p. 78) in Chapter 1, “NFM-P GUI” . Perform this procedure to add or remove buttons and to change the order of buttons on the toolbar.

2.11.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

2

Click on the Toolbar tab.

The Custom Menu panel contains a list of all the menu items in the custom workspace. The Appear on Toolbar panel displays a list of all items displayed as buttons or drop-down menu items on the toolbar.

3

Expand items as required in the Custom Menu panel and Appear on Toolbar panel.

4

Perform any of the following:

a. In the Appear on Toolbar panel, remove any items that are not required on the toolbar. Select and delete the items using the right-click menu or the keyboard.

b. Add menus and menu items to the Appear on Toolbar panel. Select items in the Custom Menu panel and drag and drop or copy the selected items to the Appear on Toolbar panel.

   When you add a menu, a drop-down is placed on the toolbar. The toolbar drop-down contains the sub-menus and menu items shown in the Appear on Toolbar panel.

   When you add a menu item, a unique icon button for the item is placed on the toolbar.

   Menu items without icons can only be pasted in a menu. They cannot be placed on the toolbar as separate icon buttons. For example, the menu items in the Window menu do not display icons.

   The usage indicator icon beside a menu in the Custom Menu panel indicates whether the menu items for that menu are all used, partially used, or not used in the Custom Menu. Usage indicator icons are described in the Legend at the bottom of the Workspace form.

c. Add a custom menu in the Appear on Toolbar panel.
   1. Right-click in the location where you need to add a custom menu and choose Add Menu. A new menu appears under the selected item.

      If you right-click on empty space in the panel, the new menu appears as the last item in the list.

   2. Type a name for the new menu and press Enter.

d. To locate a toolbar item in the Custom Menu panel, display the menu path for the item.
   1. In the Appear on Toolbar panel, select an item, right-click, and choose Show in Custom Menu Tree.

      A dialog box displays the path for the item in the Custom Menu. If the item appears more than once in the Custom Menu tree, multiple paths are displayed.

   2. Click OK to close the dialog box.

e. To locate a menu or item in the Appear on Toolbar panel, display the menu path for the item.
   1. In the Custom Menu panel, select an item, right-click, and choose Show in Appear on Toolbar Tree.

      A dialog box displays the path for the item in the Appear on Toolbar tree. If the item appears more than once in the Appear on Toolbar tree, multiple paths are displayed.

   2. Click OK to close the dialog box.

f. Arrange the order of toolbar buttons. In the Appear on Toolbar panel, drag and drop or cut and paste selected menu items to other locations.

g. Rename menus and items in the Appear on Toolbar panel.
   1. Select the item, right-click, and choose Rename
   2. Type a name and press Enter.

   Renamed menu items are displayed in blue.

h. Add grouping separators for toolbar drop-down menus. In the Appear on Toolbar panel, right-click in the location where you need to add a separator and choose Add Separator.

   You can add separators only within menus. Separators are not available between buttons on the toolbar.

5

Save your changes and close the form.

Note: When you modify a workspace currently in use, the workspace selector displays Workspace Out of Sync. Select the current workspace from the workspace selector to apply the modified settings.

END OF STEPS

2.12 To customize tree labels

2.12.1 Purpose

In the NFM-P GUI, navigation tree objects are displayed with icons and labels. Labels are displayed as text fields separated by commas.

You can customize labels for tree objects by selecting the text field definitions that appear and the order in which they appear. Labels for the Network, Unmanaged NEs, or Discovered NEs objects cannot be customized.

Perform this procedure to customize the labels for navigation tree objects.

2.12.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

2

Click on the Trees tab.

3

Choose the view type. Object types and text fields are displayed in a table.

4

For each row in the table, perform the following as required:

1. Click on the table cell below Text Field #1 and choose a text field definition from the drop down menu.The selected text field definition appears in the table cell. On the navigation tree, the corresponding value for the text field definition will appear in the Text Field #1 position for the object.

   Some text field definitions appear as options below more than one text field heading. You can specify the display sequence of these text field definitions by choosing which text field to display them in.
2. Continue for the remaining text fields in the row.

   Note:

   Click on Restore Defaults to reset text fields to the default settings for the tree type in view.

The text field definitions for each object are displayed in navigation tree tooltips.

5

Save your changes and close the form.

Note: When you modify a workspace currently in use, the workspace selector displays Workspace Out of Sync. Select the current workspace from the workspace selector to apply the modified settings.

END OF STEPS

2.13 To customize list forms

2.13.1 Purpose

For some list forms, you can customize the object type drop-down to reduce the number of menu options displayed, to create, rename, and rearrange menu items, and to choose the default option. You can also select the action buttons that appear on the form.

2.13.2 Steps

1

Open the Workspace form; see 2.6 “To create a new custom workspace” (p. 145) or 2.7 “To modify an existing workspace” (p. 146) .

Alternatively, you can click the Customize button on a list form, if available. The Workspace (Edit) form for the current workspace opens with the List Forms tab displayed. Go to Step 4 .

2

Click on the List Forms tab and choose a list form type from the Form drop-down menu.

3

Click on the Object Types tab, if required.

The All Object Types panel contains a list of all object types available for the selected list form type. The Appear on Form panel contains the object types that appear in the object type drop-down menu.

4

Expand items as required in the All Object Types panel and Appear on Form panel.

5

Perform any of the following:

a. In the Appear on Form panel, remove any items that are not required in the object type drop-down menu. Select and delete the items using the right-click menu or the keyboard.

There must be at least one object type left in the Appear on Form panel. If all object types are removed, an error message is displayed.

b. Add a new custom folder in the Appear on Form panel.
   1. Right-click in the location where you need to add a folder and choose Add Folder. A new folder appears under the selected item.

      If you right-click on an empty space in the panel and choose Add Folder, the new folder appears as the last item in the list.
   2. Type a name for the new folder and press Enter.

      Note: If you add a new folder in the Appear on Form panel and the folder is empty, the folder is not displayed in the drop-down.

c. Rename a custom folder in the Appear on Form panel.
   1. Select the folder, right-click, and choose Rename.
   2. Type a name and press Enter.

d. Add object types to the Appear on Form panel. Select items in the All Object Types panel and drag and drop or copy the selected items to the Appear on Form panel.

   The usage indicator icon beside an object type in the All Object Types panel indicates whether the object types are all used, partially used, or not used in the Appear on Form panel. Usage indicator icons are described in the Legend at the bottom of the Workspace form. Usage indicators reduce the need to compare object types between the All Object Types panel and the Appear on Form panel.

e. To locate an object type in the All Object Types panel, display the path for the object type.
   1. In the Appear on Form panel, select an item, right-click, and choose Show in All Object Types Tree.

      A dialog box displays the path for the object type in the All Object Types Tree.
   2. Click OK to close the dialog box.

f. To locate an object type in the Appear on Form panel, display the path for the object type.
   1. In the All Object Types panel, select an object type, right-click, and choose Show in Appear on Form Tree.

      A dialog box displays the path for the object type in the Appear on Form Tree. If the object type appears more than once in the Appear on Form Tree, multiple paths are displayed.
   2. Click OK to close the dialog box.

g. Arrange the order of object types. Drag and drop or cut and paste selected menu items in the Appear on Form panel to other locations in the tree.

   The order of items in the Appear on Form panel determines their order in the drop-down menu.

h. Configure the default object type. The default object type is displayed in the drop-down when the list form opens.

   To set the default object type, select an object type in the Appear on Form panel, right-click, and choose Set Default. The selected object type is identified as the default.

To set no default for the list form, select the default object type, right-click, and choose Remove Default. When no default is set, you must select an object type each time the list form opens.

When you set a new default object type, the system overwrites any previous default.

For the Manage Services list form, the system initially assigns the Service (Service Management) object type as the default. You can change or remove this default as required.

For the Manage Equipment and Manage MPLS Objects list forms, the system does not assign a default object type until a user-configured default is set.

6

Configure the button display for the list form type selected in Step 2 . Click on the Button Panel tab and select or deselect buttons, as required. Buttons that are not selected will not appear on the list form.

The available buttons vary depending on the form type.

Mandatory buttons such as Search, Properties and Customize are dimmed and cannot be removed.

7

Save your changes and close the form.

Note: When you modify a workspace currently in use, the workspace selector displays Workspace Out of Sync. Select the current workspace from the workspace selector to apply the modified settings.

END OF STEPS

2.14 To configure the workspace selector

2.14.1 Purpose

Perform this procedure to add, remove, or change the sequence of workspaces in the workspace selector, or to set a workspace as the default workspace.

2.14.2 Steps

1

Choose Application→User Preferences from the NFM-P main menu. The User Preferences form opens.

2

Click on the Workspaces tab. A list is displayed of the workspaces that are currently available in the workspace selector.

3

Perform any of the following:

a. Add a workspace to the workspace selector.

1. Click Add. The Add Workspace form opens. A list is displayed of the workspaces not currently in the workspace selector.

   The Add button is disabled when the Allow Mandatory Workspaces Only check box is selected during user group configuration. See the procedure to add or remove workspaces for a user group in the *NSP System Administrator Guide*.

2. Choose a workspace from the list and click OK.

b. Set the default workspace that displays when the NFM-P opens. Choose a workspace in the list and click Set as Default.

c. Remove a workspace from the workspace selector. Choose a workspace in the list and click Delete.

   You cannot remove a workspace when that workspace is set as the default workspace.

   You cannot remove a workspace that is listed as a mandatory workspace in your user group.

d. Change the sequence of workspaces in the workspace selector. Choose a workspace and click Move Up or Move Down.

4

Save your changes and close the form.

Note: If the current workspace is deleted from the workspace selector, the workspace selector displays “Workspace Out of Sync”.

END OF STEPS

2.15 To apply a different workspace using the workspace selector

Note: You must add a workspace to the workspace selector before you can select it. See 2.14 “To configure the workspace selector” (p. 155) .

2.15.1 Steps

1

Click on the workspace selector drop-down in the top right corner of the GUI display.

2

Choose a workspace from the list. A confirmation dialog box appears.

a. Click Yes to close all internal and external NFM-P windows.

b. Click No to close only the navigation tree window. All other open windows are retained. The new workspace settings are not applied to the open windows.

The NFM-P GUI is displayed using the settings from the selected workspace. The workspace selector displays the workspace currently in use.

END OF STEPS

2.16 To delete a custom workspace

2.16.1 Purpose

Perform this procedure to permanently delete a custom workspace from the NFM-P. You must have the required scope of command permissions to delete a public workspace. System-defined workspaces cannot be deleted.

2.16.2 Steps

1

Choose Application→Manage Workspaces from the NFM-P main menu. The Manage Workspaces form opens.

2

Choose a workspace from the list and click Delete.

3

Save your changes and close the form.

Note: When you delete the current workspace, the workspace selector displays "Workspace Out of Sync".

END OF STEPS

2.17 To export custom workspaces

2.17.1 Purpose

Perform this procedure to export custom workspaces. You must have the required scope of command permissions to export public workspaces.

2.17.2 Steps

1

Choose Application→Manage Workspaces from the NFM-P main menu. The Manage Workspaces form opens.

2

Perform one of the following:

a. To export multiple workspaces, choose the custom workspaces in the list and click Export. The Export Directory window opens.

b. To export a single custom workspace, choose the workspace in the list and click Export. The Export Directory window opens.

   Alternatively, you can export a custom workspace by clicking on the Export button on the Workspace (Edit) form. To open the workspace (Edit) form, see 2.7 “To modify an existing workspace” (p. 146) .

Note: If you do not have permissions to export a workspace the Export button is disabled. You cannot export system-defined workspaces or private workspaces owned by other users.

3

Specify the export directory in the Save In drop-down, or create a directory or folder, and click Save. The selected workspace files are exported to the specified directory.

4

Close the Manage Workspaces form.

END OF STEPS

2.18 To import a workspace

2.18.1 Purpose

Perform this procedure to import NFM-P workspaces that were previously saved as exported workspaces. For information about how to export workspaces, see 2.17 “To export custom workspaces” (p. 157) .

2.18.2 Steps

1

Choose Application→Manage Workspaces from the NFM-P main menu. The Manage Workspaces form opens.

2

Click Import. The Import Directory window opens.

3

In the Look In drop-down, navigate to the directory that contains the workspace you need to import.

4

Configure the Overwrite User Name and Overwrite Existing Workspace(s) parameters as required, based on your permissions.

5

Click Open. A dialog box appears.

6

Click Yes. The imported workspaces are added to the Manage Workspaces form.

7

Close the Manage Workspaces form.

END OF STEPS

2.19 To add new menu items to a custom workspace of an earlier NFM-P release

2.19.1 Purpose

Later releases of the NFM-P may contain newly-introduced menu items that are not in the custom workspaces of an earlier release.

Perform this procedure to view new menu items and add them as required to an existing custom workspace of an earlier release.

2.19.2 Steps

1

Open the Workspace form; see 2.7 “To modify an existing workspace” (p. 146) .

2

Click on the All Menu Items drop-down and choose New Menu Items. A list of menu items newly-introduced in the NFM-P since the earlier release is displayed. The Finished adding New Menu Items parameter is also displayed.

Note: The New Menu Items function is available only if the workspace was created in an earlier NFM-P release.

3

Expand the items in the New Menu Items panel as required.

4

Add new menu items to the Custom Menu panel and configure as required. See 2.10 “To customize menus” (p. 148) . When you finish, if you no longer need the list of new menu items, go to Step 5 . Otherwise, go to Step 6 .

Note: When the Finished adding New Menu Items check box is left unselected, the New Menu Items panel remains as an option each time the Menus tab on the Workspace (Edit) form is opened. The usage indicator icons display the current usage of new menu items in the custom workspace.
Nokia recommends that you select the Finished adding New Menu Items check box when the New Menu Items option is no longer required.

5

Select the Finished adding New Menu Items check box.

Note: After the Finished adding New Menu Items check box is selected and you click on the Apply or OK buttons, the New Menu Items option is no longer displayed.

6

Save your changes and close the form.

END OF STEPS

3 NFM-P navigation tree

NFM-P navigation tree

3.1 Overview

3.1.1 General information

The NFM-P navigation tree provides multiple views that list the equipment groups, NEs, routing instances, protocols, and other objects in the NFM-P managed network. Objects are listed in a parent-child hierarchy. You can expand the tree to access child objects or collapse the tree for a broader view. You can access property forms, create equipment groups, and perform other functions using the navigation tree window.

For more information about network objects, see Chapter 11, “Working with network objects” .

In most system workspaces, the navigation tree window opens by default when the NFM-P client is opened. See Chapter 2, “NFM-P custom workspaces” for information about workspaces. See 1.15 “To manage a window or form as an external window” (p. 107) for information about managing the display of windows in the GUI.

You can use the navigation tree to locate specific objects in the physical and logical network views. See 3.7 “To locate objects in the navigation tree” (p. 165) for information about searching for objects in the navigation tree.

The display of objects in the navigation tree is affected by the span of control settings in user preferences; see 1.2.2 “Span of control” (p. 79) in Chapter 1, “NFM-P GUI” .

You can use the following methods to navigate the tree and manage objects.

- Double-click on an object or click on the + sign to display child objects. When you double-click on an object that has no child objects, a properties form for the object opens.
- Double-click on an object or click the - sign to hide child objects.
- Select an object and use the cursor keys to navigate the object hierarchy. The up and down arrow keys move the selection up or down in the tree. The right and left arrow keys expand or collapse objects and move up or down in the hierarchy.
- Right-click on an object to open a contextual menu, and choose an option. The menu options are specific to the object type. See 3.5 “Contextual menus” (p. 164) in this section for more information.

Note: Keyboard-based navigation tree operations may not function as expected when you open the client GUI using a third-party access tool, for example, a Citrix server.

The following figure shows the NFM-P navigation tree window, with the Equipment view selected.

Figure 3-1 NFM-P navigation tree - Equipment view

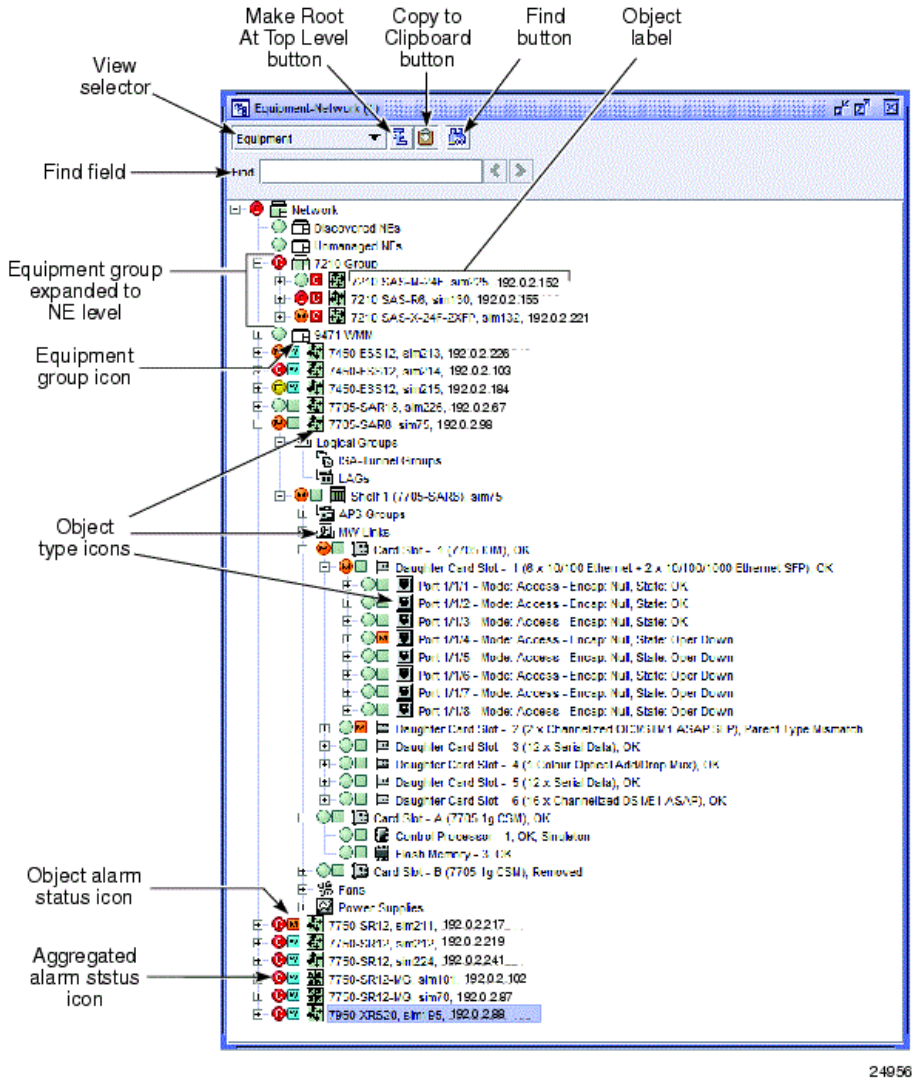

3.2 Icons and labels

3.2.1 General information

Objects in the navigation tree display on a single line with icons and labels.

Icons indicate alarm status and object type. The circle at the left indicates the aggregated alarm status. The square indicates the alarm status for the object. Object types are represented by a unique NFM-P icon for each type. Click on the Legend icon on the physical topology map toolbar for a description of alarm and equipment group icons.

Labels are displayed as text fields separated by commas, and typically provide an object description, ID number, IP address, or other information such as operational or administrative state. You can customize labels for tree objects by selecting the text field definitions that appear and the order in which they appear; see 2.12 “To customize tree labels” (p. 152) .

3.3 Equipment groups

3.3.1 General information

Equipment groups allow you to organize the network into logical groupings of NEs, for example, in a geographical area, or by equipment type. An equipment group is sometimes called a topology group. Some of the views in the view selector, for example the Equipment view, allow you to use the navigation tree to create and manage equipment groups.

An equipment group can contain up to 2000 NEs. The NFM-P client displays NEs in the navigation tree up to a limit of 500 NEs per group. An administrator can specify a system preference for the default number of NEs to display when a group is expanded on the tree (up to 500); see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*. See 3.9 “To manage NEs in equipment groups on the navigation tree” (p. 167) for information about managing NEs in equipment groups on the navigation tree.

3.4 Navigation tree toolbar

3.4.1 General information

A toolbar in the NFM-P navigation tree window contains the view selector, the Make Root At Top Level icon, the Copy to Clipboard icon, and the Find icon. The Find field is located below the navigation tree toolbar.

3.4.2 View selector

The view selector is a drop-down menu that lists the physical and logical network views available in the main navigation tree window.

You can choose the following views:

- Equipment—displays the physical objects that the NFM-P manages
- OSPF—displays all OSPF objects in the network
- ISIS—displays all IS-IS objects in the network
- Routing—displays the device routing instances and child objects, such as the network interfaces and the configured protocols
- Ring Group—displays the ring group objects that the NFM-P manages
- Cloud Network—displays DC POD or interconnect management entities

3.4.3 Make Root At Top Level icon

The root is the highest-level object in the tree. When Equipment is selected in the view selector, you can change the root of the tree to reduce the display of objects in the tree.

The Make Root At Top Level icon restores the navigation tree to the default Network root when a lower-level object is the current root.

See 3.8 “To change the root object of a navigation tree” (p. 166) for information about changing the root object of the navigation tree.

3.4.4 Copy to Clipboard icon

You can use the Copy to Clipboard icon to copy the property form identifier for one or more objects in the tree to the NFM-P clipboard. You can use the clipboard to open forms, send identifiers to other users, and configure search filters. On the topology map, you can use the clipboard to search for objects and create physical links; see Chapter 4, “Topology map management” .

See 1.20 “To use the NFM-P clipboard” (p. 112) for more information about the NFM-P clipboard,

3.4.5 Find icon and Find field

You can use the Find icon or the Find field to locate and highlight objects in the navigation tree.

See 3.7 “To locate objects in the navigation tree” (p. 165) for more information about locating objects in the tree.

3.5 Contextual menus

3.5.1 General information

When you right-click on an object in the NFM-P navigation tree, a contextual menu opens for that object. You can use contextual menus to:

- create objects
- configure object properties
- perform maintenance functions
- change the state of objects
- open a different management interface, for example, a CLI
- change the root object in the navigation tree

The available contextual menu options vary, depending on the object type.

3.6 Basic navigation tree procedures

3.6.1 General information

The following procedures describe how to use the NFM-P navigation tree.

3.7 To locate objects in the navigation tree

3.7.1 General information

You can use the Find icon or the Find field in the navigation tree header to locate objects in the navigation tree. Objects that match the search terms are highlighted sequentially in the tree. Each search can produce up to 200 results.

A system preferences threshold may limit the number of NEs displayed in the navigation tree when an equipment group is expanded. When you use the Find icon or Find field, the search results include NEs and descendant objects that were not displayed prior to the search, even if they exceed the system preferences NE display threshold.

Note: The NFM-P client displays NEs in the navigation tree up to a maximum limit of 500 NEs per equipment group. When an equipment group contains more than 500 NEs, some objects may not appear in the search results.

When the system limit of 500 NEs is already displayed under the equipment group, and a matching NE object or descendant object is not among the displayed NEs, the matching object is not displayed or highlighted.

See 3.9 “To manage NEs in equipment groups on the navigation tree” (p. 167) for more information about managing the display of NEs in equipment groups.

3.7.2 Steps

1

To locate an object using the Find icon:

1. Click the Find icon or press Ctrl-F. The search panel appears.
2. Enter search terms in the attribute fields. The available fields vary depending on the selected view.
3. Click Find in the search panel or press Enter. The navigation tree expands to show the first matching object highlighted.
4. To find additional matching objects, click the Next button or press F3. The navigation tree expands to show the next matching object highlighted in the tree.

   The Next and Previous buttons are available when more than one object matches the search term.
5. To return to the previous matching object in the navigation tree, click the Previous button or press Shift-F3.
6. To remove all search terms from the search panel, click Clear.
7. To close the search panel, click the Find icon at the top of the navigation tree window or press Ctrl-F.

2

To locate an object using the Find field:

Note: For some of the views in the view selector the Find field is not available.

1. Enter a search term in the Find field. Search terms are based on object labels. A tooltip for the Find field shows a list of the available labels. See 2.12 “To customize tree labels” (p. 152) for information about how to customize labels.
2. Press Enter. The navigation tree expands to show the first matching object highlighted.
3. To find additional matching objects, click the Next button or press F3. The navigation tree expands to show the next matching object highlighted in the tree.

   The Next and Previous buttons are enabled when more than one object matches the search term.
4. To return to the previous matching object in the navigation tree, click on the Previous button or press Shift-F3.

END OF STEPS

3.8 To change the root object of a navigation tree

3.8.1 General information

The root is the highest-level object in the tree. When Equipment is selected in the view selector, you can change the root of the tree.

Perform this procedure to make a lower-level object the root of the current navigation tree window or a new navigation tree window, or to restore the default Network root.

3.8.2 Steps

1

Open the navigation tree window and select the Equipment view.

2

Change the root object of a tree. Perform any of the following:

a. Make a selected object the root of the current tree. Right-click on the object and choose Make Root. The navigation tree is refreshed with the selected object as the root of the tree.

b. Make a selected object the root of a new navigation tree window. Right-click on the object and choose Make Root in New Tree. A new navigation tree window opens with the selected object as the root of the tree.

   Up to five navigation tree windows can be open at one time.

   The Make Root in New Tree option is also available when you right-click on an equipment group in the topology map navigation tree, or on the map background. A new navigation tree window opens with the equipment group selected or displayed in the map as the root of the tree.

   The Make Root In New Tree option is not available for the root object of a tree.

c. Restore the default Network object as the root object. Click on the Make Root At Top Level icon on the navigation tree toolbar.

The Make Root At Top Level icon is available when a lower-level object is the root of the tree.

END OF STEPS

3.9 To manage NEs in equipment groups on the navigation tree

3.9.1 General information

On the navigation tree, you can open an NEs list form for a selected equipment group and use the form to search for NEs, to show the NEs you require on the tree, and to move NEs to other groups. The form lists all of the NEs in the selected equipment group.

An administrator can specify a threshold value in system preferences for the default number of NEs to display when a group is expanded in the tree. See the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*. When an equipment group contains more NEs than the threshold value, some of the NEs are not displayed in the tree. When the threshold value is less than 500, you can use the NEs list form to temporarily display additional NEs in the group, above the threshold amount, up to 500.

An equipment group can contain up to 2000 NEs. The NFM-P client displays NEs in the navigation tree up to a maximum limit of 500 NEs per group. When an equipment group contains more than 500 NEs, some of the NEs are not displayed in the tree. You can use the NEs list form to access all of the NEs in the group, or to move NEs to other groups.

3.9.2 Steps

1

Open the NEs list form for an equipment group in the navigation tree. Right-click on the equipment group object and choose List NEs. The NEs list form opens.

When you expand an equipment group in the tree, and the number of NEs in the group equals or exceeds the threshold value set in system preferences, a message is displayed in the navigation tree under the equipment group object. You can click on the message to open the NEs list form for the equipment group.

The message remains in the tree display, even if the number of NEs in the equipment group is reduced below the system preferences threshold.

2

Manage NEs using the NEs list form. Perform any of the following:

a. Locate NEs in the tree or display additional NEs. Select an NE in the list and click Show on Tree. The selected NE is highlighted in the tree.

   If the selected NE was not previously displayed in the tree, it is added to the display for the equipment group and highlighted.

   Alternatively, you can use the Find icon or Find field in the navigation tree header to locate NEs in the tree or to display additional NEs; see 3.7 “To locate objects in the navigation tree” (p. 165) .

The display of additional NEs is not preserved. If the navigation tree window is closed or the root object of the tree is changed, the next display of the NEs in the group reverts to the default display.

Up to 500 NEs (the maximum system limit) can be displayed under an equipment group in the navigation tree. If this limit is exceeded, a message is displayed.

b. Move NEs to other groups.
   1. Select one or more NEs in the list and click Move to Group. A drop-down menu opens.
   2. Choose an equipment group in the menu. The selected NEs are moved to the selected group.

   Alternatively, you can move NEs to different groups by clicking and dragging icons on the navigation tree or topology map.

3

Close the list form.

When the NE List form for an equipment group is opened from the navigation tree, it closes automatically when you change the view in the view selector, close the navigation tree window, or change the root object of the tree.

END OF STEPS

4 Topology map management

4.1 Topology map types

4.1.1 Service tunnel topology map

A service tunnel topology map is available on the NFM-P by choosing Application→Service Tunnel Topology from the NFM-P main menu. The Service Tunnel Topology map is displayed, as shown in the following figure.

Figure 4-1 Service tunnel topology map

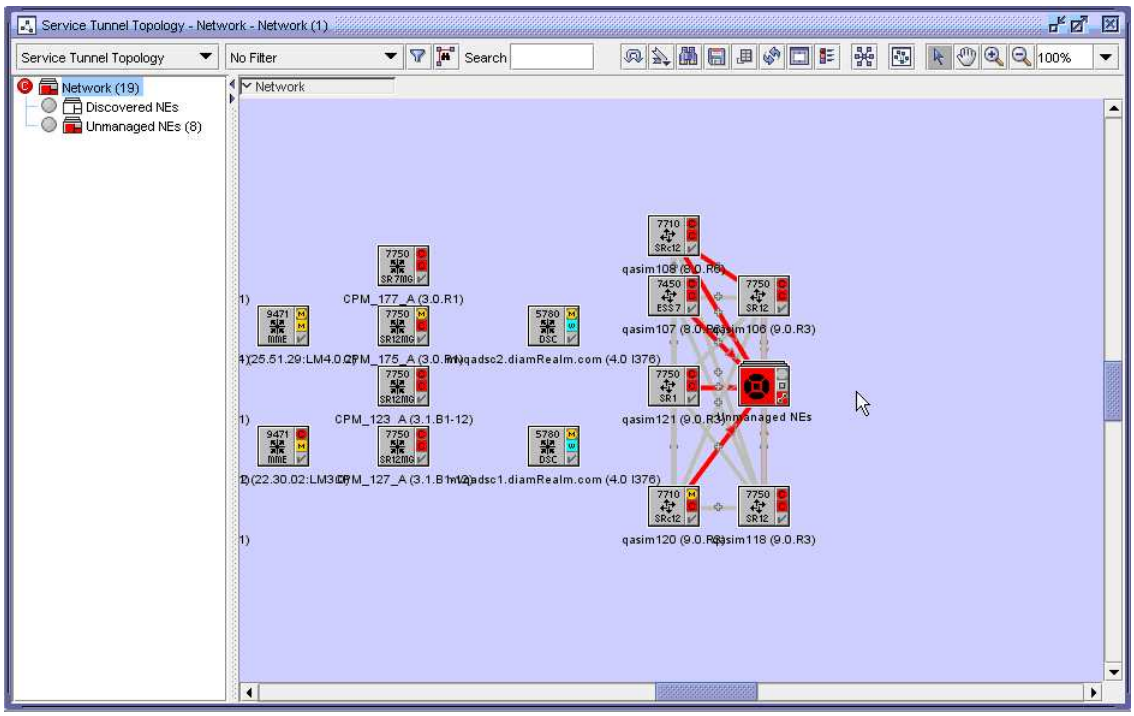

Icons in the service path topology map represent devices. The color of the device icon represents the status of the device. Red means that the device is down. Green means that the device is up. Yellow means that the device is being synchronized. Purple means that the device is in a suspended management state. See Chapter 9, “Device discovery” for information about device discovery and management states.

Link groups between devices represent service tunnels. When a link group is red, at least one tunnel in the link group is down. For link groups between managed devices, right click the link group icon to list and edit tunnels in the link group. For link groups between managed and unmanaged devices, right-click the link group icon to open contextual menus and submenus which allow you to open additional information forms for the service tunnel, including the properties form.

4.1.2 EPS path topology maps

An EPS path topology map is available on the NFM-P by choosing Application→EPS Path Topology from the NFM-P main menu. The EPS path topology map displays a static representation of mobile network objects and EPS paths. Each network object icon represents an aggregate of all network objects of that type. For SGWs and PGWs, the icon represents an aggregate of all instances of that gateway type and all network objects that contain instances of that type. Each EPS path link represents an aggregate of all EPS paths of a specific type.

4.1.3 LSP path topology map

An LSP path topology map is available from the MPLS Path form and the LSP Path form. See 4.5 “To open an MPLS provisioned path map from the MPLS Path form” (p. 179) to view the map from the MPLS Path form. See 4.6 “To open a dynamic LSP path map from the LSP Path form” (p. 179) to view the map from the LSP Path form.

The LSP path topology map is used to view a specific provisioned, actual, or CSPF LSP path in the context of its source, and transient and destination hops. The following figure shows an LSP path topology map.

Figure 4-2 LSP path topology map

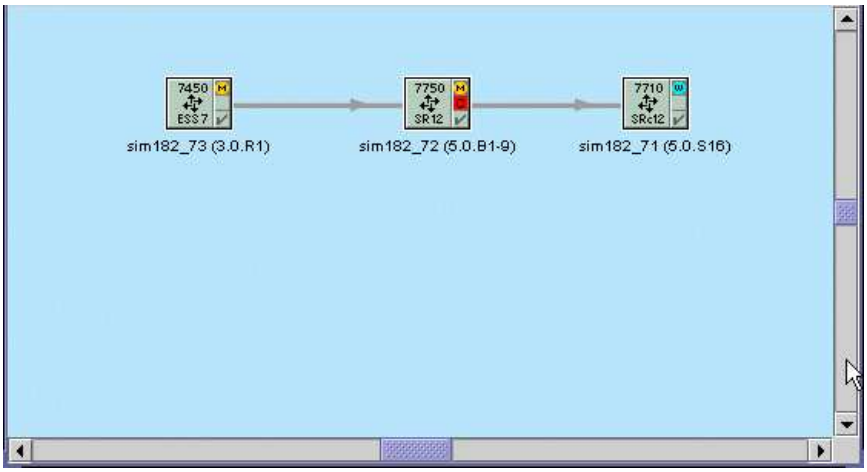

When you view the LSP path topology map, hops are linked by straight lines, where each line represents a sub-path between two hops of the LSP path. The direction of each path is indicated by an arrow. Green lines indicate provisioned paths, and gray lines indicate actual paths.

4.1.4 LSP cross-connect topology map

An LSP cross-connect topology map is available from the LSP Path form. To view the map from the LSP Path form, see 4.7 “To open a dynamic LSP cross-connect topology map” (p. 180) .

The LSP cross-connect topology map is used to view a specific LSP cross-connect in the context of its source, and transient and destination hops. The following figure shows the LSP cross-connect topology map.

Figure 4-3 LSP cross-connect map

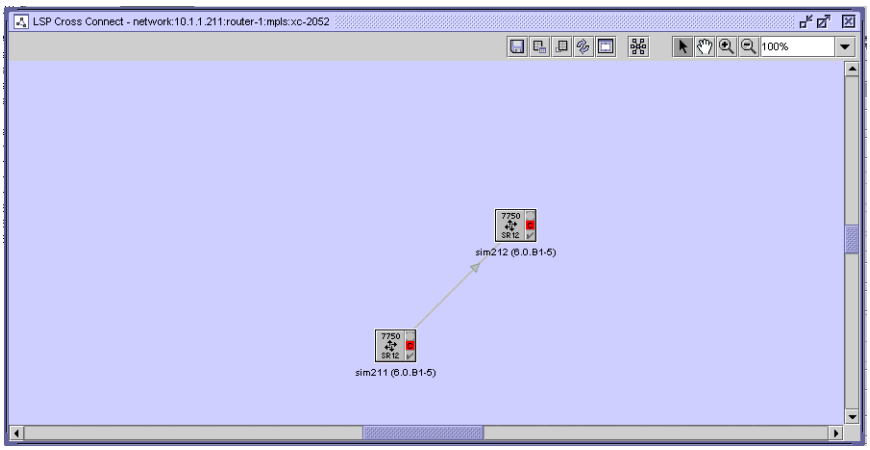

4.1.5 Service topology maps

You can open a topology map for a selected service from the Manage Services form by clicking Topology View. The following figure shows a service topology map. Large NE icons represent managed devices. Small NE icons represent unmanaged devices. The label of an NE icon indicates the service name and the NE IP address.

Figure 4-4 Service topology map

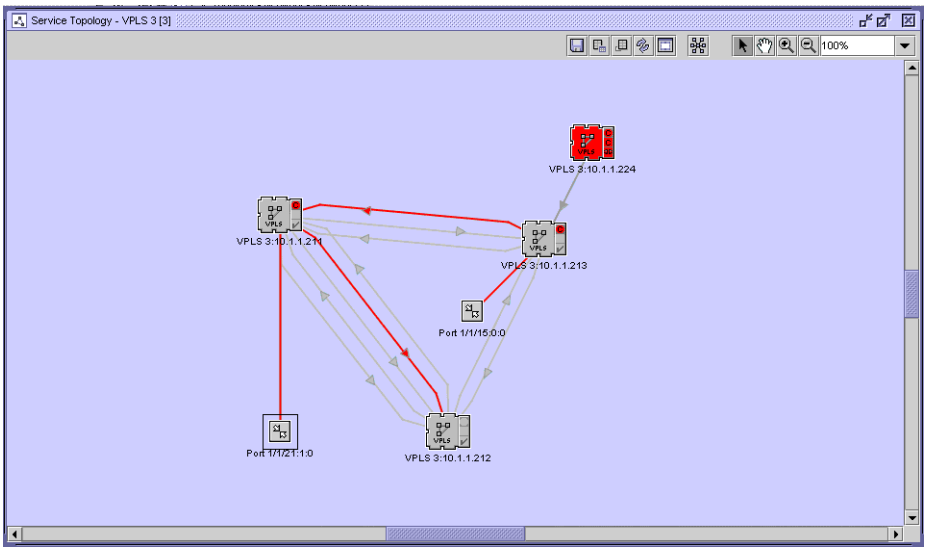

The symbol and color in the top-right corner of the managed device icon represents the aggregated alarm status, which is the most severe alarm on any service that includes the device. The symbol immediately below the aggregated alarm status and the color of the icon indicate the alarm status

for the currently displayed service site. A red icon indicates that the service site is down. A green icon indicates that the service site is up.

A port icon represents a managed access interface. A port label includes the port number and the inner and outer port encapsulation values. A service site can support up to 20 SAPs. When the number of SAPs on a service site exceeds 20, the SAPs on the site are grouped into a SAP group, represented by a SAP group icon. When the number of SAPs drops below 20, the SAP group icon is replaced with the icons for the individual SAPs. To view the list of SAPs in a SAP group, right-click on the SAP group icon and choose List L2 Access Interface or List L3 Access Interface, depending on the service type. The corresponding Site form for the service opens.

SAP aggregation group icons represent SAP aggregation groups configured for a VLL Apipe on the 7705 SAR. The SAP aggregation group is linked to the service site and each of the individual SAPs. To hide the associated SAPs for a SAP aggregation group, right-click on the SAP aggregation group icon and choose Hide SAPs. Choose Show SAPs to show them again.

A line between two map objects represents a link or group of links. Links between device icons represent service circuits. Links between device icons and port icons represent the binding of an access port or interface to a service. The symbol and color in the bottom right corner of the managed device icon represents the connectivity alarm status. During a resynchronization of the managed device, the icon represents the resynchronization status, and is yellow. The status is inherited from the link endpoints. A plus sign icon located in the centre of the link indicates a bidirectional group link. An arrow icon located on the link indicates a unidirectional link and identifies the direction of the path.

Right-clicking on a managed site, port, or link opens a contextual menu that allows you to perform functions that include the following:

- choose layout options
- open an object properties form
- turn up or shut down an object
- create a service object
- create a service object from a template, if a template is bound to the service
- delete an object
- open a CLI session
- manage scripts
- plot statistics
- display or highlight service topology features

You can view multiple services on a map at the same time if the services are selected from the Manage Services form during map creation.

Service segmentation

A service segmentation view is also available to aid in conceptualizing complex services. Segments are logical grouping of interconnected sites, services, and bindings. The segmentation view is available for VPLS and VLL services.

A service segment is considered to be a portion of a single service that extends to multiple sites connected within that segment. It is based on the service type through one of the possible

connection topologies (for example, mesh, PBB tunnels, a switching VLL, rings, and so on), without having to pass through any connectors such as spokes, CCAGs, or SCPs (SAP-to-SAP).

General examples of segments in Layer 2 service topologies include:

- A simple pair or single spoke/mesh SDP binding comprises one segment
- A mesh of a multi-NE VPLS service comprises one segment
- A multi-NE mesh with a spoke SDP to a single NE comprises two segments
- Each mesh of VSIs forms a segment (applicable to H-VPLS)

Just a few examples of the many possible service-specific segmentation configurations include scenarios such as the following:

- H-VPLS (Inter-Metro with redundant spoke SDPs):
  - One application of H-VPLS is the connection of two or more geographically-dispersed VPLS domains belonging to the same customer. Two spoke SDP connections are used to connect each VPLS between the two metros, either in a redundant PW spokes topology or under STP protocol. The redundant spokes comprise one segment, while each VPLS will also comprise one or more segments, depending on their specific configurations.
- VLL Switching:
  - For a VLL service at a switching router, a terminating PE device has at least one VLL SAP, while a switching PE device has a VLL instance which cross-connects two spoke bindings. All VLL instances of such a service must have the same service ID, and if the VLL has one or more switching sites, it must have at least two terminating sites. In this scenario, the primary and redundant spoke SDPs on the same network endpoint are considered to be in the same segment.
- PBB:
  - In a PBB configuration, the B-VPLS is considered a service tunnel, from the I-VPLS or I-Epipe perspective. Therefore, the sites connected via a B-VPLS (that is, having the same ISID) are considered to be in one segment.
  - I-Sites bound to the same PBB tunnel (B-VPLS) and having the same ISID exist in the same segment.
  - Epipe sites bound to same PBB tunnel (B-VPLS) exist in the same segment

Whenever you modify a service, the following actions can trigger segment creation, modification, or deletion in a segmented service view:

- Adding or removing service sites
- Creating or deleting spoke and/or mesh bindings (either through the NFM-P GUI or CLI)
- Creation or deletion of VLAN Uplinks by NFM-P

You can access the segmented view of a service by right-clicking on an empty portion of the background in the service's topology view. A contextual menu allows you to activate or de-activate the segmented view for the service.

In a segmented view, the outlined links of all spoke bindings, mesh bindings, or VLAN Uplinks are colored in the same distinct way for the segment they belong to. However, there are currently only 19 different colors available for use in showing service segments. If all the colors for a specific service are used, a warning message is logged to indicate this.

The following figure shows an example of how the NFM-P displays the segmented view of an H-VPLS Metro-to-Metro service.

Figure 4-5 Service segmentation example

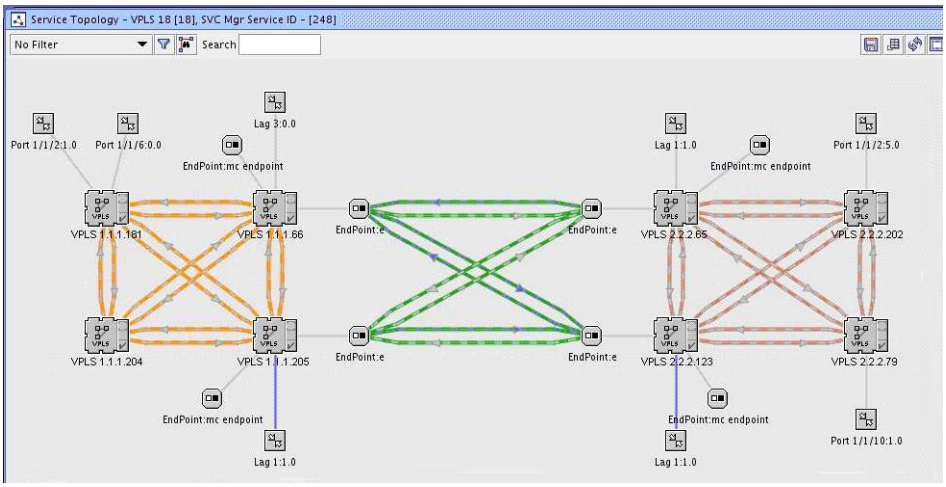

4.1.6 Composite service topology maps

You can use the NFM-P to view composite service topology and flat topology maps. The figure below shows a sample composite service topology map. When you open a flat topology view map, the navigation tree is not part of the map.

Figure 4-6 Composite service topology map

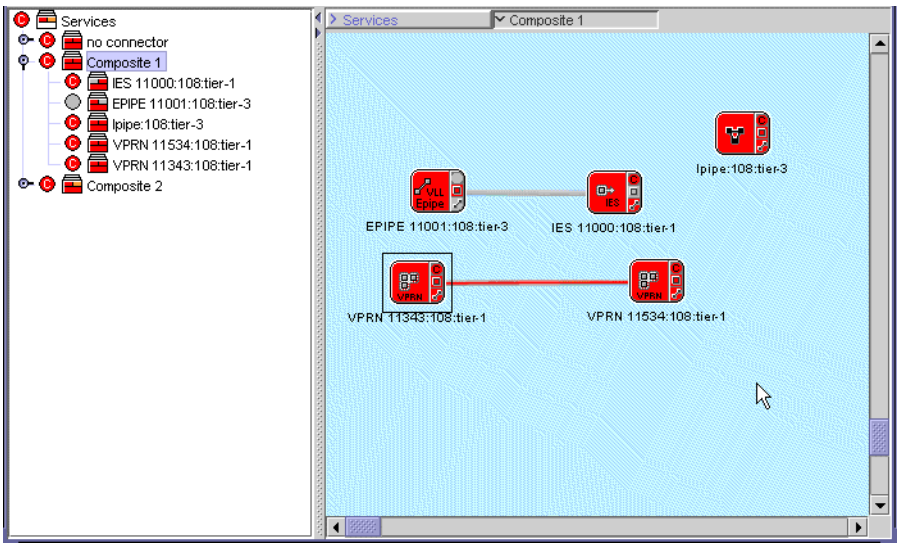

In a composite service topology map, you can use the navigation tree at the left side to display the composite service and service tier hierarchy starting from the services object. The map navigation tree displays the following:

- composite services and service tiers, including service name and tier number
- status of the composite services and their tiers

Double-click on a composite service object in the map panel to display the service objects that belong to the composite service. Double-click on the service objects to display the service sites and access interfaces. The links or groups of links between the service sites and access interfaces are also displayed.

In the composite service topology map, all service objects in the composite service are displayed simultaneously. The service sites, access interfaces, and the links or groups of links between them are also displayed. The navigation tree is thus not required in this view. The following figure shows a sample composite service flat topology map.

Figure 4-7 Composite service flat topology map

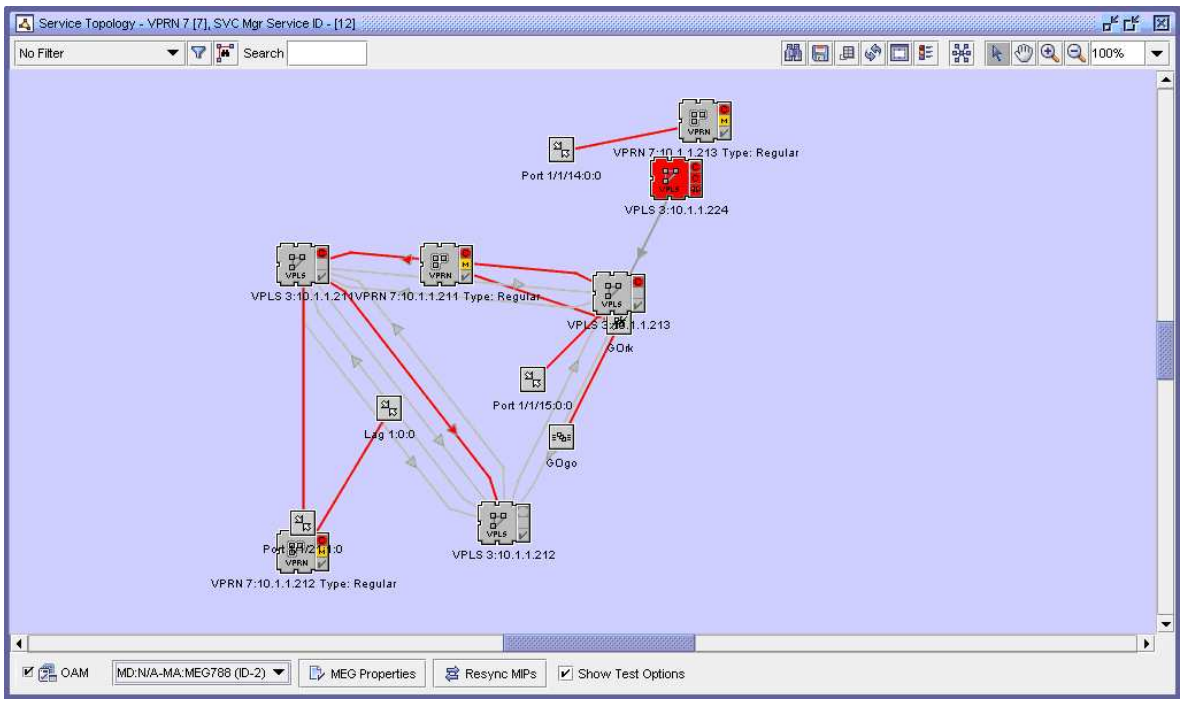

You can right-click an object icon or link group icon to turn up, shut down, or display the properties form for the item.

You can also perform Ethernet CFM diagnostics directly from the composite service flat map.

4.2 Working with topology maps

4.2.1 Modifying a service from the topology view

The topology view for a service provides a graphical representation of the various components and their interconnections. You can also use this view to add, modify, or just navigate to service components. This provides an alternative approach to performing these functions from the navigation tree view in the service configuration forms.

Working from the topology view can expedite the creation of the components, since many of the fields you would ordinarily have to set in the configuration forms will be automatically populated using this approach. The configuration forms can also be accessed directly at any time from this view by right-clicking a component. This allows quick access to conduct more detailed component configuration.

Services that can be modified from their topology views include VPLS, VPRN, VLL, IES, and composite services. See the appropriate service management chapter for procedures that describe how to modify services from the topology view.

4.2.2 Managing OAM diagnostics from the topology view

An NFM-P operator can configure and execute OAM diagnostic tests from the service topology or composite service flat topology view. A right-click contextual menu of test types and test operations allows a GUI user to perform many test-management tasks. The OAM diagnostic functions on the map are enabled and disabled using a selector in the map window.

NFM-P OAM diagnostic test management using the service topology map is supported for the following service types:

- VLAN, excluding Wavence SM VLAN
- VLL Epipe
- VPLS
- VPRN

The NFM-P OAM diagnostic test management functions are also available for composite services using the composite service flat topology map.

You can use a contextual menu option to open a test-creation form for selected map objects such as NEs, SAPs, SDPs, MEPs, and MIPs. The form is automatically populated with the information for the selected objects. After you apply the test configuration, you can use the map to execute the test and view the results. A test result summary is viewable using statically displayed or mouse-over global info tables that are shown beside the object that initiates the test. By default, the info table for a tested object displays the five most recent test results; you can enable or disable the display of individual or all results.

The general OAM contextual menu options that are available include the following:

- Create OAM Tests—lists the tests that are appropriate for the service type and selected objects; choosing a test type opens the pre-populated test-creation form
- Select OAM Tests—opens a filterable list form from which you can choose previously created tests associated with one or two objects

Note: The tests that are listed using the Select OAM Tests option do not include generated tests, or manually added first-run or last-run tests in a test suite.

See 4.8 “To use OAM diagnostic functions on service topology and composite service flat topology maps” (p. 181) for information about enabling basic OAM-related map functions.

See 90.43 “To configure and run OAM tests contextually” (p. 3052) for information about contextually configuring and executing OAM tests for entities and services on topology maps.

4.2.3 Working with Ethernet CFM objects

You can configure and manage Ethernet CFM objects such as MEGs, MEPs, and MIPs using the Ethernet CFM contextual menu option. MEPs and MIPs are represented iconically. For example, a SAP or SDP binding object can have a graphical connector to MEPs and MIPs, and a site object can be connected to a B-VPLS MEP.

Within a MEG, MEPs can be divided into logical groupings called sub-groups. A MEG which contains sub-groups is called a Global MEG. When created, a Global MEG has by default one subgroup.

A service topology map displays only one MEG at a time. The MEG is selectable using a drop-down menu in the map window when the OAM function on the map is enabled. The CFM objects associated with the currently selected MEG are displayed.

For a MEP, the following attributes are displayed:

- the MEP level, as a number inside the icon
- the MEP direction, as the up or down direction in which the icon points
- the MEP administrative state, as the icon color
- the MEG sub-group to which the MEP belongs, as the icon outline color.

Note: For composite services, MEP generation does not occur for for VLAN uplinks and overlapping service sites.

When the Show Test Options check box is enabled, the roles assigned to MEPs for Test Suite Generation are displayed beside the MEP icons. The roles are Target, Source, and Hub or Spoke. See 4.8 “To use OAM diagnostic functions on service topology and composite service flat topology maps” (p. 181) .

For a MIP, only the level is displayed.

The displayed Ethernet CFM contextual menu options for MEPs and MIPs depend on the Ethernet CFM configuration, for example, the MD level and whether MEPs and MIPs are already in the MEG.

When an operator right-clicks on the map background, the Ethernet CFM options are:

- global MEG creation, which opens a configuration form populated with the selected sites or all service sites, if none are selected.
- addition of a service to an existing global MEG
- delete selected global MEG

When the selected object is a SAP, you can create, view, change and delete MEPs. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) for information about

configuring MEPs. You can enable, disable, and view MIPs. Enabling a MIP creates a MIP only when the configuration supports this function.

When the selected object is a VLAN uplink, a solid line is displayed between the two SAPs and an OAM icon is displayed above the port icons. Right-click on the OAM icon to display the following Ethernet CFM options:

- Edit MEP
- Delete MEP
- Disable MIP
- Change MEG Sub Group

When the selected object is an SDP binding, you can enable, disable, and view the associated MEPs.

Note: When you enable a MIP on a SAP or SDP binding, the NFM-P creates a MEG site if one does not exist, and uses explicit MHF creation.

4.3 To open a map from the NFM-P main menu

4.3.1 Steps

1

Choose Application from the NFM-P main menu.

2

Choose a type of map to view from the menu options:

- Service Tunnel Topology to view service tunnels
- Flat Maps→Physical Topology
- Flat Maps→Service Tunnel Topology
- EPS Path Topology

The appropriate map opens and displays the network objects.

END OF STEPS

4.4 To open a service topology map

4.4.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select one or more services and click Topology View. A dialog box appears.

3

Click Yes to continue. The Service Topology map opens.

4

View and close the map.

END OF STEPS

4.5 To open an MPLS provisioned path map from the MPLS Path form

4.5.1 Steps

1

Choose Manage→MPLS→MPLS Paths from the NFM-P main menu. The Manage MPLS Paths form opens.

2

Select an MPLS path and click Properties. The MPLS Path configuration form for the selected path opens.

3

Click on the Provisioned Path tab. The hops of the MPLS path are listed.

4

Select a hop and click Topology View. The MPLS provisioned path map opens showing the hops between the devices.

5

View and close the map.

END OF STEPS

4.6 To open a dynamic LSP path map from the LSP Path form

4.6.1 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Select a dynamic LSP and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the LSP-Path Bindings tab. The LSP paths are listed.

4

Select an LSP path and click Properties. The LSP-Path Binding (Edit) form opens.

5

Perform one of the following:

a. View the topology for provisioned paths:
   1. Click on the Provisioned Path tab. The LSP path hops are listed.
   2. Select a hop and click Topology View. The LSP path map opens showing the LSP path hops between the devices.

b. View the topology for actual paths:
   1. Click on the Actual Path tab. The actual path hops are listed.
   2. Select a hop and click Topology View. The LSP path map opens showing the actual path hops between the devices.

c. View the topology for CSPF paths:
   1. Click on the CSPF Path tab. The CSPF path hops are listed.
   2. Select a hop click Topology View. The LSP path map opens showing the CSPF path hops between the devices.

6

View and close the map.

END OF STEPS

4.7 To open a dynamic LSP cross-connect topology map

4.7.1 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Choose an LSP and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the Cross Connects tab. The cross-connects of the LSP path are listed.

4

Choose a cross-connect and click Topology View. The LSP path map opens and displays the cross-connects between the devices.

5

View and close the map.

END OF STEPS

4.8 To use OAM diagnostic functions on service topology and composite service flat topology maps

4.8.1 Purpose

Perform this procedure to enable the OAM diagnostic functions on a service topology or composite service flat topology map, and use the basic OAM controls in the map window. See the other procedures in this chapter for information about using specific elements such as the toolbar icons and info tables.

4.8.2 Steps

1

Create a service topology map by selecting the service in the Manage Services form and clicking Topology View.

2

Select the OAM check box at the bottom left of the map window.

3

To manage one or more objects, select the objects and right-click on the set of objects to choose an option.

4

To select a MEG to display, choose the MEG from the drop-down menu of available MEGs at the bottom of the map window.

Note: You can also create a MEG directly from the topology map by right-clicking on an empty portion of the map to open a contextual menu. Select Ethernet CFM→Create Global MEG.

5

To view the properties of the selected MEG, click MEG Properties at the bottom of the window. The Global Maintenance Entity Group (Edit) form opens. The service displayed on the map is listed on the Service tab.

6

To resynchronize the MIPs in the selected MEG, click Resync MIPs at the bottom of the window.

7

To display the Test Generation Options for MEPs in the selected MEG, enable the Show Test Options check box at the bottom of the map window. See 4.2.3 “Working with Ethernet CFM objects” (p. 177) in 4.2 “Working with topology maps” (p. 176) .

8

Use the other procedures in this chapter to refine the map contents or display object properties, as required.

9

Use the right-click contextual menu options to configure and manage OAM objects, as required. See 4.2.2 “Managing OAM diagnostics from the topology view” (p. 176) for a functional description of the available contextual menu options for OAM diagnostics.

10

See 90.43 “To configure and run OAM tests contextually” (p. 3052) for information about contextually configuring and executing OAM tests for entities and services on topology maps.

END OF STEPS

4.9 To modify a service from the topology view

4.9.1 Steps

1

Perform one of the following to open a service topology view.

a. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

b. Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Click Search to list the services or composite services.

3

Perform one of the following.

a. Select a service and click Topology View. The Service Topology map opens.

b. Select a composite service and click Flat Topology View. The Composite Service Flat Topology View map opens.

4

Right-click on the map background, or on one or more selected service sites or components and choose an option. A configuration form opens.

5

See the appropriate service management chapter for information about modifying a specific service type, or Chapter 90, “OAM diagnostic tests” for information about configuring OAM diagnostic tests.

END OF STEPS

4.10 To create a physical link

4.10.1 Purpose

An operator with Device Mgmt privileges can create, modify, and delete physical links.

Note: Each physical link endpoint must terminate on the same type of MDA.
You can configure only one physical link endpoint on a port.

4.10.2 Steps

1

Choose Create→Equipment→Physical Link from the NFM-P main menu. The Physical Link (Create) form opens.

2

Configure the required parameters.

3

Perform one of the following to configure the Endpoint A parameters.

a. If you set the Endpoint A Type parameter in Step 2 to Port, go to Step 4 .

b. If you set the Endpoint A Type parameter in Step 2 to Network Element, go to Step 5 .

c. If you set the Endpoint A Type parameter in Step 2 to LAG, go to Step 6 .

d. If you set the Endpoint A Type parameter in Step 2 to Bundle, go to Step 7.

4

Configure the required parameters in the Endpoint A - Port panel by performing one of the following steps.

a. Click Select to choose a port as an endpoint.

b. Choose a port from the equipment navigation tree:
   1. Select a port and click the Copy to Clipboard icon. The port information is copied to the clipboard.
   2. Click Paste from Clipboard in the Endpoint A - Port panel of the Physical Link (Create) form.

5

Configure the required parameters in the Endpoint A - Network Element panel by performing one of the following steps.

a. Click Select to choose a managed NE as an endpoint.

b. Choose an NE from the equipment navigation tree:
   1. Select an NE and click the Copy to Clipboard icon. The NE information is copied to the clipboard.
   2. Click Paste from Clipboard in the Endpoint A - Network Element panel of the Physical Link (Create) form.

6

Configure the required parameters in the Endpoint A - LAG panel by performing one of the following steps.

a. Click Select to choose a LAG as an endpoint.

b. Choose a LAG from the equipment navigation tree:
   1. Select a LAG and click the Copy to Clipboard icon. The LAG information is copied to the clipboard.
   2. Click Paste from Clipboard in the Endpoint A - LAG panel of the Physical Link (Create) form.

7

Configure the required parameters in the Endpoint A - Bundle panel by performing one of the following steps.

a. Click Select to choose a bundle as an endpoint.

b. Click Search to populate the list of bundles:
   1. Select a bundle and click the Copy to Clipboard icon. The bundle information is copied to the clipboard.
   2. Click Paste from Clipboard in the Endpoint A - Port panel of the Physical Link (Create) form.

8

Perform one of the following to configure the Endpoint B parameters.

a. If you set the Endpoint B Type parameter in Step 2 to Port, go to Step 9 .

b. If you set the Endpoint B Type parameter in Step 2 to Network Element, go to Step 10 .

c. If you set the Endpoint B Type parameter in Step 2 to Unmanaged NE, go to Step 12 .

d. If you set the Endpoint B Type parameter in Step 2 to LAG, go to Step 13 .

e. If you set the Endpoint B Type parameter in Step 2 to Bundle, go to Step 15.

9

Configure the required parameters in the Endpoint B - Port panel by performing one of the following steps.

a. Click Select to choose a port as an endpoint.

b. Choose a port from the equipment navigation tree:

1. Select a port and click the Copy to Clipboard icon. The port information is copied to the clipboard.
2. Click Paste from Clipboard in the Endpoint B - Port panel of the Physical Link (Create) form.

10

Configure the required parameters in the Endpoint B - Network Element panel by performing one of the following steps.

a. Click Select to choose a managed NE as an endpoint.

b. Choose an NE from the equipment navigation tree:

1. Select an NE and click the Copy to Clipboard icon. The NE information is copied to the clipboard.
2. Click Paste from Clipboard in the Endpoint B - Network Element panel of the Physical Link (Create) form.

11

If the Endpoint B is an unmanaged NE that the NFM-P recognizes as an unmanaged mobile NE, go to Step 16 .

12

Configure the required parameters in the Endpoint B - Unmanaged NE panel.

Note: If the NFM-P is to manage endpoint B, configure the Unmanaged - Name parameter with the management IP address of the unmanaged NE.
When the NFM-P discovers the NE, the unmanaged endpoint of the physical link is displayed as the newly managed NE on the topology map.

13

Configure the required parameters in the Endpoint B - LAG panel by performing one of the following steps.

a. Click Select to choose a LAG as an endpoint.

b. Choose a LAG from the equipment navigation tree:

1. Select a LAG and click the Copy to Clipboard icon. The LAG information is copied to the clipboard.
2. Click Paste from Clipboard in the Endpoint B - LAG panel of the Physical Link (Create) form.

14

Go to Step 17 .

15

Configure the required parameters in the Endpoint B- Bundle panel by performing one of the following steps.

a. Click Select to choose a bundle as an endpoint.

b. Click Search to populate the list of bundles:

1. Select a bundle and click the Copy to Clipboard icon. The bundle information is copied to the clipboard.
2. Click Paste from Clipboard in the Endpoint A - Port panel of the Physical Link (Create) form.

16

Click Select in the Endpoint B - Unmanaged NE (LTE) panel to choose an unmanaged mobile NE.

17

If System Bandwidth Management is in use, configure the parameters on the Bandwidth tab as needed.

For Bundle links the bandwidth unit value is Kbps by default.

18

Click OK to save your changes and close the form. The physical link is created between the two endpoints and is displayed in the physical topology map.

19

If Service Bandwidth Management is enabled, perform a system-wide CAC audit to ensure that the bandwidth on each physical link is properly calculated.

Perform the following:

- Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
- Click CAC Audit.

END OF STEPS

4.11 To create a radio link

4.11.1 Purpose

Perform this procedure to create a radio link between NEs. You can create a link between a GNE radio interface and a managed port that has radio properties, or between two GNE radio interfaces.

You can monitor a GNE radio link endpoint when the associated GNE profile includes the appropriate GNE driver. See “GNE commissioning” (p. 246) for information about using GNE profiles. See Chapter 7, “Device management using drivers” for information about drivers.

4.11.2 Steps

1

Choose Create→Equipment→Radio Link from the NFM-P main menu. The Radio Link (Create) form opens.

2

Configure the required parameters.

3

Use the Select buttons to choose the Endpoint A and Endpoint B objects.

4

Save your changes and close the Radio Link (Create) form.

END OF STEPS

5 NFM-P-based schedules

Schedules overview

5.1 Overview

5.1.1 General information

You can create an NFM-P-based schedule for the automatic execution of tasks at a designated time as a one-time event or as an ongoing task. You can optionally specify the time at which an ongoing schedule is to stop functioning.

You can use NFM-P-based schedules for the following:

- STM test suites; see Chapter 89, “Service Test Manager”
- NE configuration backups; see Chapter 23, “NE backup and restore”
- NE checkpoint file creation and configuration rollbacks; see Chapter 24, “NE configuration rollback”
- NE software upgrades; see Chapter 26, “NE software upgrades”
- RCA service audits; see Chapter 95, “RCA audit”
- policy audits; see Chapter 49, “Policies overview”
- path optimization tasks; see Chapter 31, “MPLS”

A task such as running an STM test suite can be immediately processed, scheduled for later execution, or retained for future use. A task that is associated with a schedule is called a scheduled task. A schedule task must be created in the configuration form for the task. For example an STM test suite scheduled task must be created in the STM test suite configuration form. After scheduled tasks are created they are associated with a schedule.

5.2 Time zones and time stamps

5.2.1 General information

To simplify creating NFM-P-based schedules when the user and server are in different time zones, the NFM-P converts and displays schedule times specific to the user and server. For example, if a client user in Chicago wants to schedule a task on a server that is in New York, the NFM-P calculates the time difference and displays the local client and server times.

User time zones are configured in the user preferences form. If a time zone is not configured on the main server, the NFM-P uses the default time zone on the single-user client or client delegate server station, If the default time zone is not one of the NFM-P time zone options, the NFM-P displays the time zone ID and uses the UTC value without the time zone offset.

The NFM-P displays whether daylight saving time is in effect for the client and the server. Daylight saving time is specified for the user start time and is based on the client time zone. Daylight saving time does not specify the current time and end time.

For scheduled OAM tests, the timeCaptured attribute in a specific result is a time stamp inserted by NFM-P at the time that the test result is processed (accounting file transfer time plus test processing time). The date stamp is taken from the NFM-P server clock, with a resolution of milliseconds.

The scheduledTime attribute is a time stamp inserted by a managed NE, and corresponds to the accounting policy configuration on that NE. The current (or most recent) test result for each test is written to the accounting file with this time stamp. The actual timing is derived from the NE's clock. An accounting file can contain multiple test results for one such time stamp. The actual number of test results depends on the file policy rollover and accounting policy collection interval.

Figure 5-1 Schedule management and configuration forms

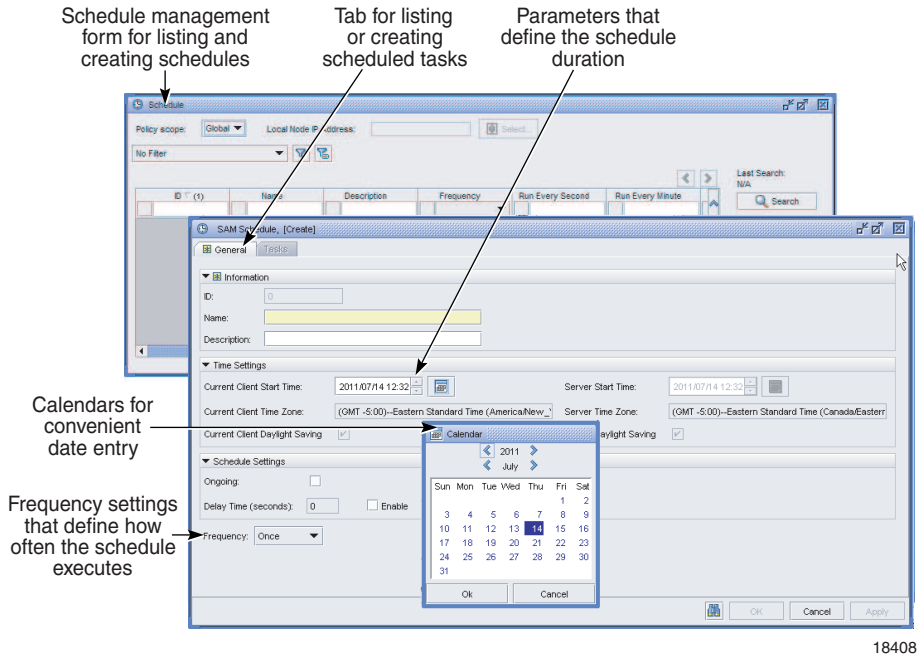

18408

5.3 NFM-P-based schedules

5.3.1 General information

An NFM-P-based schedule is only retained in the NFM-P db and is not deployed to the target NEs. When a configured schedule run time arrives, the NFM-P executes the associated task and compiles the task results for presentation to NFM-P clients.

A delay time for executing NFM-P-based scheduled runs can be configured to delay the execution of a run within a scheduled task in cases when a previously scheduled run is still executing. The NFM-P attempts to execute the new run at the configured delay time; if the run cannot execute at the delay time, the NFM-P skips the scheduled run. When a previous run completes and the next scheduled run is triggered, the run executes and is not delayed, whether a delay time is configured.

The NFM-P-based schedule functionality uses the NFM-P server time zone to trigger the scheduled tasks.

5.3.2 NFM-P-based schedules design considerations

Consider the following when you create an NFM-P-based schedule or add a scheduled task.

- Ensure that scheduled tasks are run sufficiently far apart to allow one task to finish before the next starts. Otherwise, the next occurrence of the task is skipped or delayed, if the delay time is configured.
- Do not create schedules that overlap, because no validation is performed to ensure that a newly configured schedule does not overlap with an existing schedule.
- A new scheduled task is shut down by default and must be turned up before it can be executed with one exception. Scheduled software upgrade tasks associated with all NE types are auto-enabled by default.
- Users with write access permissions to a specific schedule package can view or delete all the created schedules and scheduled tasks related to the package. Other users can only view schedules. See the section on NFM-P user security in the *NSP System Administrator Guide* for more information about user permissions.
- By default, the NFM-P associates a scheduled task with the user account that is used to create the scheduled task. You can assign a different user account to a scheduled task. The user account must have the assigned scope of command role that is appropriate for the task, or the NFM-P does not execute the scheduled task.
- NFM-P-based schedules are not distributed to the target NEs.
- You cannot delete a schedule that has a dependency, for example, one that is associated with a task. You must first delete the scheduled task from the schedule, then delete the schedule.
- One minute is added to the default start and end time values of an NFM-P-based schedule to allow time for schedule configuration.
- A monthly NFM-P-based schedule with the Run Every Month or Run Every Months parameter configured uses a 30-day interval.
- When you create a monthly NFM-P-based schedule using the Run Every parameter and specify a date that does not exist for the specified months, the last date of the month is used. For example, if you create a monthly scheduled task, starting January 31st, the scheduled task will run on February 28th, March 31st and April 30th when those months are specified in the schedule.

5.4 Workflow to create and manage NFM-P-based schedules

5.4.1 Stages

1

Create a task that uses the NFM-P-based scheduling function, such as an STM test suite. See “Schedules overview” (p. 189) in this chapter for information about which tasks can use NFM-P-based schedules.

2

Create or modify a schedule; see 5.6 “To configure an NFM-P-based schedule” (p. 193) .

3

Associate a task with the NFM-P-based schedule; see 5.7 “To associate a task with an NFM-P-based schedule” (p. 194) .

4

As required, view the scheduled tasks associated with an NFM-P-based schedule; see 5.8 “To view scheduled tasks associated with an NFM-P-based schedule” (p. 195) .

5

For a scheduled task that requires user account privileges other than the default privileges, assign a different user account to the scheduled task; see 5.9 “To assign a different user account to an NFM-P-based scheduled task” (p. 195) .

6

Turn up the scheduled task to put the schedule into a state that will execute according to the schedule; see 5.10 “To turn up or shut down an NFM-P-based scheduled task” (p. 196) . You can also use this procedure to prevent a scheduled task from executing.

7

If required, execute the scheduled task immediately; see 5.11 “To immediately execute an NFM-P-based scheduled task” (p. 197) .

NFM-P-based schedule procedures

5.5 Overview

5.5.1 General information

Use the following procedures to create and manage NFM-P-based schedules and to assign tasks to the schedules.

5.6 To configure an NFM-P-based schedule

5.6.1 General information

As required, review the 5.3.2 “NFM-P-based schedules design considerations” (p. 191) information in this chapter before starting this procedure.

5.6.2 Steps

1

Choose Tools→Schedules→Schedule from the NFM-P main menu. The Schedule form opens.

2

Click Create or choose an existing schedule and click Properties. The NFM-P Schedule (Create|Edit) form opens.

Note: You can modify a schedule only when you are logged in as the user account that is assigned to the schedule. To assign a different user account to the schedule, perform 5.9 “To assign a different user account to an NFM-P-based scheduled task” (p. 195) . You cannot modify a schedule that is in use by a scheduled task.

3

Configure the required NFM-P Schedule parameters.

Note: If you are creating a schedule to apply to policy audits, the Frequency parameter must be set to a value greater than 15 minutes.
When an NFM-P-based schedule is not Ongoing and is assigned to a task, the NFM-P raises an alarm when the Current Client End Time expires.

4

Save your changes and close the form.

END OF STEPS

5.7 To associate a task with an NFM-P-based schedule

5.7.1 Purpose

In this procedure, an STM test suite is used to illustrate the required steps to associate a task with an NFM-P-based schedule.

5.7.2 Steps

1

Open the properties form for a task that uses an NFM-P-based schedule. For example, perform 89.13 “To modify an STM test suite and view additional information” (p. 2955) to open the properties form associated with an STM test suite.

2

Click Schedule. The *task_name* Schedule Task (Create) form opens.

3

CAUTION

Service Disruption

Setting the Administrative State parameter to Enabled puts the scheduled task into effect according to the schedule parameters. Ensure that the test suite and the objects to which it applies are appropriately configured before you set the parameter to Enabled.

Contact your Nokia technical support representative for more information.

Configure the required parameters.

4

Select a schedule in the Schedule panel.

Note: The form lists only the schedules that are associated with the current NFM-P user.

5

Save your changes and close the form.

END OF STEPS

5.8 To view scheduled tasks associated with an NFM-P-based schedule

5.8.1 Steps

1

Choose Tools→Schedules→Schedule from the NFM-P main menu. The Schedule form opens.

2

Click Search and select an NFM-P-based schedule in the list and click Properties. The Schedule (Edit) form opens.

3

Click on the Tasks tab. A list of scheduled tasks associated with an NFM-P-based schedule is displayed.

4

Select a task and click Properties to view or change the scheduled task associated with the NFM-P-based schedule.

END OF STEPS

5.9 To assign a different user account to an NFM-P-based scheduled task

Note: By default, the NFM-P associates a scheduled task with the user account that is active when the scheduled task is created.

The NFM-P does not execute a scheduled task unless the scheduled task is associated with an NFM-P user account. If you delete the user account that is associated with a scheduled task, you must assign a different user account to the scheduled task. Only an NFM-P user with an assigned administrator scope of command role can assign a user account to a scheduled task.

5.9.1 Steps

1

Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task list form opens.

2

Select a scheduled task and click Reassign User. A message indicates that the new user must have the appropriate access permissions to manage the scheduled task. See the section on NFM-P user security in the *NSP System Administrator Guide* for information about scope of command roles and access permissions.

3

Click OK. The Select User form opens.

4

Select the user account that you want to reassign to the NFM-P-based scheduled task and click OK.

5

Close the form.

END OF STEPS

5.10 To turn up or shut down an NFM-P-based scheduled task

Note: A new NFM-P-based scheduled task is shut down by default and must be turned up before it can be executed at the designated time. However, scheduled software upgrade tasks associated with all NE types are auto-enabled by default.

5.10.1 Steps

1

Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task form opens.

2

Select a task entry and click Task Action.

3

Perform one of the following:

a. Turn Up—enable the scheduled task. The administrative state changes to enabled.

b. Shut Down—disable the scheduled task. The administrative state changes to disabled. The Shut Down option is enabled after the NFM-P-based scheduled task has been turned up.

4

Close the form.

END OF STEPS

5.11 To immediately execute an NFM-P-based scheduled task

5.11.1 Steps

1

Turn up the NFM-P-based scheduled task, as described in 5.10 “To turn up or shut down an NFM-P-based scheduled task” (p. 196) .

2

Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task form opens.

3

Select a task entry and click Task Action.

4

Choose Execute from the menu that appears. The NFM-P executes the NFM-P-based scheduled task.

5

Close the form.

6

Perform 5.12 “To view the current status of an NFM-P-based scheduled task” (p. 196) to view the scheduled task execution status, if required.

END OF STEPS

5.12 To view the current status of an NFM-P-based scheduled task

5.12.1 Steps

1

Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task form opens.

2

Select a task entry and click Task Action.

3

Choose View Result. The Execution Status indicator displays the current execution status of an NFM-P-based scheduled task for example, In Progress, Succeeded, or Failed.

4

Close the form.

END OF STEPS

5.13 To modify a scheduled task on an NFM-P schedule

Note: You can modify a scheduled task only when you are logged in as the user who is assigned to the scheduled task. To assign a different user account to the scheduled task, perform 5.9 “To assign a different user account to an NFM-P-based scheduled task” (p. 195) .

5.13.1 Steps

1

Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task form opens.

2

Select a scheduled task and click Properties. The Scheduled task (Edit) form opens.

3

Modify the parameters, as required.

4

Click OK. The NFM-P Scheduled Task (Edit) form closes and the Scheduled Task form reappears.

5

Close the form.

END OF STEPS

Part II: Device management

Overview

Purpose

This part provides information about using the NFM-P for device management.

Contents

6 Device support

6.1 Device support overview

6.1.1 Device list

The NFM-P can manage devices in the following device families:

- 6.3 “1830 VWM” (p. 207)
- 6.4 “210 WBX” (p. 215)
- 6.5 “7210 SAS” (p. 216)
- 6.6 “7250 IXR” (p. 220)
- 6.7 “7450 ESS” (p. 221)
- 6.8 “7705 SAR” (p. 221)
- 6.9 “7750 SR” (p. 226)
- 6.10 “7850 VSG/VSA” (p. 227)
- 6.11 “7950 XRS” (p. 227)
- 6.13 “OmniSwitch” (p. 230)
- 6.14 “ Wavence SM and Wavence SA” (p. 231)

See the *NSP NFM-P Network Element Compatibility Guide* for specific device support information.

The NFM-P also provides limited management support for 6.12 “Generic NEs” (p. 229), which are typically devices from other vendors.

By default, for most core-network device types, the NFM-P supports the current software release and a limited number of immediately preceding major releases.

Note: If the NFM-P license includes the Extended NE Software Support option, additional device releases are supported; contact Nokia for information.

The NFM-P GUI displays a device feature or function only if the device supports the feature or function. The NFM-P user documentation describes the major variations in support, including device-specific procedures for specific implementations, but does not describe the following minor GUI variations for different devices:

- partial support for the parameter list on a form
- partial support for the buttons and tabs on a form

The NFM-P supports an NSP mediation option for the management of NEs with CLI and/or NETCONF/YANG management interfaces. MDM provides mediation between the NSP and Nokia or third-party NEs. The NSP can be used in the same way with any NE, regardless of its management operational mode. For example, you can use Service Management to create a service between NEs from different vendors that are managed using different mediation protocols.

See the *NSP Installation and Upgrade Guide* for more information.

For Release 16.0 R1 and later SROS devices, you can view the Management Operational Mode parameter on the NE object properties form:

- Classic—no model-driven management
- ModelDriven—management via model-driven interfaces only
- Mixed—management, with restrictions, by both classic and model-driven interfaces; also called dual management

The Management Operational Mode parameter is read-only in the NFM-P.

6.1.2 Device-specific release information

CAUTION

Service Disruption

When a device introduces a maintenance release, you must ensure that your NFM-P software supports the new device release before you upgrade the device to the new release.

You may need to upgrade the NFM-P to support the new device release. See the NSP NFM-P Network Element Compatibility Guide for more information.

For information about NFM-P support for specific device releases and functions, see the *NSP NFM-P Network Element Compatibility Guide*.

See the device user documentation for information about device functions, parameters, and CLI commands that are outside the scope of the NFM-P documentation. Contact Nokia technical support for specific network or facility considerations.

This chapter describes NFM-P managed devices and notable information about cards, however it does not include a comprehensive list of all supported cards on the devices. See the NE documentation for hardware specifics that are outside the scope of the NFM-P documentation.

6.1.3 Device-specific guides

The *NSP NFM-P Classic Management User Guide* describes the management of the majority of Nokia devices. However, for some devices, detailed NFM-P management is described in other guides. For the devices covered by the *NSP NFM-P Classic Management User Guide*, the sample workflow in this chapter describes the high-level steps that are required to configure and manage the device.

6.2 Sample workflow to configure and manage devices

6.2.1 Overview

The following is a generic workflow of the high-level tasks that are typically used to configure and manage supported devices using the NFM-P. This workflow is common to all NFM-P-managed devices but not all tasks apply to all device types. As appropriate, review the workflow associated with each task for detailed instructions.

Note: You can use an OSS client to configure many of the functions available on the NFM-P GUI. See “Communicating with the NFM-P” in the *NSP NFM-P XML API Developer Guide* for information.

6.2.2 Stages

Prerequisite tasks

1

Plan your NFM-P deployment for managing devices by determining the following:

- the number of NEs the NFM-P is to manage, the redundancy requirements and the hardware required for the system
- the management network latency and management network bandwidth requirements
- the naming conventions for objects that you create

See the *NSP Planning Guide* for the full list of deployment considerations.

2

Integrate the NFM-P with other EMS, as required.

3

Review the *NSP NFM-P Network Element Compatibility Guide* for release-specific information about the compatibility of managed device features in NFM-P releases. Install the physical device as per the appropriate device-specific hardware user documentation.

Review GUI basics for managing devices

4

Familiarize yourself with GUI operations for configuring and managing devices such as the following:

- navigating the GUI, performing searches, and customizing the GUI user preferences; see Chapter 1, “NFM-P GUI” .
- creating custom workspaces to simplify navigation and operation; see Chapter 2, “NFM-P custom workspaces” .
- managing devices using the equipment navigation tree; see Chapter 3, “NFM-P navigation tree” .
- managing equipment using topology maps; see Chapter 4, “Topology map management” .
- creating schedules to allow for the automatic execution of tasks at designated times; see Chapter 5, “NFM-P-based schedules” .

5

Launch the on-product user documentation to access the NFM-P customer documentation and search tools.

Perform account and security tasks

6

Set up all required NFM-P user accounts and user groups with the required scope of command roles, span of control permissions, and the ongoing monitoring and management of those accounts. See the section on NFM-P user security in the *NSP System Administrator Guide* for more information.

7

Perform the required NFM-P and UNIX security management tasks for accessing and securing the managed device. See the sections on NE user and device security and TCP enhanced authentication in the *NSP System Administrator Guide* for more information.

Prepare network devices for NFM-P management

8

Commission Nokia devices for NFM-P management and ensure all non-Nokia devices have a GNE profile. Configure how and when the NFM-P polls the devices for MIB changes; see Chapter 8, “Device commissioning and management” .

9

Create a mediation security policy and discovery rule to define the criteria required to add devices to the managed network, and optionally configure SSH2 security for CLI sessions; see Chapter 9, “Device discovery” .

10

Use a Telnet or SSH CLI from the NFM-P for device access; see Chapter 10, “Device CLI sessions” .

11

Customize the global system settings, as required. See the section on NFM-P component configuration in the *NSP System Administrator Guide*.

Configure and manage the discovered device

12

Configure and manage equipment navigation tree objects including:

- general network objects; see Chapter 11, “Working with network objects”
- device objects; see Chapter 12, “Device object configuration”
- logical group objects; see Chapter 13, “Logical group object configuration”
- ESA objects; see Chapter 14, “ESA object configuration”
- shelf, card, and daughter card objects; see Chapter 15, “Shelf and card object configuration”
- port and channel objects; see Chapter 16, “Port and channel object configuration”

13

Perform additional device management functions, as required such as the following:

- equipment inventories of managed devices; see Chapter 17, “Inventory management”
- card migrations; see Chapter 18, “Card migration”
- configure threshold-crossing alarms; see Chapter 19, “TCA”
- bulk configuration changes; see Chapter 20, “Bulk operations”

Manage the network functions

14

Configure NE routing and forwarding on NEs; see Chapter 27, “NE routing and forwarding” . As required, configure specific routing and forwarding functions and services such as:

- NE protocols; see Chapter 28, “Routing protocol configuration”
- OpenFlow; see Chapter 29, “OpenFlow”
- IS-NAT groups; see Chapter 30, “NAT”
- MPLS and LSPs; see Chapter 31, “MPLS”
- MPLS-TP; see Chapter 32, “MPLS-TP”
- service tunnels to carry service traffic; see Chapter 33, “Service tunnels”
- IPsec sessions; see Chapter 34, “IPsec”
- ISA-Video groups and members, see Chapter 35, “ISA-Video”

15

Configure network redundancy, protection, and offloading as required to ensure network availability in case of a network device or path failure using functions such as:

- VRRP; see Chapter 37, “VRRP”
- APS; see Chapter 38, “APS”
- Wi-Fi offloading; see Chapter 39, “lightRadio Wi-Fi”
- MC peer groups; see Chapter 40, “MC peer groups”
- MC IPsec; see Chapter 41, “MC IPsec”
- MC endpoint groups; see Chapter 42, “MC endpoint groups”
- MC LAG groups; see Chapter 43, “MC LAG groups”
- MC synchronization groups; see Chapter 44, “MC synchronization groups”
- MC ring groups; see Chapter 45, “MC ring groups”

Create policies to manage devices

16

Depending on your service delivery model, create, distribute, and assign the appropriate service or routing policies that define the conditions for managing network devices using the NFM-P; see Chapter 49, “Policies overview” for general information about policy management.

See Chapter 50, “QoS policies” to Chapter 68, “7705 SAR Security policies” for specific policy information and the associated workflows to configure the policy.

Create services over devices

17

Familiarize yourself with service management and QoS concepts prior to configuring services over the devices; see Chapter 70, “Service management and QoS” . Configure the following adjunct service delivery methods to support the creation of network services as required:

- shared queue groups to allow SAP or IP interface forwarding classes to be redirected; see Chapter 71, “Queue groups” .
- virtual ports to support clustering and enforce the aggregate rate of each destination BSAN; see Chapter 72, “Virtual ports” .
- create and manage customers and the services they subscribe to; see Chapter 73, “Customer configuration and service management” .
- tunnel administrative groups to enable services to use a PW template; see Chapter 88, “Tunnel administrative groups” .

18

Configure the following network services as required:

- services for residential subscribers; see Chapter 74, “Residential subscriber management” .
- VLAN services; see Chapter 75, “VLAN service management” .
- VLL services; see Chapter 76, “VLL service management” .
- VPLS; see Chapter 77, “VPLS management ” .
- IES; see Chapter 78, “IES management” .
- VPRN services; see Chapter 79, “VPRN service management” .
- SPB services; see Chapter 80, “SPB service management” .
- PW routing and dynamic MS-PW services; see Chapter 81, “PW routing and dynamic MS-PW service management” and Chapter 83, “Service PW template policies” .
- Composite services; see Chapter 85, “Composite service management” .
- Dynamic services; see Chapter 86, “Dynamic service management” .

Monitor, maintain, and troubleshoot devices

19

Configure alarm policies, and monitor incoming alarms to check the type and characteristics of the alarms, and to resolve the network problems or physical equipment failures identified by the alarms; see Chapter 36, “Alarm management” .

20

Perform the following service assurance functions as required, to ensure that services offered over network devices meet the pre-defined service quality level expected by customers:

- use the NFM-P Service Test Manager to group OAM diagnostic tests into suites of manual and automatically-generated tests; see Chapter 89, “Service Test Manager” .
- use OAM diagnostic tools to troubleshoot network problems and for SLA verification; see Chapter 90, “OAM diagnostic tests” .
- configure Ethernet CFM to detect, isolate, and report connectivity faults in Ethernet networks; see Chapter 91, “Ethernet CFM” .
- create mirror services to troubleshoot customer traffic issues or for use with Lawful Intercept; see Chapter 93, “Mirror services” and Chapter 94, “Lawful Intercept” .
- test and report network delay and loss using CFM messaging in Layer 2 networks; see Chapter 92, “Performance Monitoring tests” .
- perform on-demand verifications of the configuration of services and physical links using an RCA audit; see Chapter 95, “RCA audit” .
- create a copy of the SAP configuration of an end-user service to emulate the bandwidth, throughput, and QoS requirements of a service; see Chapter 96, “Service throughput configuration” .
- configure Application Assurance to provide deep-packet inspection and NSP subscriber traffic management; see Chapter 87, “Application assurance” .

21

Collect NFM-P and NE statistics to monitor network and service performance, compile equipment usage and billing data, and ensure SLA compliance; see the *NSP NFM-P Statistics Management Guide*.

22

Perform device maintenance functions, as required, for example:

- configuration backups and restores; see Chapter 23, “NE backup and restore”
- checkpoint file creation and configuration rollbacks; see Chapter 24, “NE configuration rollback”
- deployment management; see Chapter 25, “NE deployment”
- software upgrades; see Chapter 26, “NE software upgrades”

23

Identify and resolve performance issues in an NFM-P-managed network or on an NFM-P system as required.

6.3 1830 VWM

6.3.1 Overview

The 1830 VWM constitutes the backbone of the wireless fronthaul solution by establishing an end-to-end transport line between RRHs and the centralized BBUs.

At the cell site: The 1830 VWM TLU provides the colorization of the black and white signals and the 1830 VWM PMU provides the add-drop functionality. The 1830 VWM ITP is an integrated TLU and PMU device that provides wavelength translation and passive wavelength division multiplexing capabilities.

At the central office: The 1830 VWM TLUs and 1830 VWM PMUs are connected to an 1830 VWM OSU to form the central management hub for OSC signals transported using the out-of-band channels from the peripheries. The 1830 VWM OSU provides the management interface to the NFM-P.

6.3.2 Supported functions

The NFM-P provides the following FCAPS functions for 1830 VWM devices, as listed in Table 6-1, “NFM-P support for 1830 VWM devices” (p. 207).

Table 6-1 NFM-P support for 1830 VWM devices

Function	Support	References
NE discovery	1830 VWM device discovery Enable SNMPv3 management Configure device mediation Configure a discovery rule	6.3.3 “1830 VWM discovery” (p. 209) 9.12 “To enable SNMPv3 management and discover an 1830 VWM device” (p. 294)
	IPv6 device management	9.1.2 “Device discovery using IPv6” (p. 277)
	CLI session	10.4 “To open and close an NFM-P device CLI session” (p. 329)
Device object	Shelf craft port configuration	“Shelf craft port” (p. 210) 12.64 “To configure shelf craft port IP address on an 1830 VWM device” (p. 394)

Table 6-1 NFM-P support for 1830 VWM devices (continued)

Function	Support	References
Shelf, card, and port objects	Shelf configuration Restart shelf Lamp test	15.12 “To configure an 1830 VWM shelf” (p. 470) “Warm and cold restart” (p. 210) 15.51 “To restart an 1830 VWM shelf” (p. 512) “Lamp test” (p. 210)
	Card object configuration 1830 VWM OPS – OSM protection switching	6.3.4 “Shelf, card, and port objects” (p. 210) 6.3.12 “1830 VWM OPS – OSM protection switching and bi-directional protection switching” (p. 214)
	Port object configuration RFLM port configuration CPRI monitoring using the 1830 VWM TLU 9M MON ports CPRI rate configuration for a CDR channel CDR channels Mapping of ports and CDR channels associated with it. SFP frequency	16.21 “To configure 1830 VWM ports” (p. 596) 16.3 “Remote fiber link monitoring in 1830 VWM devices” (p. 568) 16.78 “To configure an OSC port of an 1830 VWM OSU as a RFLM port” (p. 680) “CPRI channel monitoring” (p. 211) 16.79 “To perform CPRI monitoring using 1830 VWM TLU 9M MON ports” (p. 681) 15.75 “To configure a CPRI rate and channel management for a CDR channel on an 1830 VWM TLU or ITP card slot” (p. 532) 16.76 “To view the channels associated with a 1830 VWM TLU port” (p. 679) “CDR channels” (p. 211) “SFP frequency” (p. 211)
Link management	Equipment connection topology	6.3.8 “Equipment connection topology” (p. 212)
Alarm management	Alarm severity configuration	See Procedure “To configure alarm severity and deletion behavior” in the *NSP System Administrator Guide*
	SNMP traps	6.3.5 “SNMP traps” (p. 212)
	SNMP trap destination configuration	*1830 VWM node documentation*
	SNMP trap restoration	6.3.6 “SNMP trap restoration” (p. 212)
Service test manager	PRBS test	“PRBS test” (p. 3000) 90.41 “To create and run a PRBS test” (p. 3048)
NTP support	Configure NTP	15.35 “To configure NTP on 1830 VWM OSU devices” (p. 494)
Backup and restore	NE backup and restore	6.3.9 “Backup and restore considerations for 1830 VWM” (p. 213) Chapter 23, “NE backup and restore”
Software upgrade	1830 VWM software upgrade Image software database	6.3.10 “Software upgrade” (p. 213) 26.5 “To configure a software upgrade policy” (p. 776) 26.24 “To perform an 1830 VWM on-demand software upgrade” (p. 812) “Image software database” (p. 213)

Table 6-1 NFM-P support for 1830 VWM devices (continued)

Function	Support	References
Statistics management	Statistics management	6.3.13 “Statistics management” (p. 215) See the procedure “To assign the default 1830 VWM OSU performance management policy to 1830 VWM devices” in the *NSP NFM-P Statistics Management Guide*
	Counters	*NSP NFM-P Statistics Management Guide*
	DDM data retrieval	“DDM data retrieval” (p. 215) 16.2.2 “1830 VWM − DDM data retrieval” (p. 568) 16.77 “To retrieve 1830 VWM DDM data” (p. 680)
Inventory management	Inventory management	6.3.11 “1830 VWM inventory management” (p. 213) Chapter 17, “Inventory management”
License management	View the NFM-P license information	See the section on software and license configuration procedures in the *NSP System Administrator Guide* for information about creating a license point inventory.

6.3.3 1830 VWM discovery

The 1830 VWM discovery rule requires the SNMPv3 read, write, and security access mediation policies and SNMPv2 trap access mediation policy. The NFM-P supports 1830 VWM device management using IPv4 and IPv6. See 9.12 “To enable SNMPv3 management and discover an 1830 VWM device” (p. 294) for more information about 1830 VWM SNMPv3 management, device mediation, and discovery.

6.3.4 Shelf, card, and port objects

Shelf objects

NFM-P supports the configuration of the 1830 VWM shelves. See 15.12 “To configure an 1830 VWM shelf” (p. 470) for more information about 1830 VWM shelf configuration. The 1830 VWM OSU shelf components are sorted on the equipment tree, based on the alphabetical order of the shelf name and not based on the shelf number.

Warm and cold restart

You can perform either a warm or cold restart of the 1830 VWM OSU using the NFM-P. A warm restart restarts the software without unnecessary reconfiguration of the hardware. Therefore the warm restart operation does not affect data transmission. The cold restart restarts the software and reconfigures all hardware from the internal database. See 15.51 “To restart an 1830 VWM shelf” (p. 512) for more information about performing a warm or cold restart.

Lamp test

The RMUs are automatically assigned a shelf ID based on detection by the 1830 VWM OSU. When you look at a stack of installed units in a rack, it is not obvious which shelf ID belongs to which physical unit. You can perform a lamp test by configuring the Lamp Test parameter as Active in the Shelf Specifics tab of the Shelf (Edit) form. The LEDs on the selected unit flash, enabling quick identification of a particular unit in a large array.

Shelf craft port

NFM-P supports the configuration of the craft port parameters at the NE level or shelf level.

On the equipment tree, right-click on the 1830 VWM device object and choose Properties. The Network Element (Edit) form opens. The Shelf Craft Port tab in the Network Element (Edit) form lists the craft port parameters of all the shelves except the OPS shelf. You can configure the craft port IP address, craft port gateway, and the craft port IP prefix length.

The Shelf Craft Port tab in the Shelf (Edit) form displays craft port parameters of the specific shelf in the Shelf Craft IP Details panel. You can configure the craft port IP address, craft port gateway, and the craft port IP prefix length. The Shelf Craft Port tab is not available for the OPS shelf.

See 12.64 “To configure shelf craft port IP address on an 1830 VWM device” (p. 394) for more information about configuring the craft port parameters.

Card objects

The card objects are configured automatically when the 1830 VWM shelf is configured. To view the card properties, you can right-click on the Network→1830 VWM-OSU NE→Shelf→Card Slot object and choose Properties. The Card Slot (Edit) form opens.

Port objects

The port objects are configured automatically when the 1830 VWM shelf is configured. See 16.21 “To configure 1830 VWM ports” (p. 596) for more information about port configuration.

The operational state of the ports can change due to traffic impact. For example, fiber pull and so on. Because of the missing traps, the NFM-P cannot automatically update the state change values from the network. You need to perform manual resynchronization of the ports to get the real time value.

Port label – SNMP and CLI

NFM-P displays both the Name and CLI Name parameters to identify the ports in the General tab of the Physical Port (Edit) form. The Name parameter displays the value derived from the SNMP; for example, 3/1/15 where 15 is the port number. The CLI Name parameter displays the value in the alphanumeric format derived from the CLI; for example, 3/1/C8 where C8 is the port number.

NFM-P displays the CLI Name parameter to identify the ports in the Info tab of the Alarm Info form. The Alarmed Object Name parameter value appears in the alphanumeric format derived from the CLI and identifies the port object.

SFP frequency

NFM-P supports the configuration of SFP transmission frequency on the line and client ports of the TLU 9 and TLU 9M cards. You can configure the Sfp transmission Frequency parameter on the Sfp Details panel in the Sfp Specifics tab of the Physical Port (Edit) form.

CPRI channel monitoring

NFM-P supports the following CPRI monitoring function using the 1830 VWM TLU 9M MON ports:

- in-service monitoring of interference on a specific CPRI channel

- in-service monitoring of a downlink signal on a specific CPRI channel
- allows injection and extraction of C and M traffic to commission RRH
- perform loopback tests on RRH and monitor round-trip integrity

See 16.79 “To perform CPRI monitoring using 1830 VWM TLU 9M MON ports” (p. 681) for more information about performing the CPRI monitoring function.

CDR channels

The CPRI interface supports various line rates. When you configure a CPRI rate, the 1830 VWM treats the corresponding client and line ports (C1 and L1, C2 and L2, and so on) as a CDR Channel. A CDR channel is a traffic connection between two ports on the first slot of a TLU. The TLU slot has nine CDR channels. The ports and the CDR channel mapping is preconfigured on the 1830 VWM TLU device and you cannot change the mapping.

You can configure a CPRI rate for a CDR channel using the NFM-P. See 15.75 “To configure a CPRI rate and channel management for a CDR channel on an 1830 VWM TLU or ITP card slot” (p. 532) for more information about configuring a CPRI rate for a CDR channel. The 1830 VWM TLU ports show the mapping of the ports and the CDR channels associated with it. See 16.76 “To view the channels associated with a 1830 VWM TLU port” (p. 679) for more information about viewing the channel-to-port mapping.

6.3.5 SNMP traps

The NFM-P supports traps associated with alarm management and the traps related to configuration changes on the 1830 VWM Release 8.3 and later. You can perform the configuration changes for the 1830 VWM devices using the NFM-P or CLI and the data is synchronized.

6.3.6 SNMP trap restoration

The NFM-P retrieves the last trap sequence number sent from all network elements at a configurable interval. This interval is configurable on a per resource group basis. Resource groups allow the user to configure the communications behavior of a group of network elements. By default, the core resource group includes all network elements, and verifies the trap sequence number every 4 minutes. The NFM-P compares that sequence number with the sequence number of the last trap it received from that network element. If they do not match, NFM-P will request only the missing traps from the network element. If NFM-P is missing more than 200 traps from a network element, or if the network element no longer has the missed trap, the NFM-P will request a full resynchronization on that network element rather than just request the missing traps. This behavior occurs by default and is not configurable. See the *NSP Planning Guide* for more information.

6.3.7 Synchronization of operational state

NFM-P updates the operational state dynamically to display the Operational State parameter as Down when LOS or LOF alarms are generated from the 1830 VWM device and Up when the LOS or LOF alarms are cleared from the 1830 VWM device.

6.3.8 Equipment connection topology

The equipment connection topology provides internal connections between two ports within a shelf or external connections between two ports in two different shelves.

There are two types of equipment connection topologies:

- Layer 1 connections
- Layer 2 connections

Note: The far end address is not displayed because it is applicable to the far end type, External. You cannot select the port which is already participating in any of the L1 or L2 topology links. The used ports are not listed in the selection window of the far end port.

Layer 1 connections

Layer 1 connections are the physical connections of fibers or LAN cables. NFM-P supports the configuration of layer 1 connections.

Layer 2 connections

Layer 2 connections are derived from the Ethernet RSTP attributes. NFM-P supports the discovery of layer 2 connections.

The characteristics of layer 2 connections are:

- Although the NFM-P allows the configuration and discovery of misconfigured links, the validity of the links is not guaranteed.
- If incomplete links are discovered as part of 1830 VWM device discovery, and the links are listed in the Physical Links→VwmOSU Links tab of the Network Element (Edit) form, delete the link and re-configure the link with valid interfaces. The incomplete links are discovered to identify erroneous links and rectify them.

Device-level physical topology map

NFM-P supports device-level physical topology map at the 1830 VWM device-level that displays the 1830 VWM shelves, the internal links within the shelves, and the external links between the shelves.

You can use one of the following options to open the device-level physical topology map:

- Right-click on the 1830 VWM device object on the navigation tree and choose Topology View from the contextual menu.
- Click Topology View on the Network Element (Edit) form.

6.3.9 Backup and restore considerations for 1830 VWM

- Configuration save is not supported.
- Only SFTP protocol is supported.

6.3.10 Software upgrade

NFM-P supports in-service software upgrades (ISSU) of the 1830 VWM.

See 26.5 “To configure a software upgrade policy” (p. 776) for more information about how to configure a software upgrade policy.

See 26.24 “To perform an 1830 VWM on-demand software upgrade” (p. 812) for more information about how to perform an 1830 VWM on-demand software upgrade.

Image software database

NFM-P lists the respective active and inactive 1830 VWM software images in the Shelf ISD tab of the OSU, PMU, TLU, and ITP Shelf (Edit) forms. If the ISD status is active, it is the software load currently active on the system. If the ISD status is inactive, it is the previous software load.

6.3.11 1830 VWM inventory management

You can list the 1830 VWM card inventory information using the VWM MS Card object. You can list the 1830 VWM CDR channel inventory information using the CDR Channel object. See Chapter 17, “Inventory management” for more information about inventory management.

6.3.12 1830 VWM OPS − OSM protection switching and bi-directional protection switching

The OPS OSM card of the 1830 VWM OPS shelf provides optical protection switching and the OPS OSM-DSV card provides bi-directional protection switching for both Tx and Rx. You can configure up to four OPS OSM cards or OPS OSM-DSV cards, or a mix of both on an 1830 VWM OPS shelf. The OPS OSM-DSV card is intended for use in a C-RAN hub. See 15.52 “To configure optical protection switching on an 1830 VWM OPS shelf” (p. 513) and 15.53 “To configure bi-directional protection switching on an 1830 VWM OPS shelf” (p. 513) for more information.

The OPS shelf supports the following external switch commands from the 1830 VWM OSU:

- No Command
- Clear - clears all current external switch requests
- Forced Switch to Worker
- Forced Switch to Protection
- Manual Switch to Worker
- Manual Switch to Protection

OPS protection audit entity

The Protection Audit Entity (PAE) checks the consistency of the receive selectors in the near-end and far-end optical switches.

The following events are audited:

- elapsed time since the previous check is 10 minutes
- wait-to-restore timer expiry
- if one of the associated optical switches is re-initialized
- if one of the associated optical switches generates an event when it transits from a state in which both the inputs are failed to a state in which one or no input is failed

NFM-P supports the configuration of a PAE on the 1830 VWM OSU. See 15.54 “To configure an OPS protection audit entity on an 1830 VWM OSU shelf” (p. 515) for more information about configuring a PAE.

You must activate a PAE configured on an OSU. Upon activation, the PAE attempts to read the data from the near-end and far-end selector. If the data is not accessible, the PAE generates an AUDITBLOCK alarm on the optical protection switch module. You can deactivate an active PAE, which stops the auditing functionality and clears all involved alarms. You can delete an inactive PAE, which removes all associated configurations from the OSU database. See 15.55 “To activate or deactivate a PAE” (p. 515) for more information.

1830 VWM OPS – NTP

You can view the NTP parameters and server for the 1830 VWM OPS shelf by navigating to the NTP→General and NTP→Server tabs of the Shelf (Edit) form. The NTP Enabled check box is selected, by default, and cannot be configured. The Shelf Time parameter is read-only.

6.3.13 Statistics management

Table 6-2, “Supported statistics” (p. 214) lists the supported statistics.

Table 6-2 Supported statistics

1830 VWM devices	Supported statistics
1830 VWM PMU	• DDM statistics • Ethernet statistics • Optical statistics
1830 VWM TLU 9 and TLU 9M	• DDM statistics • Optical statistics • PCS statistics
1830 VWM TLU 200	• DDM statistics • Optical statistics • PCS statistics • Ethernet statistics • FEC statistics
1830 VWM OPS	• Ethernet statistics • Optical statistics
1830 VWM SMM	• DDM statistics • Ethernet statistics
1830 VWM ITP	• DDM statistics • Optical statistics • PCS statistics

See the procedure “To assign the default 1830 VWM OSU performance management policy to 1830 VWM devices” in the *NSP NFM-P Statistics Management Guide* for more information about assigning the default PM policy.

DDM data retrieval

See 16.2.2 “1830 VWM − DDM data retrieval” (p. 568) and 16.77 “To retrieve 1830 VWM DDM data” (p. 680) for more information about DDM data retrieval.

6.4 210 WBX

6.4.1 Overview

The 210 WBX is a compact datacenter gateway designed for leaf and spine deployments in data centers and IP fabric networks, or as an overlay VPN PE. The 210 WBX performs dataplane management using a Nuage SROS VM in a Linux hypervisor environment, allowing the installation of third-party tools and applications in the Linux layer. For more information about supported features and limitations, see the *NSP NFM-P Network Element Compatibility Guide* and the NFM-P Release Notice.

The NFM-P supports the following SR management features for the 210 WBX:

- Equipment management
- NE maintenance (backup and restore, software update)
- Services management
- Statistics & Accounting management
- ICMP Ping

Note: The Product parameter for WBX 210 NEs displays as DC 210 in the NFM-P.

210 WBX software upgrade

The NFM-P supports a 210 WBX hypervisor software upgrade. In order to upgrade the 210 WBX, you must upgrade the Linux hypervisor, which supports an ONIE image. A new ONIE image is downloaded to the NE SD card and the 210 WBX is rebooted with a hypervisor upgrade, which results in an upgrade of both the 210 WBX hypervisor and the SROS VM.

See 26.5 “To configure a software upgrade policy” (p. 776) for more information about how to configure a software upgrade policy.

6.5 7210 SAS

6.5.1 7210 SAS-D support

The 7210 SAS-D is an intelligent Ethernet edge-demarcation device that extends enhanced Carrier Ethernet VPN service delivery to the CE. The license for a 7210 SAS-D is based on the chassis type.

Note: The NFM-P supports only access mode for 7210 SAS-D 6F 4T ETR ports.

6.5.2 7210 SAS-Dxp support

The 7210 SAS-Dxp is an intelligent Ethernet demarcation switch that extends enhanced Carrier Ethernet VPN service delivery to the CE and can operate in access-uplink mode. The 7210 SAS-Dxp router uplinks to the network using Layer 2 Ethernet VLAN switching (without IP/MPLS).

The 7210 SAS-Dxp 16 and 24 port variants supports PoE connections. See 16.49 “To configure PoE ports on a 7210 SAS” (p. 641) for more information.

6.5.3 7210 SAS-E support

The 7210 SAS-E is a Carrier Ethernet CE device that is typically owned and managed by a customer. Using software based on the 7750 SR OS and managed by the NFM-P, the 7210 SAS-E extends Carrier Ethernet VPN services to the CE. The 7210 SAS-E can also be deployed as an aggregation device for smaller sites.

7210 SAS-E daughter cards

The 7210 SAS-E supports an integrated 2 × 12-Gig IOM card on a single chassis. The equipment navigation tree displays a card slot with one daughter card that contains 12 × 100/1000 Ethernet SFP ports and 12 × 10/100/1000 Ethernet ports.

6.5.4 7210 SAS-K support

The 7210 SAS-K is an intelligent Gigabit Ethernet switch that provides aggregation and demarcation for services managed to the customer edge.

6.5.5 7210 SAS-M support

The 7210 SAS-M is a CE device that provides MPLS-enabled metropolitan and WAN Carrier Ethernet service delivery, Ethernet-based mobile backhaul, and residential service access. In access uplink mode, the 7210 SAS-M provides Ethernet aggregation and demarcation for services managed to the customer edge.

The chassis mode is changed using a CLI to modify the device BOF. A reboot of the device is required for the change to take effect. When you change the chassis mode of a 7210 SAS-M, you must un-manage and re-manage the NE. The NFM-P displays the chassis mode as the State parameter, in the Uplink Mode panel on the General tab of the shelf properties form for the device.

See the NE documentation for more information.

6.5.6 7210 SAS-Mxp support

The 7210 SAS-Mxp is an IP/MPLS-enabled 1GigE and 10GigE access and aggregation device that can operate in two modes:.

- Standalone mode, in which the NE is managed as an IP/MPLS-enabled service aggregation device at the customer edge
- Satellite mode, in which the NE is connected by the uplink port to an SR device, to provide port expansion

Note: The 7210 SAS-Mxp is not a variant of the 7210 SAS-M.

The 7210 SAS-Mxp ETR supports PoE connections on ports 23 and 24. See 16.49 “To configure PoE ports on a 7210 SAS” (p. 641) for more information.

6.5.7 7210 SAS-R support

The 7210 SAS-R Ethernet switch is suitable for aggregation in access Ethernet networks, and is also capable of MPLS and MPLS-TP service transport. The 7210 SAS-R has multiple IMM card slots and dual redundant CPM/SF slots.

IMM card support

First-generation IMM cards (also called KT1 or imm-sas-r) are supported on the 7210 SAS-R6. Access, network, and hybrid port modes are supported. For 1-Gig ports, the port speed is variable.

The NFM-P supports a third IMM card family, called TR3 (or imm-sas-r-c) on supporting releases of the 7210 SAS-R6 and 7210 SAS-R12. Ports on TR3 cards are supported in network mode only. The port speed is not configurable.

For the 22 CSFP card, the NFM-P displays 22 ports on the navigation tree, even when SFP transceivers are used.

Automatic provisioning of IMM cards

The NFM-P automatically provisions and configures IMM cards on the 7210 SAS-R. Auto-provision occurs for unprovisioned cards during node discovery, or when the NFM-P receives notification that a card is inserted in a managed 7210 SAS-R.

The NFM-P auto-provisions cards that match the operational IMM family type displayed on the shelf properties form for the device. See 15.16 “To configure the IMM card type on a 7210 SAS-R” (p. 474). The 7210 SAS-R6 supports KT1, KT2, and TR3 IMM cards. If no IMM family type is configured, then by default only KT1 cards are auto-provisioned. The 7210 SAS-R12 supports KT2 and TR3 IMM cards. If no IMM family type is configured, then by default only KT2 cards are auto-provisioned.

6.5.8 7210 SAS-S support

The 7210 SAS-S is a 1-Gig and 10-Gig Ethernet switch that can operate in two modes:

- Standalone mode, in which the NE is managed as an IP/MPLS-enabled service aggregation device at the customer edge
- Satellite mode, in which the NE is connected by the uplink port to an SR device, to provide port expansion

The 7210 SAS-S is similar to the 7210 SAS-Sx, but has a reduced set of hardware features.

6.5.9 7210 SAS-S/Sx virtual chassis support

The 7210 SAS includes support for virtual chassis/stacking with 7210 SAS-Sx/S 1/10GE platforms for EPIPE, VPLS, RVPLS services with IP/MPLS. To increase port density at an access aggregation site, virtual stack of nodes can be configured using the stacking ports available on the rear of the chassis. To simplify the operations and management, the stack of nodes is presented as a virtual chassis, with a single IP address to use for managing the platform.

With stacking, operators have an option to add ports at an access aggregation site as connectivity requirements increase, while maintaining the operational simplicity of managing a single node.

7210 SAS-S/Sx VC in the NFM-P equipment tree

When the 7210 SAS-S/Sx VC is displayed in the NFM-P equipment tree, the IMM functionality of the CPM-IMM cards is shown as Card Slot- *n*, based on the configuration specified while bringing up the node. The CPM functionality of the CPM-IMM cards is shown as Card Slot- A and Card Slot- B.

Ethernet IMMs can be provisioned in any of the six IMM slots on the 7210 SAS-S/Sx VC, and can be managed by either of the CPM-IMM cards, depending on configuration and the operational state of the CPM-IMMs. If a CPM-IMM is removed or goes down, control of the MDAs is taken over by the other CPM-IMM to provide redundancy.

The power supplies and fan tray parameters are displayed under navigation tree equipment view for all the stack members. For each power supply and fan tray, the corresponding card number is listed. Power supply 1 on card 1 is displayed as 1/1 under power supplies.

6.5.10 7210 SAS-Sx support

The 7210 SAS-Sx is an Ethernet switch that can operate in two modes:

- Standalone mode, in which the NE is managed as an IP/MPLS-enabled service aggregation device at the customer edge
- Satellite mode, in which the NE is connected by the uplink port to an SR device, to provide port expansion

The 7210 SAS-Sx is similar to the 7210 SAS-S, but has an enhanced set of hardware features. NFM-P support for the 7210 SAS-Sx and 7210 SAS-S is not identical.

6.5.11 7210 SAS-T support

The 7210 SAS-T is an Ethernet access device that provides MPLS-enabled metropolitan and WAN Carrier Ethernet service delivery, Ethernet-based mobile backhaul, and residential service access. In access uplink mode, the 7210 SAS-T provides Ethernet aggregation and demarcation for services managed to the customer edge.

The NFM-P supports two chassis modes for the 7210 SAS-T:

- Access uplink: the device provides network service transport using Ethernet QinQ encapsulation. Ports are configured in access or L2 uplink mode.
- Network: the device provides network service transport using IP/MPLS encapsulation. Ports are configured in access, network, or hybrid mode.

The chassis mode is changed using a CLI to modify the device BOF. A reboot of the device is required for the change to take effect. When you change the chassis mode of a 7210 SAS-T, you must un-manage and re-manage the NE. The NFM-P displays the chassis mode as the State parameter, in the Uplink Mode panel on the General tab of the shelf properties form for the device.

See the NE documentation for more information.

The 7210 SAS-T ETR supports PoE connections on ports 19, 20, 21, and 22. See 16.49 “To configure PoE ports on a 7210 SAS” (p. 641) for more information.

6.5.12 7210 SAS-X support

The 7210 SAS-X is an MPLS-enabled Ethernet aggregation device for small and medium-sized networks that provides business, mobile backhaul, and residential services. The 7210 SAS-X is similar to the 7210 SAS-M, but has 10Gb/s uplink ports, enhanced traffic management, greater scalability, and hierarchical QoS functions.

6.5.13 System resource profile

The NFM-P supports system resource allocation on 7210 SAS devices using the system resource profile. You can configure system resources to suit your network requirements by assigning more resources to functions that are used extensively, and fewer resources to functions that are used less. Unused functions can be disabled, and their resources made available for other functions. The functional areas that are configurable vary depending on the chassis type.

System resources are allocated in units called chunks. The number of chunks and the chunk size vary by device chassis.

Resources are allocated by assigning a value from the applicable range. A value of zero disables the function. The sum of resources allocated must not exceed the total available resources for the system. If you assign more than the allocated number of chunks, functions are disabled.

For SAP ACL and SAP QoS functions, configuring system resources is accomplished in two steps. First, divide overall system resources among the main functional areas: SAP ingress internal ACL, SAP egress internal ACL, and SAP ingress QoS. Second, allocate the resources specified for each functional area to the criteria within that function, such as MAC, IPv4, and 64 or 128-bit IPv4/v6. The total number of chunks assigned to the criteria must not exceed the number of chunks allocated to the function. When you specify a value of MAX, the system allocates resources as required on a first-come, first-served basis, up to the available limit.

For all 7210 SAS chassis types except the 7210 SAS-R and 7210 SAS-VC, system resource profile allocations are configured globally on the NE properties form; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

For the 7210 SAS-R and 7210 SAS-VC, some resource allocations are configured globally on the NE properties form, and some are configured using system resource profile policies assigned to card slots on the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382) and 15.71 “To select system resource profile policies for the 7210 SAS-R or 7210 SAS-S/Sx VC” (p. 529) .

See the 7210 SAS documentation for more information about system resource allocation.

6.6 7250 IXR

6.6.1 Overview

The 7250 IXR interconnect router is deployed for interconnect applications in core, metro, and data center networks, with optimized power efficiency, expandability, and modularity.

See the *NSP NFM-P Network Element Compatibility Guide* for supported IXR variant information.

Note: You can enable USB storage on all IXR variants except for IXR-s, where USB support is enabled by default. In the equipment tree and on Alarm Info forms, the NFM-P displays only Flash Memory. The display does not distinguish between compact flash memory and USB flash memory.

6.6.2 GNSS receiver support

The GNSS receiver supports GPS and GPS plus GLONASS. The GNSS port is listed under Management Port (expand Network→ NE→ Properties→ Inventory→ Management Port (Physical Equipment)→ gnss port. See 12.68 “To configure GNSS receiver functions on supported IXR and SR NEs” (p. 397) for more information.

6.6.3 System resource profile

The NFM-P supports system resource allocation on 7250 IXR devices using the system resource profile. System resource profile allocations are configured globally on the NE properties form; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

6.6.4 Forwarding path resource allocation

Supporting releases of the 7250 IXR allow custom forwarding path resource allocation for resource-intensive functions such as multicast, ACL classification criteria, policers, etc. The NFM-P supports this resource allocation using the NE properties form; see 12.57 “To configure forwarding path options or resource allocation on a 7250 IXR” (p. 388).

6.7 7450 ESS

6.7.1 Overview

The Nokia 7450 Ethernet Service Switch supports service-aware Ethernet aggregation across IP/MPLS-based networks, providing VPLS, VLL and advanced IP services. The 7450 ESS supports a wide range of interfaces.

A Network Domain feature allows users to determine which network ports are eligible to transport traffic of individual SDPs. This information is used for the SAP-ingress queue allocation algorithm applied to VPLS SAPs. See Chapter 27, “NE routing and forwarding” for more information.

The NFM-P provides network management functions for the 7450 ESS.

The 7450 ESS OAM toolkit includes ITU-TY.1731 with Synthetic Loss Measurement (SLM), IEEE 802.1ag, IEEE 802.3ah, Ethernet local management interface (E-LMI), VPLS OAM and service mirroring. This toolkit is also integrated into the NFM-P.

6.7.2 Mixed-mode

Mixed-mode is the default behavior for all 7450 ESS chassis. This mode of operation allows a 7450 ESS chassis to support all 7750 SR functionality. All supported IOMs and IMMs for the 7450 ESS chassis are 7750 SR-capable, so the 7450 ESS chassis will always operate in this mode.

6.7.3 Automatic provisioning for CPM5 SFMs

The NFM-P automatically provisions and configures the integrated SFMs of CPM5s that feature built-in modules. The automated provisioning occurs during node discovery.

6.7.4 7750 SR-B chassis support

The 7750 SR-12-B and 7750 SR-7-B chassis types are supported on both 7450 ESS and 7750 SR, Release 15.0 R5 or later. The chassis employ the FP4 chipset, which supports multi-terabit-per-second routing.

All existing features, IOM/IMMs, SFMs, CPMs, PEMs, and fan trays are supported in the B chassis.

6.8 7705 SAR

6.8.1 Overview

The 7705 SAR is an IP/MPLS aggregation and mobile backhaul router for the mobile RAN. Located at cell sites, the 7705 SAR uses PW over MPLS to aggregate mobile 2G and 3G traffic that it backhauls to the core network. The 7705 SAR supports ATM, TDM, and Ethernet.

6.8.2 7705 SAR MWA support

A packet microwave card is required to support MWA. The first four ports of the card are microwave aware. An MW link must be configured to enable MWA. MW link support includes 1+0 and 1+1 HSB protection. When configuring 1+1 HSB, each of the two radios is assigned a role of main radio or spare radio. By default, the main radio assumes the active status and the spare radio assumes the standby status. For 1+0 protection, MPT-HC, MPT-HCv2, MPT-HLC, MPT-HQAM, and MPT-MC are supported. For 1+1 HSB protection, MPT-HC, MPT-HCv2, MPT-HLC, and MPT-HQAM are supported.

1+1 HSB protection is displayed in the NFM-P physical topology map. The protection state is shown on the links, and 1+1 links are grouped in a common link group. The expansion of radio links on the physical map displays eight links. See 4.1 “Topology map types” (p. 169) for more information about the NFM-P physical topology map.

Note: If the Wavence SA is discovered before the 7705 SAR (single NE with 1+0 protection), the radio link is not shown on the topology map. Only after discovering the 7705 SAR will the radio link be shown on the topology map between the 7705 SAR and the Wavence SA.

1+1 HSB protection is only supported between two 7705 SAR NEs.

The following protection types are supported:

- Radio Protection Switching (RPS)

  RPS is always supported in 1+1 radio configurations (HSB). RPS is implemented directly on a pair of MPTs.
- Transmission Protection Switching (TPS)

  TPS is always supported in 1+1 HSB radio configurations.
- Equipment Protection Switching (EPS)

  EPS protects the MPT and cables connecting the MPT to the IDU. RPS protection always assumes the MPT EPS.

See 15.38 “To configure a 7705 SAR MW link” (p. 499) for more information about configuring an MW link. See 15.39 “To configure a 7705 SAR MW link member” (p. 501) for more information about configuring an MW link member.

6.8.3 7705 SAR NEtO support

The NFM-P supports an integrated MCT from the port or NE level on the 7705 SAR-8 and 7705 SAR-18. The MCT is distributed with a NEtO client in the NFM-P. The configuration of radio parameters is through the MCT. See 12.61 “To launch an MCT on a Wavence SA NE connected to a 7705 SAR” (p. 392) for more information about starting the MCT from the NE level.

6.8.4 7705 SAR SCADA support

The NFM-P supports SCADA on the 7705 SAR-8 and 7705 SAR-18, Release 6.1 R1 or later. SCADA is an industrial data management system that monitors and controls IEDs. See 15.1.1 “Shelf objects in the NFM-P navigation tree” (p. 453) for more information about SCADA. See 15.36 “To configure SCADA on a 7705 SAR” (p. 495) for more information about configuring SCADA on the 7705 SAR.

6.8.5 7705 SAR ISC

The 7705 SAR-8 and 7705 SAR-18 support an ISC that is equipped with a large FPGA and a Winpath network processor. The FPGA supports multiple applications on the same card by using separate images while providing the ability to perform IP/MPLS functionality such as pseudowire termination.

The ISC only supports one application running per card at a time. Multiple ISCs can be supported on a 7705 SAR-8 or 7705 SAR-18. At least one networking card is required to provide access to the ISCs.

Network or access mode configuration is not applicable on the ISC; however, the card provides access capabilities.

An ISC is required to configure SCADA and VCB on the 7705 SAR-8 and 7705 SAR-18. For more information about SCADA, see 15.1.1 “Shelf objects in the NFM-P navigation tree” (p. 453) . For more information about VCB, see 15.37 “To configure voice conference bridging on a 7705 SAR” (p. 496).

Note: MDDB and PCM multidrop bridge applications are supported on the ISC.

6.8.6 7705 SAR-A

The 7705 SAR-A is equipped with integrated twelve-port Ethernet MDA daughter cards.

6.8.7 7705 SAR-Ax

The 7705 SAR-Ax transports all types of data from a mobile cell site to a higher aggregation point of presence. The 7705 SAR-Ax also targets fixed and vertical networks.

The 7705 SAR-Ax maintains the same data plane forwarding capacity and control plane scalability as the 7705 SAR-A. The IP MPLS features are also similar to the 7705 SAR-A.

The security features such as IPsec, NAT, Firewall, and NGE are similar to the 7705 SAR Wx.

6.8.8 7705 SAR-F

The 7705 SAR-F integrates eight-port Ethernet v3 and 16-port DS1/E1 v2 ASAP MDA functions in one chassis. When the NFM-P discovers a 7705 SAR-F, the integrated MDAs are automatically configured and displayed in the equipment tree.

6.8.9 7705 SAR-H

The 7705 SAR-H is a hardened router and is based on the 7705 SAR-8, 7705 SAR-18, 7705 SAR-F, and 7705 SAR-M functionality.

The NFM-P allows you to configure the MDA mode of some MDA types on the 7705 SAR-H. There are two types of MDAs:

- Compact MDAs
- Combo MDAs

The 7705 SAR-H supports the 4 × 10/100 Ethernet module card in slots 2 and 3. The module card supports four fast Ethernet ports in access or network mode. The ports can be bound to a network interface, all routing protocols, and a VPRN, IES and VLL service.

The 7705 SAR-H supports power supply tray objects, which are automatically provisioned.

6.8.10 7705 SAR-Hc

The 7705 SAR-Hc is equipped with six Ethernet ports:

- two SFP Gigabit Ethernet ports
- two RJ-45 10/100/1000 Ethernet ports
- two PoE-capable RJ-45 10/100/1000 Ethernet ports

The 7705 SAR-Hc is equipped with a single internal power supply that operates at +/–20 V dc to 75 V dc. When the 7705 SAR-Hc is used for high-voltage applications, an external high voltage power supply is available.

The 7705 SAR-Hc supports power supply tray objects, which are automatically provisioned.

6.8.11 7705 SAR-Hm

The 7705 SAR-Hm is a small-form wireless router that extends IP/MPLS services over secure 3G/LTE wireless networks. The 7705 SAR-Hm is equipped with six Ethernet ports, two RS-232 ports, two SIM card slots, a WLAN interface, GPS, a GNSS receiver, and an alarm port. The 7705 SAR-Hm system software is a variant of the SROS, and supports IP/MPLS features that are supported on 7750 SR nodes.

7705 SAR-Hm configurations

Information about SIM cards installed on the 7705 SAR-Hm is displayed on the Cellular tab of the cellular port configuration form. The 7705 SAR-Hm is discovered through ADP-Hm; see 47.1.10 “Device discovery and deployment using ADP” (p. 1427). For information on cellular domain management; see Chapter 47, “Cellular domain management”.

The following table lists where to find information about configuring the 7705 SAR-Hm:

Table 6-3 Configuring the 7705 SAR-Hm

To configure:	See:
Cards, ports and interfaces	
Cellular ports	16.27 “To configure a cellular port on a 7705 SAR-Hm” (p. 609)
Dual SIM switchover parameters	15.78 “To configure an MDA” (p. 536) 15.85 “To perform a manual SIM switchover on a 7705 SAR-Hm ” (p. 546)
PDN policies	Chapter 69, “PDN profile policies”
PDN interfaces	27.4 “To configure a cellular interface on a 7705 SAR-Hm” (p. 839)
Serial raw sockets	Chapter 21, “Serial raw sockets for IP transport services”
GNSS for IP transport	21.9 “To configure IP transport on a VPRN site” (p. 729) 21.8 “To configure IP transport on an IES site” (p. 727)
GNSS for data streaming, location details, and listing visible satellites	15.84 “To configure GNSS receiver functions on a 7705 SAR-Hm” (p. 545)
WLAN ports (Access Point)	16.28 “To configure a WLAN port on a 7705 SAR-Hm” (p. 610)
Security	
New security PIN	47.10 “To configure a new PIN value for a cellular carrier” (p. 1452)

6.8.12 7705 SAR-Hmc

The 7705 SAR-Hmc is a compact variant of the 7705 SAR-Hm. The 7705 SAR-Hmc is equipped with three Ethernet ports, two RS-232 ports, two SIM card slots, GPS, a GNSS receiver, and an alarm port. The 7705 SAR-Hmc system software is a variant of the SROS, and supports IP/MPLS features that are supported on 7750 SR nodes. Besides the 7705 SAR-Hmc having three Ethernet ports (compared to six Ethernet ports on the 7705 SAR-Hm), and no WLAN interface, the 7705 SAR-Hmc supports the same features as the 7705 SAR-Hm.

You must configure CBSD information for the 7705 SAR-Hmc to execute the CBSD state machine, register with the CBRS Spectrum Access System (SAS) and be granted access before ADP-Hm is allowed to run. Once “Authorized”, the ADP-Hm can start sending SNMP traps over the router interface. For CBSD-related configuration, see 16.27 “To configure a cellular port on a 7705 SAR-Hm” (p. 609).

In the NFM-P documentation, mentions of 7705 SAR-Hm also apply to 7705 SAR-Hmc unless otherwise stated, or unless a WLAN interface is required.

The 7705 SAR-Hmc supports both ADP and non-ADP radio card firmware upgrade.

6.8.13 7705 SAR-M

The 7705 SAR-M variants are equipped with an integrated 7-port Ethernet card. Depending on the variant, a 7705 SAR-M may have an integrated 16-port channelized DS1/E1 ASAP card or a module slot.

6.8.14 7705 SAR-W

The 7705 SAR-W is equipped with an integrated daughter card that supports five data plane Ethernet ports: three SFP ports and two RJ-45 ports. The Ethernet ports are classified as permanent connections with primary surge protection.

6.8.15 7705 SAR-Wx

The 7705 SAR-Wx transports all types of traffic over a packet-switched network from a mobile cell site to a higher point of presence or to a mobile telephone switching office.

The Ethernet data ports support the following types of connections:

- metrocell connections
- microwave device connections—NLOS or LOS microwave connections on the RJ-45 PoE+ ports
- network connections—Gigabit Ethernet fiber connections on the three SFP ports

6.8.16 7705 SAR-X

The 7705 SAR-X is a fixed variant with three integrated MDAs.

6.8.17 Dynamic system IP address updates for 7705 SAR nodes

When a 7705 SAR node is configured to acquire an address from a DHCP pool, the system IP address may change. You can configure the NFM-P to react automatically to the change, and adjust to use the new IP address for communication and for SDPs destined for the node.

See the procedure to enable dynamic system IP address updates for 7705 SAR nodes in the *NSP System Administrator Guide* for information about enabling this feature.

6.9 7750 SR

6.9.1 Overview

The 7750 SR is a multi-service edge router designed for service providers, cable MSO, and enterprise customers that deliver residential, business and mobile services through an IP/MPLS network.

See the *NSP NFM-P Network Element Compatibility Guide* for supported SR variant information.

6.9.2 Ethernet port utilization summary

Ethernet ports on NEs using SR-OS display a summary of input and output port utilization on the Ethernet tab of the port configuration form. The summary is expressed as a percentage, and summarizes utilization over the interval specified in the Utilization Stats Interval parameter.

6.9.3 Automatic provisioning for CPM5 SFMs

The NFM-P automatically provisions and configures the integrated SFMs of CPM5s that feature built-in modules. The automated provisioning occurs during node discovery.

6.9.4 7450 ESS emulation

A 7750 SR node can be configured to operate in 7450 ESS emulation mode. While in emulation mode, the 7750 SR appears in the NFM-P as a 7450 ESS node, and supports 7450 ESS features and cards. See the 7750 SR node documentation for information about how to configure 7450 ESS emulation mode, and which variants of the 7750 SR support emulation.

6.9.5 ESA support

The Extended Services Appliance (ESA) operates as a resource server that provides packet buffering and processing, and is logically part of the 7750 SR router system. The ESA connects to the 7750 SR using standard SR interface ports, so all communication passes through the IOM, making use of the network processor complex on the host IOM for queuing and filtering functions, as with other MDAs and ISAs. The ESA hosts up to four Virtual Machine (VM) instances for multiservice processing.

See the *NSP NFM-P Network Element Compatibility Guide* for information about which SR devices support ESAs.

6.9.6 GNSS receiver support

All FP5-enabled SR NEs support integrated dual band GNSS receivers. The GNSS receiver supports GPS and GPS plus GLONASS. See 12.68 “To configure GNSS receiver functions on supported IXR and SR NEs” (p. 397) for more information.

6.9.7 Discovering VSR-I using an AIM agent

The AirFrame Infrastructure Manager (AIM) is a proxy SNMP agent running on an AirFrame box in a VSR-I stack.

You need to initiate the discovery of the AIM in order to automatically discover the VSR-I . After the AIM for the VSR-I is discovered, the VSR-I is automatically added to the AIM discovery rule, so that the VSR-I is automatically discovered.

See 9.13 “To configure the AIM mediation and discovery for management of the VSR-I” (p. 298).

6.10 7850 VSG/VSA

6.10.1 Overview

The Nokia 7850 Virtual Switch Gateway (7850 VSG) is the main hardware component of the data center solution. The 7850 VSG is a top-of-rack switch that provides connectivity to servers which store virtual machines (VMs). It is based on the SROS platform and includes standard Ethernet and L3 routing functionality.

The 7850 Virtual Switch Aggregator (7850 VSA) is a variant of the 7850 VSG. The 7850 VSA functions as an IP router as an end-of-rack switch. It includes the SROS feature set, like the 7850 VSG. It can be deployed with the standalone VSC in a two-part router setup.

6.11 7950 XRS

6.11.1 Overview

The Nokia 7950 XRS is a large-scale routing system that is designed for core deployments. The NFM-P supports the management of the 7950 XRS including device discovery, equipment management, routing management, OAM, MPLS, QoS policies, and services.

6.11.2 Shelf components

You can monitor and manage the following 7950 XRS shelf components using the navigation tree:

- CCMs
- XCM or XCM2 card slots; each slot has two XMA daughter card slots
- CPM card slots
- SFM or SFM2 card slots
- fan trays
- PEQ power supply trays
- PCM trays

The number of the preceding components depends on the variant of the 7950 XRS.

The shelf components are N+1 redundant. The fan tray and power supply tray objects are automatically discovered and provisioned, and require no configuration using the NFM-P. See Chapter 12, “Device object configuration” for more information about configuring shelf objects.

CCMs

The 7950 XRS-20 shelf includes CCMs that support operator access to the routing system. You can view and monitor CCM and flash memory module properties from the CCM properties form. The CCMs are automatically detected and provisioned.

Each CCM includes three flash memory modules which can be viewed and monitored by the NFM-P. The “cf1” module is an internal flash module that is embedded in the system. The “cf2” and “cf3” modules are external compact flash slots.

XCM cards

The 7950 XRS-20 shelf includes ten XCM card slots that can be configured with XCM X20 cards. The 7950 XRS-16c shelf includes eight XCM card slots that can be configured with XCM X16 cards. Each XCM card can house up to two XMA cards for a total of twenty XMAs per 7950 XRS-20 shelf or sixteen XMAs per 7950 XRS-16c shelf. XMA card slots can also be configured with a compact XMA card, or C-XMA card. The 7950 XRS-16c can be configured only with C-XMA cards.

Note: To configure the 40-port 10 GE SFP+ XMA or 4-port 100 GE CXP XMA cards on the 7950 XRS-20, you must configure SFM X20 B cards on the shelf. You cannot configure the XMA cards if the shelf is configured with SFM X20 cards.

XCM2 cards

These cards enable redundant switching capacity up to 2.4T FD per XMA or 4.8T FD per XCM, and can accommodate any FP3-based or FP4-based XMA.

CPM cards

The 7950 XRS shelf includes two redundant CPM card slots. Each slot is automatically provisioned with CPM X20 or CPM X16 cards, depending on the variant of the 7950 XRS.

On the 7950 XRS-20, each CPM card slot has an association with a CCM. If one of the components is removed, the appropriate status is applied to the corresponding CPM or CCM. In the navigation tree, the CPM in Card Slot A is associated with CCM A. The CPM in Card Slot B is associated with CCM B.

On the 7950 XRS-16c, each CPM X16 card includes three flash memory modules, which can be viewed and monitored with the NFM-P. There is no associated CCM for the CPM X16.

SFM cards

The 7950 XRS shelf includes eight SFM card slots. SFM card slots on the 7950 XRS-20 can be configured with SFM X20 or SFM X20 B cards. SFM card slots on the 7950 XRS-16c can be configured with SFM X16 or SFM X16 B cards. You cannot configure a mix of SFM card types cards on one shelf. All configured slots must contain the same type of SFM card. Seven SFM cards must be operational to provide full capacity to the XCM cards, and one SFM card is redundant. At least one SFM card must be configured for the system to function.

SFM2 cards

After SFM2 slots are configured, the new XCM2 cards can be configured. The SFM2 and XCM2 cards enable redundant switching capacity up to 2.4T FD per XMA or 4.8T FD per XCM. SFM2 cards support FP3-based and FP4-based XCM cards. FP3-based XCMs can be intermixed with FP4-based SFM2s.

Automatic provisioning for CPM5 SFMs

The NFM-P automatically provisions and configures the integrated SFMs of CPM5s that feature built-in modules. The automated provisioning occurs during node discovery.

6.12 Generic NEs

6.12.1 Overview

The NFM-P provides limited management support of generic NEs, or GNEs, which are typically non-Nokia devices.

NFM-P GNE support is limited to the following:

- display on topology maps and in the navigation tree, interfaces excluded
- physical link creation and representation as a graphical and logical object
- trap translation into alarms
- status polling using enterprise MIB object interface alarms

- Telnet, SSH, and web sessions invocation
- provisioning using CLI scripts
- inclusion in a service instance using NFM-P scripts

6.12.2 Statistics support

The NFM-P supports the collection of a limited set of statistics counters from standard system, interface, and routing MIBs on GNEs. These statistics are processed and presented in the same manner as statistics from other devices. You can view GNE statistics on the Statistics tab of a GNE interface properties form, retrieve them using the XML API, and display them graphically using the NFM-P Statistics Plotter.

Note: If persistent SNMP indexes are not enabled on a GNE, one or more GNE interface indexes may change after a GNE reboots. This can cause a mismatch between the statistics records collected before the reboot and the current interface indexes. The NFM-P takes no action to identify or correct such a mismatch.

6.12.3 Alarm support

By default, the NFM-P supports a limited number of standard system and interface SNMP traps for GNEs. You can configure the NFM-P to raise user-defined alarms in response to specific GNE traps using alarm catalogs. See “GNE commissioning” (p. 246) for information about configuring user-defined alarms for GNEs.

6.12.4 GNE drivers

For some specific GNEs, drivers have been developed to extend management of the NE beyond the basic GNE management functions. Drivers translate proprietary MIBs to the normalized NFM-P model. See Chapter 7, “Device management using drivers” for information about drivers.

6.13 OmniSwitch

6.13.1 Overview

CAUTION

Service Disruption

OmniSwitch devices do not support automatic synchronization with the NFM-P database when you use the CLI to make configuration changes.

To ensure that you are viewing accurate OmniSwitch information in the NFM-P, you must resynchronize the NE by clicking on the appropriate Resync button.

OmniSwitch NEs do not send trap notifications for all MIB changes on the NE. To ensure that you are viewing the current configuration information in the NFM-P, you must resynchronize the NE by clicking on the appropriate Resync button.

The OmniSwitch family of devices is a group of Enterprise and Metro Ethernet LAN switches. The NFM-P supports the OmniSwitch chassis types listed in the *NSP NFM-P Network Element Compatibility Guide*.

6.13.2 OmniSwitch LAG objects

OmniSwitch NEs support static and dynamic LAGs. When you create a static LAG, you can add ports as LAG members. The ports that you select as members of a dynamic LAG are first placed into an Unassigned Dynamic LAG Members group. The OmniSwitch uses LACP to dynamically assign ports from an Unassigned Dynamic LAG Members group to the appropriate LAG. When a port is assigned to a dynamic LAG, the port is removed from the Unassigned Dynamic LAG Members group.

You can create VLANs, use 802.1Q framing, configure QoS conditions, and enable other networking functions on a LAG, which is treated the same as a physical link.

OmniSwitch LAG support includes:

- up to 32 LAGs on an OS 6250, OS 6400, OS 6450, OS 6465, OS 6850, OS 6850E, OS 6855, or on a stack of OmniSwitch devices
- up to 128 LAGs on an:
  - OS 9600, OS 9700, and OS 9800
  - OS 9700E, or OS 9800E
  - OS 10K
  - OS 6860, OS 6860E, OS 6860N, OS 6865, or OS 6900
- two, four, or eight Ethernet links in a LAG
- access or network LAGs

6.13.3 Using WebView to manage an OmniSwitch

You can configure and manage an OmniSwitch using the WebView application, a web-based EM tool that runs on an OmniSwitch. See the*NSP Planning Guide* for information about browser compatibility.

Note: You must enable cookies and Java in the web browser that you use to open WebView.

See the appropriate OmniSwitch *Switch Management Guide* for information about configuring and using WebView. See Chapter 12, “Device object configuration” for information about opening WebView using the GUI.

6.13.4 OmniSwitch port modes

The ports of OmniSwitch NEs can operate in either access mode or network mode. Access ports can be added to stacked VLANs as UNI ports, or to standard VLANs as tagged or untagged ports. Network ports can automatically bind to stacked VLANs as NNI ports when the service is created. OmniSwitch NEs using AOS Release 6.4.6 or later, or AOS Release 6.6.5 or later, can operate in hybrid mode. Hybrid ports can be added to standard VLANs as tagged or untagged ports, and also bind to all stacked VLANs on an NE as NNI ports.

Note: When a port is changed to hybrid mode, any previous UNI associations will be deleted.

6.13.5 OmniSwitch splitter ports

The NFM-P supports splitter ports and sub-ports where equipped on OmniSwitch devices (typically, QSFP+ ports). Sub-ports are identified by the CLI Name parameter on the port properties form.

6.14 Wavence SM and Wavence SA

6.14.1 Overview

The Nokia Wavence SM is a microwave digital radio that supports PDH and Ethernet to migrate from TDM to IP. The Wavence SM provides a generic, modular IP platform for multiple network applications, such as 2G, 3G, HSDPA, and WiMAX to accommodate broadband services.

The Wavence SA is the standalone, fully outdoor application of the MPT ODU, with no shelf unit. The Wavence SA provides fixed or mobile Ethernet site backhauling and supports converged MPLS metro networks.

The following documents describe NFM-P discovery, management, configuration, and administration of Wavence SM/Wavence SA devices:

- *NSP NFM-P Classic Management User Guide* — Describes general device discovery, configuration, and management
- *NSP Wavence Device Support Guide*— Describes how to commission, configure, and manage Wavence SM and Wavence SA devices, policies and services
- *NSP System Administrator Guide* — Describes device security management

7 Device management using drivers

7.1 Overview

7.1.1 Support for drivers in the NFM-P

Note: NFM-P multi-vendor drivers are no longer supported; instead, multi-vendor support is managed using model-driven adaptors. See the Network Developer Portal for more information.

Drivers are optional software modules corresponding to specific devices. They can be used to increase or extend management of NEs currently managed as GNEs. Drivers mediate between the NE and the NFM-P, allowing the NFM-P operator to manage non-standard equipment and services from the GUI or the NFM-P XML OSS interface.

See section 6.12 “Generic NEs” (p. 229) for more information about managing other vendors’ NEs without a driver.

With the addition of a driver, management capabilities are added in the following areas, depending on the driver and release:

- configuration management
- software management
- network assurance
- service assurance
- service provisioning
- device life cycle management

7.1.2 Driver development

A number of GNE drivers are developed, delivered, and installed on the same product release cycle as the NFM-P as part of the overall product offering. Many of these drivers have been designed to support specific capabilities relevant to a particular managed device, such as radio port properties configuration for Wavence MPT devices.

7.2 Driver framework capabilities

7.2.1 Capabilities summary

The NFM-P driver framework offers a robust feature set that all drivers can employ to attain management parity with NEs natively managed by the NFM-P. The extent to which each driver takes advantage of the capabilities of the NFM-P driver framework varies from driver to driver, however.

In conjunction with the NFM-P driver framework, drivers enable users to discover, provision, and manage other vendors’ equipment and services using the same XML API, script and GUI workflows

as are used for natively managed Nokia devices and services. When a driver is employed, the NEs become available in NFM-P map and service topology views, and configuration management is synchronized through the GUI or through an XML API interface.

The driver architecture supports SNMP, NETCONF, and CLI management protocols.

Note: Manual or scheduled resynchronization of a device is required if any direct CLI changes are made while the device is managed by the NFM-P.

The NFM-P discovers the NEs that will be managed by drivers using a GNE Profile. The NE and the driver are associated with a GNE Profile through the standard sysObjectId MIB attribute of the NE and the “Driver module” selection, respectively. Once this association is made, the NFM-P expects a driver to be used for all NEs of that type. You can either manage all nodes of a particular SNMP sysOID with a driver or all nodes as a GNE. In most cases, a GNE Profile for the driver is automatically created in the NFM-P.

7.2.2 Extended management features

The NFM-P driver framework supports drivers designed to extend NE management in the following areas:

- node, shelf and port discovery and display in the GUI topologies
- MPLS dynamic LSP and SDP discovery and provisioning
- VLL Epipe service discovery and provisioning
- VPRN service discovery and provisioning (PE-CE IP, OSPF, BGP, static routes, loopback on L3 access interface, L3 access interface on physical ports)
- VPLS service discovery and provisioning (RSVP-TE and LDP)
- LLDP link endpoint specification at the physical port level rather than GNE interface level
- alarm integration, correlation and aggregation
- QoS and policer control policy definition
- interface statistics and plotting

For information about RCA audits for Epipe, VPLS, and VPRN services with multi-vendor NE sites, see 95.3.3 “RCA audit for services with multi-vendor NE sites” (p. 3232).

7.3 Driver availability

7.3.1 Timing

GNE drivers and new driver versions are developed and delivered on the same product release cycle as the NFM-P.

7.3.2 Location

GNE drivers are installed with the NFM-P software. It is no longer necessary to obtain these drivers separately from the NSP software delivery site.

7.3.3 Documentation

Driver compatibility information is located in the *NSP NFM-P Network Element Compatibility Guide*. Driver capabilities and driver-specific discovery and management information are provided in the appropriate domain user guides.

Table 7-1 Location of driver information

Driver name	Reference
MPTGM MPTGS MPTSUB6 MPTBWA is installed with the name MPTSUB6	*NSP Wavence Device Support Guide*
CloudBand Application Manager	*Network and Service Assurance Guide*

Note: The 9500 MPT-BWA device is a hardware variant of the MPT-SUB6. The NFM-P uses the MPT-SUB6 driver to manage the MPT-BWA.

7.4 Driver installation and upgrade

7.4.1 Overview

Drivers that manage the devices listed in Table 7-1, “Location of driver information” (p. 235) are automatically installed during NFM-P installation and are preserved during upgrade.

7.5 View installed drivers on the NFM-P

7.5.1 Steps

Perform this procedure to determine which drivers are installed on the NFM-P.

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Driver (Generic NE) from the drop-down menu and click Search. A list of installed drivers appears.

3

To view the properties of a driver, select a driver and click Properties. The Properties form opens. You can view the driver and driver version.

The Maximum NFM-P Version parameter shows the highest NFM-P release that the driver supports.

The Operational State parameter shows the current operational status of the GNE driver. The operational state can be Operational, Unsupported, or File Missing.

Note: GNEs cannot be managed using a driver which is not Operational. If a driver is in a File Missing state, the driver must be re-installed. The driver cannot be replaced until the original driver is re-installed. If a driver is in an Unsupported state, a supported driver is required to replace the unsupported driver.

END OF STEPS

7.6 View the automatically created Generic NE profile

7.6.1 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Profile (Generic NE) and click Search. A list of Generic NE Profiles appears.

3

Select a Generic NE Profile and click Properties. The Generic NE Profile (Edit) form opens with the General tab displayed.

4

Click on the Interface Types, Other MIBs, and Translators tabs to view information about the Generic NE Profile.

5

Close the form.

END OF STEPS

7.7 View the automatically created alarm catalog

7.7.1 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Alarm Catalogue (Trap to Alarm Mapper) and click Search. A list of alarm catalogs appears.

3

Select an alarm catalog and click Properties. The Generic Alarm Catalogue (Edit) form opens.

4

Click on the Mappings, Transform Functions, and Generic NE Profiles tabs to view information.

To modify the alarm catalog, see 8.17 “To create a GNE alarm catalog” (p. 269).

5

Close the form.

END OF STEPS

8 Device commissioning and management

Device commissioning

8.1 Overview

8.1.1 Preconfiguration

A device requires preconfiguration before NFM-P management of the device is possible. When the preconfiguration is complete, the NFM-P can discover the device, as described in Chapter 9, “Device discovery” . See 8.3 “Workflow to commission Nokia devices” (p. 244) for the device commissioning workflow.

8.1.2 In-band and out-of-band management

The NFM-P supports the in-band and out-of-band management of devices. The following figure shows an example of in-band and out-of-band management.

Figure 8-1 Example of in-band and out-of-band management

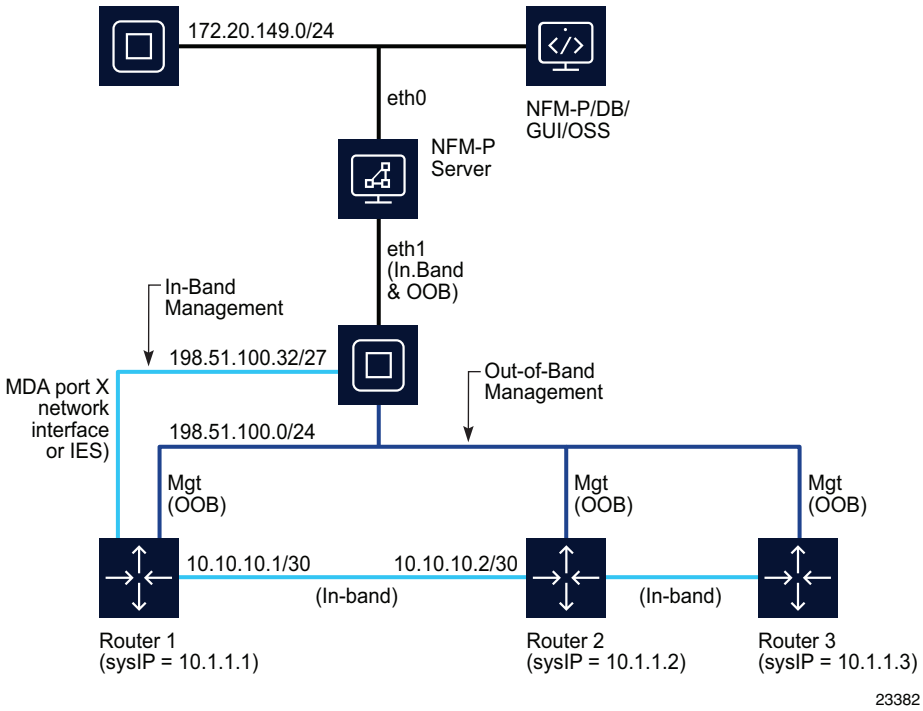

When you configure only in-band management, management traffic between the NFM-P and a device is transmitted through any port that is configured for network access, except the

management port. For in-band management, the NFM-P sends management traffic to the device system IP address, or to an optionally specified L3 management interface.

When you configure only out-of-band management, management traffic between the NFM-P and a device is transmitted through the management port of the device. For out-of-band management, the NFM-P sends management traffic to the management IP address of the device.

When a device is configured for both in-band and out-of-band management, one method provides management redundancy for the other. However, if the in-band and out-of-band IP addresses are the same, the NFM-P cannot determine which NE interface is active, and management redundancy is compromised. If the in-band and out-of-band IP addresses are the same on a newly discovered NE, the NFM-P raises an alarm.

Note: The OmniSwitch does not support management redundancy.

The NFM-P uses one SNMP trap target for in-band and out-of-band management. The NFM-P SNMP IP address must be routable from each IP address when in-band and out-of-band management are used. See the device documentation for information about configuring an SNMP trap target.

The type of management is determined during device discovery. When the device is discovered using its management IP address, system IP address, or L3 interface IP address, the in-band and out-of-band addresses are synced in the NFM-P. This enables configuration of management redundancy, if both in-band and out-of band are available. In each case, a valid route to the device must exist. See Chapter 9, “Device discovery” for more information about the discovery process.

In Figure 8-1, “Example of in-band and out-of-band management” (p. 239) , an out-of-band management route allows a ping from the NFM-P to the management IP address of Router 1 (192.168.10.1/32). The in-band connection sends management packets to the system IP address on Router 2 (12.12.12.2/32) on MDA port X.

Important information

CAUTION

Service Disruption

Do not use the NFM-P to modify the configuration of the in-band port. If the port is shut down, network visibility is lost.

Note: In Figure 8-1, “Example of in-band and out-of-band management” (p. 239) , the IP address of Routers 1 and 2 ends with /*number*, which is the subnet mask. Common subnet mask values are:

- /31 for two hosts
- /30 for two hosts with network and broadcasts numbers
- /24 for 254 hosts

Some device types require configuration in addition to the SNMP configuration before they can be managed by the NFM-P. See the appropriate part of this section for device-specific management information.

8.1.3 Firewalls and management bandwidth

The ports between NFM-P components, and between the NFM-P system and the managed devices, must be open through firewalls to allow proper operation of the software. See the*NSP Planning Guide* for more information about requirements for the following:

- firewalls and open ports
- communication bandwidth between NFM-P components
- communication bandwidth between the NFM-P and managed network

8.1.4 IPv6 device management

The NFM-P supports device management using IPv6. When a device management port or system interface is configured with an IPv6 address, and the associated discovery rule is configured to scan for the IPv6 address, the NFM-P discovers and manages the device using IPv6. To switch between IPv4 and IPv6 management of a device, you must unmanage the device, create a new discovery rule that specifies the other protocol, and then rediscover the device.

The NFM-P supports the configuration of IPv4 and IPv6 in-band and out-of-band management addresses on the same device.

Note: To use IPv6 for out-of-band management of a 7210 SAS-E, 7210 SAS-M, 7210 SAS-Mxp, 7210 SAS-R, 7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T, 7210 SAS-X, or 7705 SAR, you must configure an IPv6 address on the management port and an IPv4 address on the system interface.

The 7210 SAS-D, 7210 SAS-K, 7210 SAS-M, 7210 SAS-Mxp, 7210 SAS-R,7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T, 7210 SAS-X, and 7705 SAR support in-band management using IPv6.

8.1.5 Secure file transfers

The NFM-P supports secure and non-secure file transfers for backups, restores, software upgrades, and statistics collection. The device mediation policy determines whether FTP or SCP is used to perform file transfers to and from the managed devices. SCP requires SSH2. See Chapter 9, “Device discovery” for more information about configuring SSH2.

8.2 Device-specific commissioning information

8.2.1 7705 SAR management bandwidth

Some additional bandwidth management configuration is required on the NFM-P, if you want to transmit on in-band management on the 7705 SAR. Currently, if in-band and out-of-band management are both used, then the 7705 SAR will only transmit on the out-of-band management out-of-preference if that route exists.

You need to configure NFM-P so that each 7705 SAR with out-of-band management connectivity has the following parameter values:

- Active Management IP - In Band
- Auto Revert to Preferred - True or False
- Management IP selection - in Band Preferred

See 8.10 “To configure polling for a 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, VSR, or Wavence SM” (p. 260) for the steps on how to configure these values.

8.2.2 7210 SAS in-band and out-of-band management

The NFM-P supports the in-band and out-of-band management of 7210 SAS NEs, except for the 7210 SAS-D and 7210 SAS-K, which do not support out-of-band management. In-band management uses the system address on a network interface. Out-of-band management uses the management address on the management port. Addresses are assigned using a CLI on the console port; see 8.7 “To commission a device for NFM-P management” (p. 251).

IPv6 is supported for management port addresses on all 7210 SAS chassis types except the 7210 SAS-D, 7210 SAS-K, 7210 SAS-S, and 7210 SAS-Sx. The management port address can be IPv4 or IPv6, but not both.

IPv6 is supported for system interface addresses on the following 7210 SAS chassis types:

- 7210 SAS-D
- 7210 SAS-M
- 7210 SAS-Mxp
- 7210 SAS-R
- 7210 SAS-T
- 7210 SAS-X

Where IPv6 is supported, you can assign a system interface address using IPv4, IPv6, or both. When both IPv4 and IPv6 addresses are assigned, the address used is determined by the Management Protocol parameter configured in 9.23 “To configure a discovery rule” (p. 310).

Note: For 7210 SAS NEs, when both the management address and the system address use only IPv6, the following network functions and protocols are not supported:

- MPLS
- LDP
- Services with SDPs
- OAM tests related to services
- IPv6 peers for BGP, RIP-NG, MC-LAG, and MC-APS

The 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K support static routes, but do not support an IGP or SDPs. The example network shown in the following figure and the associated configuration steps describe the configuration required to enable 7210 SAS-E in-band management using static routes that are distributed to an IGP, which, in the example, is OSPF.

Figure 8-2 Example 7210 SAS-E in-band management network

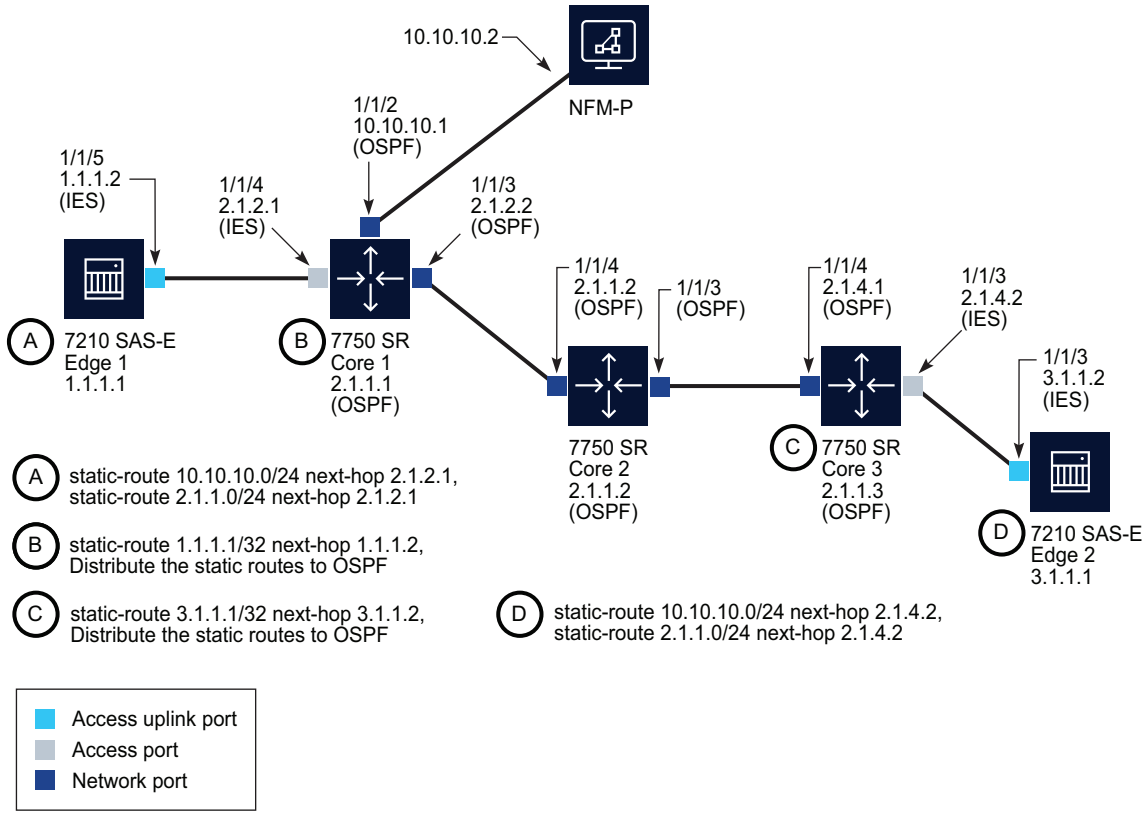

The following configuration steps are required to set up the example network:

1. Using the CLI, configure the SNMP security parameters on the devices that you need to discover. See 8.7 “To commission a device for NFM-P management” (p. 251) for more information.
2. Perform the following steps on Core 1.
   - Configure an interface on port 1/1/2 that connects to the NFM-P server interface.
   - Configure an interface on port 1/1/3, which is connected to Core 2.
   - Create an L3 interface on port 1/1/4 by creating an IES.
   - Create a static route to Edge 1.
   - Create an ASBR OSPF (area 0.0.0.0) instance on the system and enable OSPF on the interfaces to Core 2 and the NFM-P server.
   - Create a routing policy to distribute the static route to OSPF.
3. Perform the following steps on Core 2.
   - Configure an L3 interface on port 1/1/4, which is connected to Core 1.
   - Configure an L3 interface on port 1/1/3, which is connected to Core 3.

- Enable OSPF (area 0.0.0.0) on the system and the interfaces connected to Core 1 and Core 3.

4. Perform the following steps on Core 3.
   - Create an L3 interface on port 1/1/4, which is connected to Core 2.
   - Create an L3 interface on port 1/1/3, which is connected to Edge 2.
   - Create a static route to Edge 2.
   - Create an ASBR OSPF (area 0.0.0.0) instance on the system and enable OSPF on the interface to Core 2.
   - Create a routing policy to distribute the static route to OSPF.
5. Perform the following steps on Edge 2.
   - Create an IES L3 interface on port 1/1/3, which is an uplink port.
   - Create static routes that direct traffic to the IES L3 interface.
6. Perform the following steps on Edge 1.
   - Create an IES L3 interface on port 1/1/5, which is configured as an uplink port.
   - Create static routes to direct traffic to the IES L3 interface.
7. Ensure that each 7210 SAS-E can ping the network interface IP address which is configured on the NFM-P main server.
8. Ensure that the NFM-P main server can ping the system IP address of each 7210 SAS-E.
9. Configure an in-band polling policy using the NFM-P. See 8.11 “To configure polling for a 7210 SAS” (p. 262) for more information.

8.2.3 Model driven devices

Model driven devices are read-only in the NFM-P.

8.3 Workflow to commission Nokia devices

8.3.1 Stages

1

Enable and configure SNMP on the device.

a. See 8.7 “To commission a device for NFM-P management” (p. 251) for NFM-P natively-managed devices.

b. See 8.8 “To commission an OmniSwitch for NFM-P management” (p. 254) for the OmniSwitch.

2

If required, configure the source of device SNMP traps; see 8.9 “To configure the NFM-P SNMP trap listener” (p. 258) .

3

If required, use a CLI to enable NFM-P in-band management of the device by to configure a second trap destination and trap log on the device. See “GNE commissioning” (p. 246) for information about in-band management. See the device documentation for configuration information.

4

If required, establish a route for in-band traffic. For example, to configure a static route or use OSPF. See Chapter 27, “NE routing and forwarding” for more information on static routes. See Chapter 28, “Routing protocol configuration” for more information about OSPF.

5

Configure the NFM-P to use in-band polling, out-of-band polling, or in-band and out-of-band polling for the device. An NFM-P mediation policy specifies the polling interval. See Chapter 9, “Device discovery” for information about mediation policies.

a. For 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS polling configuration information, see 8.10 “To configure polling for a 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, VSR, or Wavence SM” (p. 260)

b. For 7210 SAS polling configuration information, see 8.11 “To configure polling for a 7210 SAS” (p. 262) .

c. For 7705 SAR-Hm polling configuration information, see 8.12 “To configure polling for a 7705 SAR-Hm” (p. 264) .

d. To edit polling settings for multiple managed devices, see 8.13 “To edit polling settings for multiple devices” (p. 265) .

GNE commissioning

8.4 Overview

8.4.1 Preconfiguration

A GNE requires preconfiguration for NFM-P management. When the preconfiguration is complete, the NFM-P can discover the device, as described in Chapter 9, “Device discovery” .

The NFM-P uses configurable profiles and CLI configuration scripts to discover and manage GNEs. You can use one GNE profile for multiple GNEs of the same type.

By default, only the NFM-P admin user, or a user with an assigned admin scope of command role, can manage GNE objects. To create and configure GNE objects, a non-admin user requires a scope of command role that has Create and Update/Execute permissions on packages and classes that include the following:

- genericne—to create GNE interfaces and configure profiles
- netw—to create discovery rules
- script—to manage scripts
- security.MediationPolicy and snmp.PollerManager—to create mediation policies
- trapmapper—to create alarm catalogs that map SNMP traps to NFM-P alarms

The NFM-P can discover a device as a GNE only when the following conditions are in place.

- The GNE preconfiguration is complete.
  - The System Object ID of the GNE is required for the GNE Profile, for example 1.3.6.1.4.1. When multiple GNE profiles contain possible matches for a system object ID, the profile that contains the system object ID with the longest or most specific match is chosen.
  - SNMP and/or NETCONF is configured and enabled on the device, as appropriate.
  - Telnet or SSH is enabled on the device.
- An NFM-P GNE profile for the device exists.
- An NFM-P mediation policy is configured with SNMP and NETCONF settings that match the device configuration.
- A discovery rule for the device exists.

See 8.6 “Workflow to commission GNEs” (p. 250) for the GNE commissioning workflow. See Chapter 9, “Device discovery” for information about mediation policies and device discovery.

8.4.2 NFM-P upgrade and native GNE management

If an NFM-P system that is to be upgraded manages a device as a GNE, and the new NFM-P release supports native management of the device, you must unmanage the device and delete it from the NFM-P database before the upgrade.

After the upgrade, you can use the NFM-P to discover and manage the device natively, rather than as a GNE.

Note: GNE management cannot be used to extend the release support of a device. The same release compatibility between NEs and NFM-P applies whether the device is natively managed or managed as a GNE.

8.4.3 Invoking alternate element managers

If an EMS client application for a GNE is installed on an NFM-P GUI client station, you can use a right-click GNE menu option in the network navigation tree to open the EMS client application. The OS-level command that opens the client must be specified in the associated GNE profile, or in the individual GNE configuration, to enable the function. A command in the configuration of a specific GNE overrides a command in the associated GNE profile.

If the command path is not included, the NFM-P attempts to locate the command using the paths listed in the PATH environment variable of the client station OS.

If the command that opens an EMS client accepts an NE IP address as an argument, you can specify the address using a keyword in the command entry. The NFM-P replaces the keyword with the target GNE IP address as it runs the command.

An alternate GNE manager is configurable using the GUI or OSS. See 8.14 “To prepare a GNE for NFM-P management” (p. 266) for information about configuring an alternate EMS in a GNE profile. See 8.15 “To configure an alternate EMS for a specific GNE” (p. 268) for information about specifying an alternate EMS in the GNE configuration.

8.4.4 GNE profiles

NFM-P management of a device as a GNE requires a GNE profile, which defines how the NFM-P communicates with the device. A GNE profile includes the following elements:

- a unique GNE name
- the system object ID from the device MIB
- the GNE category
- the interface types that the GNE supports
- the interfaces that can be specified as the endpoints of NFM-P physical links
- an SNMP trap management configuration that associates trap configuration and deconfiguration scripts with the profile
- CLI profile information
- an optional alarm catalog that defines the alarms that the NFM-P raises in response to SNMP traps from the device

8.5 Configuring user-defined alarms for GNEs

8.5.1 Standard alarms

The NFM-P monitors SNMP reachability and interface status, and raises a standard alarm for each of the following events:

- coldStart—The GNE restarts.
- linkDown—An interface goes out of service.

- linkUp—An interface returns to service.

8.5.2 GNE alarm catalogs

You can map GNE SNMP traps to user-defined NFM-P alarms in an alarm catalog, and associate the catalog with a GNE profile. A GNE profile can have one alarm catalog that contains up to 150 trap-to-alarm mappings.

A mapping is one of the following types:

- static—A specific SNMP trap is associated with a specific alarm.
- dynamic—The mapping includes one or more transform functions that define alarm properties based on values in SNMP trap PDUs.

Each mapping in a catalog defines an alarm that the NFM-P raises or clears when it receives a specific SNMP trap. When the NFM-P receives a high trap volume and must discard traps that it cannot process, it does not distinguish between standard and user-defined traps. To conserve system resources, Nokia recommends that you configure a GNE to send only the required traps to the NFM-P.

Traps that map to user-defined alarms require extra processing by the NFM-P and are managed in a separate, resource-limited queue. When this queue fills, the NFM-P discards some of the traps and raises an alarm. You can monitor the queue length using the NFM-P Resource Manager. When a mapping is administratively disabled, the NFM-P raises no alarm in response to an associated trap from a GNE.

GNE trap sequencing and throttling support are configurable in a GNE profile. After the NFM-P finishes throttling traps from a GNE or encounters a trap sequence error, it resynchronizes the discovered device MIBs.

A mapping includes standard elements such as the trap OID, alarm type, probable cause, and severity, but can include the following optional elements:

- trap name
- self-clearing designation—specifies that the alarm clears when a specific clearing trap is received, and has the following requirements:
  - If the severity of the raising alarm is a static value, the clearing trap must have a static mapping in the same catalog as the raising trap, and can be linked to only one raising trap.
  - If the severity of the raising alarm is defined using a transform function, the transform function must include a raising value pair and a clearing value pair.
- FDN extension, which is an alarm-name extension that can include the following:
  - static text
  - scripting functions—expressions that specify the trap PDU values to include; these allow the same alarm type to be raised in response to different traps while uniquely identifying the trap origin in the alarm name.
- additional text—used to provide information of value related to the trap event, for example, troubleshooting actions; the additional text consists of the trap OID by default, but can include the following:
  - static text
  - scripting functions—expressions that specify the trap PDU values to include; these are used to generate a more precise description of the alarm condition

- System Address and/or Interface Index varbind positions for a GNE on which to raise an alarm

Note: The FDN extension of a GNE alarm is not appended to the Alarm Name field in the NFM-P GUI, but is included in the Additional Text field. To create a filter for GNE alarms that have FDN extensions, you must filter on the Additional Text field.

You cannot filter on the Additional Text column on the dynamic alarm list, Faults tab, and Correlated Alarms tab.

When the NFM-P drops or fails to receive an SNMP trap from a GNE, the trap is lost. The NFM-P is unable to request that a GNE resend an SNMP trap.

If reliable communication via SNMP protocol is required, use “inform event” messaging instead of trap notifications.

A change to an alarm catalog or to a mapping in a catalog takes effect immediately after you commit the change.

You can use an NFM-P GUI or OSS client to configure GNE alarm catalogs and alarm mappings. The GUI supports the following methods:

- configuration forms—for object creation, modification, viewing, and deletion
- XML API script—for object creation and modification only

  A script template for alarm catalog configuration is available at the following location:

  /opt/nsp/nfmp/server/nms/sample/xmlapi/AlarmCatalogue-Template.txt

An OSS client can also retrieve a catalog or a subset of the alarm mappings in a catalog using the standard methods.

8.5.3 Transform functions

A transform function is an optional catalog component that associates one or more values in an SNMP trap PDU with an alarm property such as the alarm name, probable cause, or severity. For example, you can create a transform function that assigns an alarm severity of Critical when the value in a specific varbind is 1, Major when the value is 2, and Minor when the value is 3.

When an alarm mapping includes one or more transform functions, the NFM-P can raise multiple alarms in response to the same SNMP trap. A trap value and the associated alarm property value are specified as a value pair in a transform function. You can also specify a default alarm property value that the NFM-P assigns to an alarm when a received value is not defined in a value pair.

A transform function defines the input value type, such as integer, and the output alarm property type, such as severity. You can modify these parameters only when the transform function does not contain a value pair and is not used by a mapping.

A transform function returns an empty string when a received value is not defined in a value pair and no default alarm property value is assigned. When the transform function defines the alarm name, probable cause, or severity, the NFM-P logs an error in response and does not raise an alarm.

8.6 Workflow to commission GNEs

8.6.1 Stages

1

Create a CLI script to enable trap forwarding to each main server in the NFM-P system. See the device documentation for information about configuring SNMP trap targets. Workflows could be used for this task.

2

Create a CLI script to disable trap forwarding to each main server in the NFM-P system. See the device documentation for information about configuring SNMP trap targets.

3

Prepare the device for NFM-P discovery and management using a GNE profile; see 8.14 “To prepare a GNE for NFM-P management” (p. 266) .

4

If required, configure or modify a GNE alarm catalog for use with a GNE profile.

- Create a GNE alarm catalog; see 8.17 “To create a GNE alarm catalog” (p. 269) .
- Create a transform function for the alarm catalog; see 8.18 “To create a transform function for a GNE alarm catalog” (p. 273) .
- Add an alarm mapping to the alarm catalog; see 8.19 “To add an alarm mapping to a GNE alarm catalog” (p. 274) .

Procedures for device commissioning

8.7 To commission a device for NFM-P management

8.7.1 Commissioning tasks not covered in this procedure

Table 8-1 Location of device commissioning information

Device	Information location
Devices for model-driven management	*NSP System Administrator Guide*

CAUTION

Service Disruption

Do not apply an SNMP log filter to the NFM-P SNMP log. The NFM-P cannot manage an NE that has an SNMP log filter applied to the log used by the NFM-P, which is typically log ID 98. In rare cases, Nokia may apply filters.

8.7.2 Steps

1

Open a console window on the device.

2

Enter the following command to configure the system address:

`configure router interface system address nnn.nnn.nnn.nnn/mask` ↵

where *nnn.nnn.nnn.nnn* is the system IP address

mask is the subnet mask

3

Enter the following command to enable Telnet:

`configure system security telnet-server` ↵

4

Enter the following command to enable FTP:

`configure system security ftp-server` ↵

5

If required, enter the following command to enable SSH2:

`configure system security ssh version 2` ↵

6

Enter the following command to enable console, FTP, and SNMP access for the appropriate user account on the device:

configure system security user *user_account* access console ftp snmp ↵

where *user_account* is the appropriate user account for Telnet, FTP, and SNMP access, for example, admin

7

If required, enter one of the following commands to enable hash encryption for passwords and authentication keys during device configuration save or list operations:

a. For 7x50 devices, Release 16.0 R3 and earlier, enter the following:

configure system security hash-control read-version *read-version* write-version *write-version* ↵

where

read-version is the version of encryption accepted during read operations, for example, 1, 2, or all to indicate that both are accepted

write-version is the version of encryption used during write operations, for example, 1 or 2

Version 1 encryption uses a simple key algorithm that generates the same character string each time it hashes a specific password or authentication key.

Version 2 encryption uses a more complex key algorithm that generates a different character string each time it hashes a specific password or authentication key.

b. For 7x50 devices, Release 16.0 R4 and later, enter one of the following:

configure system security management-interface classic-cli read-algorithm *read-algorithm* ↵

configure system security management-interface classic-cli write-algorithm *write-algorithm* ↵

configure system security management-interface md-cli hash-algorithm *hash-algorithm* ↵

configure system security management-interface netconf hash-algorithm *hash-algorithm* ↵

configure system security management-interface grpc hash-algorithm *hash-algorithm* ↵

where

read-algorithm is the version of encryption accepted during read operations. Nokia recommends using hash2.

write-algorithm/hash-algorithm is the version of encryption accepted during write operations. Nokia recommends using hash2.

8

Enter the following commands in sequence to set the time zone and time:

configure system time zone *time_zone -offset_from_UCT* ↵

admin set-time *YYYY/MM/DD hh:mm:ss* ↵

where

time_zone is the appropriate time zone, for example, EST

offset_from_UCT is the offset, in hours, from Universal Co-ordinated Time, also known as Greenwich Mean Time, for example, if you specify EST, *offset_from_UCT* is -5, as EST lags UCT by five hours

YYYY/MM/DD hh:mm:ss is the current local time

9

If required, perform one of the following to enable a time protocol.

a. Enter the following command to enable NTP:

configure system time ntp server-address *server_IP_address* ↵

where

server_IP_address is the IP address of the NTP server

b. Enter the following command to enable SNTP:

configure system time sntp server-address *server_IP_address* ↵

where

server_IP_address is the IP address of the SNTP server

10

Enter the following in sequence to enable the SNMPv2 engine and to configure an SNMP community:

configure system snmp no shutdown ↵

configure system snmp packet-size 9216 ↵

configure system security snmp community *community_name* rwa version both ↵

where *community_name* is the SNMPv2 community name

Caution: When configuring a mediation policy, you need use a community string that uses rwa in order for the trap destination is set on the NE. This ensures fault conditions and alarms are reported.

Note: The command is used for the NFM-P write mediation policy. If you are using SNMPv2, you must use this mediation policy for read as well, or create another mediation policy that is also configured for rwa.
The SNMPv2 community string name rwa attributes must be enabled for the NFM-P to properly manage a device, even if the NFM-P is only used to monitor a network.

11

Enter the following commands in sequence to ensure that the device uses persistent SNMP indexes:

bof persist on ↵

```
bof save ↲
```

12

Enter the following command to save the configuration changes:

```
admin save ↲
```

13

If the device has redundant CPMs, enter the following command to synchronize the CPMs:

```
admin redundancy synchronize boot-env ↲
```

14

Enter the following command to reboot the device:

```
admin reboot now ↲
```

The device initializes with SNMP communication enabled.

15

Enter the following to clear the log ID and trap group ID:

```
configure log ↲
log-id 98 ↲
shutdown ↲
exit ↲
no log-id 98 ↲
no snmp-trap-group 98 ↲
exit all ↲
```

16

Use an NFM-P client to discover the device and to verify that the device configuration allows management of the device. See Chapter 9, “Device discovery” for information about device discovery.

END OF STEPS

8.8 To commission an OmniSwitch for NFM-P management

8.8.1 Before you begin

See the appropriate OmniSwitch documentation for more information about the CLI command syntax and SNMP.

Note: The NFM-P cannot discover an OmniSwitch that is configured with the factory default settings.

Note: You must use a direct console port connection to access an OmniSwitch for the first

time. All other management methods such as SNMP, Telnet, FTP, and HTTP, are disabled until you enable them.

8.8.2 Steps

1

Open a console window using a direct console port connection to the OmniSwitch.

2

Create a Loopback0 interface and assign an IP address to the interface by entering the following:

ip interface Loopback0 address *xxx.xxx.xxx.xxx* ↵

where

xxx.xxx.xxx.xxx is the IP address of the interface

Note: Loopback0 is the name assigned to an IP interface to identify an address that is used for network management purposes. The Loopback0 interface is not bound to any VLAN, therefore it always remains operationally active.
The Loopback0 interface name is case-sensitive. Ensure that you enter the name exactly as shown.

3

Enable SNMP sessions on the switch by entering the following:

aaa authentication snmp local ↵

4

Enable FTP, Telnet, HTTP, or SSH sessions, if required, by entering the following:

aaa authentication ftp local ↵

aaa authentication telnet local ↵

aaa authentication http local ↵

aaa authentication ssh local ↵

5

Perform one of the following to configure the required version of SNMP on the switch and the NFM-P:

Note: The OmniSwitch default user, admin, does not have SNMP access. Before the NFM-P can discover an OmniSwitch, you must create at least one user on the switch with SNMP access.

a. For SNMP v2c:
   1. Configure an SNMP v2 user by entering the following:

      user *user_name* password *password* no auth ↵

where

user_name is a username that corresponds to an SNMP v2 user who the NFM-P can identify; Nokia recommends that you use the name sam, which is the NFM-P default name

password is a password associated with the username; the password is between 8 and 47 characters

Note:

If you need to use a different SNMPv2 default user name, create an SNMPv2 default user name on the NFM-P. See the procedure to create a default SNMPv2 OmniSwitch user in the *NSP System Administrator Guide* for information about creating an SNMPv2 default user name.

2. Configure SNMP v2 trap forwarding to the NFM-P by entering the following:

   snmp station *xxx.xxx.xxx.xxx v2 user_name* ↵

   where

   xxx.xxx.xxx.xxx is the IP address of the NFM-P main server

   user_name is the username that you created in 1

   Note:

   Trap forwarding configuration occurs automatically when the NFM-P discovers a switch and uses the default SNMP v2 user name sam or the user name configured, if any.

3. Configure an SNMP security level that allows the switch to accept all SNMP queries by entering the following:

   snmp security no security ↵

4. Configure an SNMP v2 community string by entering the following:

   snmp community map *community_string* user *user_name* ↵

   where

   community_string is the name of an SNMP v2 community string that the NFM-P can identify

   user_name is the SNMP v2 username created in 1

5. Create a mediation security policy on the NFM-P that uses a community string that matches the string created in 4 . See 9.17 “To configure device mediation” (p. 301) for information about creating a mediation security policy.

6. Create a discovery rule on the NFM-P to discover the switch and reference the policy created in 5 . See 9.23 “To configure a discovery rule” (p. 310) for information about creating a discovery rule.

b. For SNMP v3:

1. Configure an SNMP v3 user on the switch by entering the following:

   user *user_name* password *password security_level* ↵

   where

   user_name is a username that matches an SNMP v3 USM username configured on the NFM-P

password is a password associated with the username; the password is between 8 and 47 characters. The password is the plain text ASCII MD5/SHA authentication key and DES privacy key.

security_level is MD5, MD5 + DES, SHA, or SHA + DES

2. Configure SNMP v3 trap forwarding to the NFM-P by entering the following:

   `snmp station xxx.xxx.xxx.xxx v3 user_name` ↵

   where

   xxx.xxx.xxx.xxx is the IP address of the NFM-P main server

   user_name is the username created in 1

   Note:

   Trap forwarding occurs automatically when the NFM-P discovers a device with a username that matches the SNMP v3 USM username specified in the NFM-P mediation policy.

3. Configure the SNMP v3 switch security option that you need by entering the following:

   `snmp security security_option` ↵

   where *security_option* is one of the security options described in the following table

Option	Description
no security	All SNMP queries are accepted.
authentication set	Includes: • SNPM v1 and v2 Gets • Non-authenticated v3 Gets and Get-Nexts • Authenticated v3 Sets, Gets, and Get-Nexts • Encrypted v3 Sets, Gets, and Get-Nexts
authentication all	Includes: • Authenticated v3 Sets, Gets, and Get-Nexts • Encrypted v3 Sets, Gets, and Get-Nexts
privacy set	Includes: • Authenticated v3 Gets and Get-Nexts • Encrypted v3 Sets, Gets, and Get-Nexts
privacy all (default)	Includes: • Encrypted v3 Sets, Gets, and Get-Nexts
traps only	Includes: • All SNMP requests are rejected

4. Create an SNMP v3 user on the NFM-P using the NE User Configuration manager. See the section on NE user and device security in the *NSP System Administrator Guide* for information about NE user configuration.
   - Enable SNMP to give the SNMP v3 user SNMP access.
   - Choose a username that matches the name created on the switch in 1 .

- Choose the same SNMP v3 authentication protocol, privacy protocol, and password that is configured on the switch.

5. Create an SNMP v3 mediation security policy. See 9.17 “To configure device mediation” (p. 301) for information about configuring a mediation security policy.
   - Choose the SNMP v3 (USM) security model option.
   - Choose a username that matches the name created on the switch in 1 .
6. Create a discovery rule that uses the mediation security policy created in 5 . See 9.23 “To configure a discovery rule” (p. 310) for information about creating discovery rules.

6

Use an NFM-P client to discover the switch and to verify that the switch configuration allows you to manage the switch.

END OF STEPS

8.9 To configure the NFM-P SNMP trap listener

8.9.1 Purpose

As an SNMP trap destination, the NFM-P listens to certain SROS event logs. Perform this procedure to configure the NFM-P as an SNMP trap listener to the SNMP trap sources, which are certain SROS event logs associated with log ID 98. If no SNMP log ID on a managed SROS device matches the log ID in the NFM-P configuration, the NFM-P recreates the log ID and SNMP trap group on the device, which results in a single log file that receives traps from the following sources:

- main
- security—always included as an SNMP source
- change—supported by the NFM-P for the 7950 XRS and for the 7750 SR

CAUTION

Service Disruption

This procedure requires a restart of each NFM-P main server, which is service-affecting. Perform this procedure only during a scheduled maintenance window.

8.9.2 Steps

1

If the NFM-P is deployed in a standalone configuration, go to Step 3 .

2

Perform the following steps to stop the standby main server.

1. Log in to the standby main server station as the nsp user.
2. Open a console window.

3. Enter the following:

   ```
   bash$ cd /opt/nsp/nfmp/server/nms/bin ↵
   ```

4. Enter the following to stop the main server application:

   ```
   bash$ ./nmsserver.bash stop ↵
   ```

5. Enter the following to display the main server status:

   ```
   bash$ ./nmsserver.bash appserver_status ↵
   ```

   The command displays a status message.

6. The main server is stopped when the command displays the following status message:

   ```
   Application Server is stopped
   ```

   If the command displays a different message, wait 5m and repeat Step 2 5 . Do not proceed to the next step until the server is stopped.

3

Perform the following steps on the primary main server station or the standby main server station:

1. Log in to the main server station as the nsp user.
2. Open the /opt/nsp/nfmp/server/nms/config/nms-server.xml file using a plain-text editor.

   Caution: Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file. Modifying the nms-server.xml file can have serious consequences that can include service disruption.

3. Locate the following section:

   ```
   <snmp
   ip="IP_address"
   port="nnn"
   trapLogId="98"
   ```

   where

   IP_address is the IP address of the NFM-P main server

   nnn is the SNMP port number

   This default configuration applies to traps from main, security, and change, if available.

4. Perform one of the following:

   To restrict the trap source to main only, add the following line:

   ```
   logIdSourceBits="80"
   ```

   To restrict the trap source to main and change, add the following line:

   ```
   logIdSourceBits="A0"
   ```

5. Save and close the nms-server.xml file.
6. Open a console window.
7. Navigate to the /opt/nsp/nfmp/server/nms/bin directory.
8. Enter the following to restart the main server:

```
bash$ ./nmsserver.bash force_restart ↵
```

The main server restarts, and the configuration change takes effect.

4

Close the console windows.

END OF STEPS

8.10 To configure polling for a 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, VSR, or Wavence SM

8.10.1 Purpose

Perform this procedure to configure polling and the schedule resynchronizing settings for a specified NE type. You can also optionally configure the system initialization commands configured in the BOF such as the primary/secondary/tertiary DNS addresses, Lawful Intercept settings, and the scaling profile that specifies the type of network scale usage that is applied to the hardware configured in the network. The parameters that display vary depending on the NE type, release, and the settings of other parameters.

8.10.2 Steps

1

Choose Equipment or Routing from the navigation tree view selector, right-click on a supported NE, and choose Properties. The Network Element (Edit) form opens.

2

Record the system IP address and the management IP address.

3

If required, configure the Location parameters.

4

Click on the Polling tab.

5

Enable or disable the Scheduled Polling parameter as required. Scheduled polling is configured using the Mediation configuration form. See 9.17 “To configure device mediation” (p. 301) for more information.

6

If required, configure the Redundant Synchronization Mode parameter.

7

Verify that the Persistent SNMP Indices parameter is set to true to ensure persistent SNMP indices are used.

8

Configure the parameters in the Bof Configuration panel as required.

Note: To configure LI parameters, or view LI configuration information, you require LI user privileges. Right-click on a discovered device in the navigation tree and choose Properties from the contextual menu, and click on the LI Configuration Status tab.

Note: If you enable encryption, you must run bof save.

9

View the read-only parameters to determine the current polling status:

- Resync Status indicates whether the last poll was successfully completed.
- Last Resync Start Time and Last Resync End Time indicate the start and finish of the last poll.
- Scheduled Resync Status indicates whether the last scheduled poll was successfully completed.
- Last Scheduled Resync Start Time and Last Scheduled Resync End Time indicate the start and finish of the last scheduled poll.

10

Click on the Management tab and configure the required parameters.

The Auto Revert to Preferred parameter is configurable when the Management IP Selection parameter is set to Out Of Band Preferred or In Band Preferred.

Note: To avoid a communications outage between the NFM-P and managed NEs, the parameter settings in the Management Preference panel must match the parameter settings in the Notifications Preferred Management panel. See the Management IP Selection parameter description for more information.
If you want to transmit on in-band management on the 7705 SAR when both in-band and out-of-band management are both used, configure the Active Management IP, Auto Revert to Preferred, and Management IP Selection parameters as described in 8.2.1 “7705 SAR management bandwidth ” (p. 241) .

11

Save your changes and close the form.

END OF STEPS

8.11 To configure polling for a 7210 SAS

8.11.1 Steps

1

Choose Equipment or Routing from the navigation tree view selector, right-click on a 7210 SAS, and choose Properties. The Network Element form opens.

2

Record the system IP address and the management IP address.

3

If required, configure the Location parameters.

4

Click on the Polling tab, then the General tab.

5

Enable or disable the Scheduled Polling parameter. Scheduled polling is configured using the Mediation configuration form. See 9.17 “To configure device mediation” (p. 301) for more information.

6

Verify that the Persistent SNMP Indices parameter is set to true to ensure that persistent SNMP indices are used.

7

Configure the parameters in the Bof Configuration panel as required.

8

View the read-only parameters to determine the current polling status:

- Resync Status indicates whether the last poll was successfully completed.
- Last Resync Start Time and Last Resync End Time indicate the start and finish of the last poll.
- Scheduled Resync Status indicates whether the last scheduled poll was successfully completed.
- Last Scheduled Resync Start Time and Last Scheduled Resync End Time indicate the start and finish of the last scheduled poll.

9

Click on the Management tab and configure the required parameters.

The Auto Revert to Preferred parameter is configurable when the Management IP Selection parameter is set to Out Of Band Preferred or In Band Preferred.

Note: The Active Management IP, Auto Revert to Preferred, and Management IP Selection parameters are not supported on the 7210 SAS-D or 7210 SAS-K. The 7210 SAS-D and 7210 SAS-K support only in-band management.

10

Click on the 7210 BOF tab to configure uplinks in the BOF for port redundancy.

Note: If the image path and configuration file path are local, you do not need to configure the IP address and routing information for uplink A and uplink B. You can optionally obtain IP parameters using DHCP when you configure a value of 0 for the uplink port IP address. The DHCP server should be configured to provide the IP address and the default gateway information that is used to reach the server where the image and configuration files are stored.

11

Select an uplink port in the Uplink A panel.

12

Configure the required parameters in the Uplink A panel.

13

Select an uplink port in the Uplink B panel.

14

Configure the required parameters in the Uplink B panel.

15

Click on the Uplink Routes tab and click Create. The Uplink Route Configuration form opens.

16

Configure the required parameters and click Apply.

17

Repeat Step 16 to add another uplink route. You can add up to ten routes for each uplink.

18

Save your changes and close the form.

END OF STEPS

8.12 To configure polling for a 7705 SAR-Hm

8.12.1 Purpose

Perform this procedure to configure the SNMP polling interval for 7705 SAR-Hm NEs. The default polling cycle is one day (1440 minutes). You can change the polling to up to once a week.

8.12.2 Steps

1

If the NFM-P is deployed in a standalone configuration, go to Step 3 .

2

Perform the following steps to stop the standby main server.

1. Log in to the standby main server station as the nsp user.
2. Open a console window.
3. Enter the following:

   ```
   bash$ cd /opt/nsp/nfmp/server/nms/bin ↲
   ```

4. Enter the following to stop the main server application:

   ```
   bash$ ./nmsserver.bash stop ↲
   ```

5. Enter the following to display the main server status:

   ```
   bash$ ./nmsserver.bash appserver_status ↲
   ```

   The command displays a status message.

6. The main server is stopped when the command displays the following status message:

   ```
   Application Server is stopped
   ```

   If the command displays a different message, wait 5m and repeat Step 2 5 . Do not proceed to the next step until the server is stopped.

3

Perform the following steps on the primary main server station or the standby main server station:

1. Log in to the main server station as the nsp user.
2. Open the /opt/nsp/nfmp/server/nms/config/nms-server.xml file using a plain-text editor.

   Caution: Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file. Modifying the nms-server.xml file can have serious consequences that can include service disruption.

3. Locate the following parameter:

   ```
   <snmp reachabilityCheckInterval="
   ```

4. Modify the number of minutes as needed.
5. Save and close the nms-server.xml file.
6. Open a console window.

7. Navigate to the /opt/nsp/nfmp/server/nms/bin directory.
8. Enter the following to restart the main server:

   ```
   bash$ ./nmsserver.bash force_restart ↵
   ```

   The main server restarts, and the configuration change takes effect.

4

Close the console windows.

END OF STEPS

8.13 To edit polling settings for multiple devices

8.13.1 Purpose

You can use the list of managed devices from the Discovery Manager Resync Status tab to modify polling settings for a device or devices; for example, when you want to enable or disable polling on numerous managed devices.

8.13.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click on the Resync Status tab, select multiple devices from the list, and click Properties. The Network Element (Multiple Instances) form opens.

Note: You can only configure those parameters that are common to all devices selected from the list.

3

Click on the Polling tab and enable or disable the Scheduled Polling parameter.

4

Save your changes and close the forms.

END OF STEPS

Procedures for GNE commissioning

8.14 To prepare a GNE for NFM-P management

8.14.1 Purpose

Perform this procedure to prepare a GNE, which is typically a non-proprietary device, for NFM-P discovery and management.

8.14.2 Steps

1

Preconfigure the device using a CLI:

1. Enable FTP.
2. Enable Telnet.
3. Enable SNMP and/or NETCONF, as appropriate.
4. Configure at least one SNMPv2c or SNMPv3 community.
5. Set the SNMP PDU size to 9216.
6. Enable persistent SNMP indexes.

2

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

3

Click Create and choose Generic NE Profile. The Generic NE Profile (Create) form opens.

4

Configure the required parameters.

The Sys Object ID parameter must begin with the following, which identifies the SMI enterprises subtree defined in RFC 1213: .1.3.6.1.4.1

Do not enclose the Alternate Element Manager value in quotation marks; the NFM-P automatically encloses the value in quotation marks at run time.

You cannot modify the Generic NE Type parameter after the GNE profile is created; you must delete the GNE profile and create a profile using the new Generic NE Type parameter value.

Note: To specify the target GNE IP address in the Default Element Manager URL or as a Default Alternate Element Manager command argument, use the %IP% keyword.

5

If applicable, select a driver to associate with the GNE profile from the Driver Module panel.

6

Click on the CLI Profile tab and configure the required parameters.

7

If required, configure the login confirmation parameters.

The Prompt and Answer parameters are configurable when the Enable Confirm Prompt parameter is selected.

8

If required, configure the second level login security parameters.

The Enable Login Command and Enable Login Prompt parameters are configurable when the Enable Second Login parameter is selected.

9

To configure an interface for the GNE, perform the following steps:

1. Click on the Interface Types tab and click Add. The Select Generic NE Interface Type form opens.
2. Choose an interface type and click OK. The form closes.

10

To associate one or more routing MIBs with the GNE profile, perform the following steps.

Note: After you discover the device using the GNE profile, you can use the Routing view in the network navigation tree to view the discovered routing objects.
You must configure the GNE interface types before you associate a routing MIB with a GNE profile. Otherwise, you must perform a customized NE resync of all MIBs and select the Ignore Timestamps parameter.

1. Click on the Other MIBs tab and click Add. The Select Generic NE Routing MIB form opens.
2. Select one or more MIBs and click OK.

11

Click Apply and confirm your actions. The Trap Configuration tab is enabled.

12

Click on the Trap Configuration tab and perform one of the following, if required select a script in the Trap Configuration Script panel. During device discovery, the NFM-P runs this script on the GNE to enable trap forwarding to the NFM-P.

13

Select a script in the Trap De-Configuration Script panel. During device discovery, the NFM-P runs this script on the GNE to disable trap forwarding to the NFM-P.

14

Select an alarm catalog in the Alarm Catalogue panel.

15

Save your changes and close the forms.

END OF STEPS

8.15 To configure an alternate EMS for a specific GNE

8.15.1 Purpose

Perform this procedure to specify an alternate EMS client application for a specific NE.

Note: The EMS that you specify overrides an EMS specified in the associated GNE profile.

8.15.2 Steps

1

Right-click on a GNE object in the network navigation tree and choose Properties. The GNE properties form opens.

2

Configure the Alternate Element Manager parameter.

Note: Do not enclose the Alternate Element Manager value in quotation marks; the NFM-P automatically encloses the value in quotation marks at run time.
To specify the target GNE IP address as an Alternate Element Manager command argument, use the %IP% keyword.

3

Save your changes and close the form.

END OF STEPS

8.16 To modify a GNE profile

8.16.1 Purpose

Perform this procedure to change a GNE profile configuration.

8.16.2 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Profile (Generic NE), select a GNE profile in the list, and click Properties. The Generic NE Profile form opens.

The following tabs contain configurable parameters:

- General—contains parameters that identify the GNE type
- CLI Profile—contains parameters that define the console window prompts and commands
- Trap Configuration—contains parameters that define the SNMP trap management configuration
- Interface Types—allows the association of multiple types of interfaces with the GNE profile

  Caution: If you modify the list of interfaces on the Interface Types tab, you must resynchronize the entire GNE, as described in Step 4 .

3

Modify the required parameters for the GNE profile, click OK, and confirm your actions. The form closes.

4

If you modified the list of interfaces on the Interface Types tab, perform the following steps:

1. Right-click on the GNE in the NFM-P topology map and choose Resync→Customized Resync. The Resync Site(s) form opens.
2. Click Next and select the Ignore Timestamps parameter.
3. Click Finish and close the Resync Site(s) form. The NFM-P resynchronizes the GNE.

END OF STEPS

8.17 To create a GNE alarm catalog

8.17.1 Purpose

Perform this procedure to create an alarm catalog for use with a GNE profile.

8.17.2 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Click Create and choose Generic NE Alarm Catalogue. The Generic NE Alarm Catalogue (Create) form opens.

3

Configure the required parameters.

4

If required, perform Step 3 to Step 12 of 8.18 “To create a transform function for a GNE alarm catalog” (p. 273) to create a transform function for the catalog.

5

Click on the Mappings tab and click Create. The Generic NE Alarm Mapping (Create) form opens.

6

Configure the required parameters.

7

Perform one of the following:

a. Configure the Alarm Name parameter to specify a static alarm name.

b. Specify a dynamic alarm name:

   1. Enable the Specify Transform Function parameter.
   2. Select a transform function beside the Transform Varbind Function parameter.
   3. Configure the Varbind Position parameter.

8

Perform one of the following:

a. Select a probable cause beside the Probable Cause parameter to specify a static probable cause.

b. Specify a dynamic probable cause:

   1. Enable the Specify Transform Function parameter.
   2. Select a transform function beside the Varbind Transform Function parameter.
   3. Configure the Varbind Position parameter.

9

Perform one of the following:

a. Specify a static severity by configuring the Severity parameter.

b. Specify a dynamic severity:

Note: When you specify the use of a transform function for the alarm severity, the Mapping Type changes to Raising/Clearing. You cannot create a clearing mapping for this type of mapping; you must use a transform function to clear an alarm when the alarm severity is defined using a transform function.

1. Enable the Specify Transform Function parameter.
2. Select a transform function beside the Varbind Transform Function parameter.
3. Configure the Varbind Position parameter.

10

Configure the varbind positions for specific objects on a GNE in the Alarmed Object panel:

- System Address Varbind Position
- Interface Index Varbind Position

Note: In order to configure the System Address Varbind Position and Interface Index Varbind Position parameters, you must set the Alarm Name parameter to GNE MTIE Alarm.

11

Click OK and confirm your actions. The form closes.

12

Repeat Step 5 to Step 11 to create an additional raising alarm mapping, if required.

13

Perform one of the following:

a. If you need to create a clearing alarm, continue to Step 14 .

b. If you do not need to create a clearing alarm, go to Step 20 .

14

Select a raising alarm mapping and click Create Clearing. The Generic NE Alarm Mapping form opens.

Note: A raising alarm mapping is indicated by the word Raising in the Mapping Type list column.
You can create a clearing alarm mapping only when the corresponding raising alarm is in the same alarm catalog.
You can associate a clearing alarm mapping with only one raising alarm mapping.
You cannot create a clearing mapping for a mapping that uses a transform function to define the alarm severity; you must use a transform function to clear an alarm when the alarm severity is defined using a transform function.

15

Configure the required parameters:

Note: A static clearing alarm mapping inherits the following values from the associated raising alarm:

- Alarm Name—must match raising alarm Alarm Name
- Probable Cause—must match raising alarm Probable Cause
- FDN Extension—resulting text string must match text string generated by raising alarm mapping
- Additional Text—resulting text string must match text string generated by raising alarm mapping

  See 8.17 “To create a GNE alarm catalog” (p. 269) for information about modifying the Alarm Name or Probable Cause value.

  The explicit FDN Extension and Additional Text values can differ from the values in the raising mapping, but the generated text strings must match. For example, if the object name is in varbind 2 of the raising trap and in varbind 3 of the clearing trap, the parameter values name different varbinds but the script output is identical.

  The alarm severity in a clearing alarm mapping is set to Cleared and cannot be changed.

16

If required, specify a dynamic alarm name:

1. Enable the Specify Transform Function parameter.
2. Select a transform function beside the Varbind Transform Function parameter.
3. Configure the Varbind Position parameter.

17

If required, specify a dynamic probable cause:

1. Enable the Specify Transform Function parameter.
2. Select a transform function beside the Varbind Transform Function parameter.
3. Configure the Varbind Position parameter.

18

Click OK and confirm your actions. The form closes.

19

Repeat Step 14 to Step 18 to create additional clearing alarm mappings, if required.

20

Save your changes and close the forms.

END OF STEPS

8.18 To create a transform function for a GNE alarm catalog

8.18.1 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Click Create and choose Create Generic NE Alarm Catalogue. The Generic NE Alarm Catalogue form opens.

3

Click on the Transform Functions tab and click Create. The Generic NE Alarm Transform Function form opens

4

Configure the required parameters.

5

If you enabled the Specify Default Out Value parameter, select a value beside the Default Out Value parameter.

6

Click on the Pairs tab and click Create. The Pair form opens.

7

Configure the In Value parameter.

8

Select a value beside the Out Value parameter.

9

Click OK and confirm your actions. The form closes.

10

Repeat Step 6 to Step 9 to add another value pair, if required.

11

Click OK and confirm your actions. The form closes.

12

Repeat Step 3 to Step 11 to create another transform function, if required.

13

Save your changes and close the forms.

END OF STEPS

8.19 To add an alarm mapping to a GNE alarm catalog

Note: A modification to a GNE alarm catalog takes effect immediately after you commit the change.

8.19.1 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Alarm Catalogue (Trap to Alarm Mapper), select an alarm catalog, and click Properties. The Generic NE Alarm Catalogue form opens.

3

If required, configure the Description and Version parameters.

4

Perform Step 5 to Step 18 in 8.17 “To create a GNE alarm catalog” (p. 269) .

END OF STEPS

8.20 To delete a GNE alarm catalog

Note: You cannot delete an alarm catalog that is associated with a GNE profile. You must remove the alarm catalog from each associated GNE profile before you can delete the catalog.

Note: When you delete a GNE alarm catalog, you delete the alarm mappings that the catalog contains.

8.20.1 Steps

1

Choose Administration→Generic NE Manager from the NFM-P main menu. The Generic NE Manager form opens.

2

Choose Generic NE Alarm Catalogue (Trap to Alarm Mapper), select a GNE alarm catalog, and click Properties. The Generic NE Alarm Catalogue form opens.

3

Click on the Generic NE Profiles tab and perform one of the following:

a. If no profiles are listed, go to Step 4 .

b. If profiles are listed, perform the following steps to remove the profile associations:

1. Select a profile and click Properties. The Generic NE Profile form opens.
2. Click on the Trap Configuration tab and click Clear in the Alarm Catalogue panel.
3. Click OK and confirm your actions. The form closes.
4. Repeat 1 to 3 for each listed profile.

4

Close the Generic NE Alarm Catalogue form.

5

Select the alarm catalog and click Delete.

6

Confirm your actions and close the form.

END OF STEPS

9 Device discovery

Discovering devices using the NFM-P

9.1 Overview

9.1.1 Functional description

The NFM-P discovers devices using SNMP and stores the device properties in the NFM-P database. To discover one or more devices in your network, you create a discovery rule and then scan the network for devices according to the IP addresses or address ranges specified in the discovery rule.

A discovery rule contains one or more rule elements that specify which devices or subnets are to be included in, or excluded from, the discovery process. For example, you can configure one rule element to discover a subnet, and another to exclude specific IP addresses from the subnet.

The NFM-P periodically scans the network for new devices according to a discovery rule scan interval. By default, the scan interval specified in the NFM-P mediation configuration applies to all NFM-P discovery rules. However, in a discovery rule, you can specify a scan interval that overrides the scan interval in the mediation configuration. For example, to reduce the amount of network-management traffic created by the periodic scans in a busy subnet, you can increase the scan interval in the discovery rule for the subnet.

When the system IP address is used to discover a device, the management is in-band. When the management IP address is used to discover a device, the management is out-of-band. See Chapter 8, “Device commissioning and management” for information about in-band and out-of-band management.

Note: The NFM-P does not attempt to discover tests or test suites that are configured on a device using a CLI.

Note: If a discovered chassis exceeds the NFM-P license count during the execution of a discovery rule, the NFM-P marks the device as Unmanaged.

Note: After you update the NFM-P license to accommodate an additional chassis, you must manage the device manually. See 9.27 “To manage, suspend, or unmanage a device” (p. 316) for information.

Note: If you change the system name of a device, you must unmanage the device and then remanage the device in order for the NFM-P to recognize the new system name. See 9.27 “To manage, suspend, or unmanage a device” (p. 316) for information.

Note: The NFM-P blocks the discovery of mixed mode SROS devices.

9.1.2 Device discovery using IPv6

The NFM-P supports the discovery of devices that use IPv6 in-band or out-of-band management IP addresses. In order for the NFM-P to discover and manage a device that uses IPv6, the device

must have an IPv6 address on the management port, system interface, or both. The NFM-P main server must also be given an IPv6 address during installation.

The IP version that the NFM-P uses to discover a device is specified in each discovery rule. If the NFM-P discovers both IPv4 and IPv6 addresses on the system interface of a device, it discovers the device using the address that corresponds to the IP version specified in the discovery rule.

9.1.3 Dying Gasp trap target support for IXR-E

NSP automatically enables Dying Gasp trap target on IXR-E when new IXR is discovered by NSP. If NE manual changes are made on Dying Gasp trap target after discovery, the IXR-E will need to be unmanaged and remanaged for NSP to enable the Dying Gasp trap target again.

Note: The discovery of IXR-E node in more than 3 NFM-P is not possible due to node limitation on setting Dying Gasp. If a user wants to discover IXR-E node on 4th NFM-P, they will need to either remove the feature of Dying Gasp in an already discovered node (from node side) or they will need to unmanage the node which is already discovered NFM-P.

9.2 Device SNMP management

9.2.1 SNMP overview

Simple Network Management Protocol, or SNMP, is an application-layer protocol that provides a message format to facilitate communication between SNMP managers and agents. SNMP provides a standard framework for NE monitoring and management from a central location.

An SNMP manager controls and monitors the activities of network hosts that use SNMP. An SNMP manager uses a get operation to obtain a value from an SNMP agent, and uses a set operation to store a value in the agent. The manager uses definitions from a management information base, or MIB, to perform operations on the managed device, for example, retrieving data values, replying to requests, and processing SNMP notifications called traps.

SNMPv1 and SNMPv2c provide no security, authentication, or encryption. Without authentication, an unauthorized user is able to perform SNMP network management functions and eavesdrop on management information as it passes from one system to another.

SNMPv3 requires that an authentication and encryption method such as SSH is assigned to each user for validation by the NE. SNMPv3 authentication and encryption enable an NE to validate the system that issues an SNMP message and to determine whether another system has tampered with the message. When creating an SNMPv3 NE user on the NFM-P, ensure that the NEs support a valid combination of authentication and privacy protocols. See 9.11 “To enable SNMPv3 management of a device” (p. 291) for more information.

For information about device-specific SNMP support, see the SNMP chapter of the appropriate *System Management Guide* for the device. For information about SSH security, see 9.6 “Configuring SSH security on devices” (p. 283) .

9.2.2 SNMP packet size considerations for device discovery

The network infrastructure carrying traffic between the NFM-P main and auxiliary servers and the managed NEs must support packet fragmentation and reassembly when the PDU size is greater than the MTU of the network path. The 7210 SAS, 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS

require an SNMP PDU size of 9216 bytes to be configured and fragmentation support is required in the network path between the NFM-P and the NE.

Consider the following:

- When an intermediate network device receives SNMP traffic, it must be able to process fragmentation of packets. If the SNMP packet exceeds the MTU size of the intermediate device, and the device does not support fragmentation of packets, the packets may be dropped and device discovery may fail.
- Verify that data payloads of the SNMP PDU size can travel from the managed devices to the NFM-P by using a CLI to ping the IP address of the NFM-P main server, using a packet size of the configured SNMP PDU.

9.2.3 SNMP streaming

During the initial discovery of supported 7250 IXR platforms, or a 7450 ESS, 7750 SR, or 7950 XRS, the NFM-P enables SNMP streaming on the device, if it is not currently enabled. SNMP streaming facilitates the bulk transfer of MIB-based configuration data using a streaming mechanism. SNMP streaming may substantially reduce the time that operations such as NE discovery and resynchronization require in a high-latency network. See 9.14 “To enable or disable SNMP streaming on an NE” (p. 299) for information about enabling and disabling SNMP streaming on an NE.

9.3 Device management states

9.3.1 Description

A device can be in one of the following management states:

- discovered—the NFM-P manages the device
- unmanaged—an operator has unmanaged the formerly managed device; the management information is removed from the NFM-P database. See see 9.3.2 “Unmanaging and deleting devices” (p. 280) for more information
- suspended—an operator has suspended management of the device; see 9.3.3 “Suspending device management” (p. 280) for more information

NEs not associated with a discovery rule

After a discovery rule is deleted, any managed devices associated with the discovery rule are not subsequently included in certain network administration tasks such as NE version change detection.

In the event that a discovery rule is deleted, perform 9.28 “To associate a device with a discovery rule” (p. 317) to ensure that the managed devices associated with the discovery rule are included in subsequent network administration activities,

9.3.2 Unmanaging and deleting devices

CAUTION

Service Disruption

Unmanaging or deleting a device results in the loss of all historical AA statistics for the device.

Using the NFM-P to unmanage a device excludes the device from the managed network, but a reference to the device remains in the NFM-P discovery system. The unmanage function may be used for unusual conditions such as when the NFM-P requires a complete refresh of device data because of data corruption. Unmanaging a device results in a loss of management data for the selected device, which includes, but is not limited to, the following:

- object names and descriptions
- statistics
- alarms
- physical links
- policies
- script results
- scheduled activities
- device configuration backups

Deleting a device results in the complete loss of management data for the device and completely removes the device from the managed network.

9.3.3 Suspending device management

CAUTION

Service Disruption

When you suspend the management of a device, the NFM-P raises a critical alarm against the device.

When you need to exclude a device from NFM-P management but do not want to lose the NFM-P information about the device, for example, during a maintenance period that may generate an exceptional number of SNMP traps, you can use a GUI or OSS client to suspend the management of the device.

When you suspend management of a device, the NFM-P removes the associated trap targets from the device configuration, and blocks the following operations:

- SNMP deployment and trap handling; all traps from the device are dropped
- NETCONF deployments and event subscriptions
- statistics collection
- automated FTP and other file transfers
- GUI or OSS client configuration of device parameters; a configuration attempt results in a failed deployment
- scheduled or periodic operations that the NFM-P initiates, for example:
  - connectivity checks
  - resynchronizations
  - device configuration backups

- device software upgrades
- OAM test and test suite execution

Note: You cannot suspend the management of a device while a software upgrade or configuration backup is in progress on the device.

By default, the NFM-P does not raise alarms against suspended NEs. You can override the default setting using a parameter in the main-server configuration file. Contact technical support for more information.

You can use the NFM-P to perform the following on a suspended device:

- open a Telnet or SSH CLI session
- execute CLI scripts
- open an FTP session
- configure NFM-P parameters that affect the device management but are not deployed, such as the in-band or out-of-band management preferences

Note: By default, the NFM-P suppresses alarms for suspended NEs. If you attempt to modify an OAM test or test suite that is deployed to a suspended NE, the NFM-P raises an alarm. The alarm does not automatically clear when the NE is remanaged.

As shown in Figure 9-1, “GUI representation of suspended device” (p. 281) , when a device is suspended, the device icons in the network navigation tree and physical topology map are colored purple. When a suspended device is in a group other than the Network group, only the lower right corner of the Network group icon is colored purple.The following icons are also colored purple to indicate that at least one device is suspended:

- parent group icon in the network navigation tree Equipment and Routing views
- parent group icon in the physical topology map

Note: The group icon colors are prioritized; when a group contains one or more unreachable NEs and one or more suspended NEs, the group icon is colored red.

Figure 9-1 GUI representation of suspended device

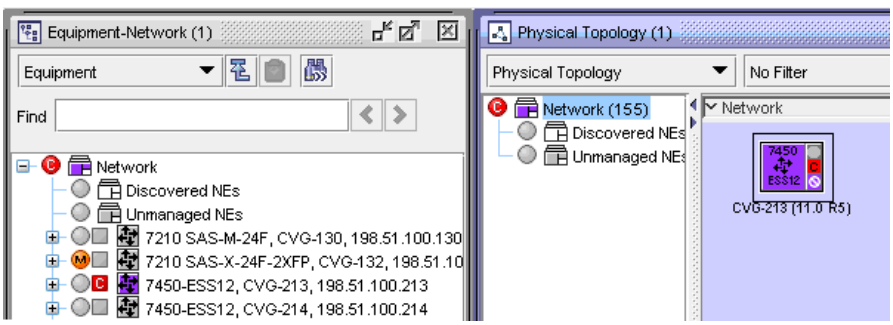

When you return a suspended device to the managed state, the NFM-P rediscovers the device to capture events such as device reboots or software release changes, and initiates a full device resynchronization based on the response to a trap sequence ID check.

Note: The NFM-P also checks the device management parameters to adapt to a switch from one type of management to another, for example, from in-band to out-of-band.

9.4 Using multiple management interfaces

9.4.1 Overview

A default SNMP management interface on an NFM-P main or auxiliary server is defined during NFM-P system deployment. To segregate the management traffic of different groups of NEs, for example, IPv4 and IPv6 networks, you can configure additional main and auxiliary server management interfaces.

After you create the required interfaces, you must define the management network and subnets to associate with the interfaces. The NFM-P subsequently attempts to discover the NEs in each network and subnet.

Note: In a segregated network, the same NFM-P management interface must be used for the in-band and out-of-band management of an NE.

When an NE is managed in-band and out-of-band using a non-default interface, you must assign the in-band and out-of-band subnets to the interface.

9.5 Post-discovery actions on discovered NEs

9.5.1 Overview

In a discovery rule, you can create and specify a post-discovery action to perform on the successfully discovered NEs included in the rule. The NFM-P performs the post-discovery action automatically at the end of rule execution.

A post-discovery action requires an NFM-P control script that manages the auto-configuration actions performed on discovered NEs.

For example, an operator specifies a control script as a post-discovery action in a discovery rule for a subnet. The control script co-ordinates the execution of device-specific child control scripts. The child control scripts call XML API scripts that each configure a group of discovered NEs, such as aggregation ring or IP backhaul NEs. After the NFM-P executes the discovery rule, it invokes the control script to auto-configure the successfully discovered NEs in the subnet.

Note: Because a post-discovery action is intended to be autonomous, the NFM-P displays an error message and halts a post-discovery action that requires user interaction.

Workflows could be used to perform post-discovery actions..

From the Managed State tab of the Discovery Manager form, you can do the following:

- List the post-discovery action associated with each managed NE.
- View the status of the script execution as it occurs.
- View the script run result.

The Node Discovery Control form of an NE displays additional information and allows post-discovery action management. See 9.26 “To manage a post-discovery action failure on an NE” (p. 315) for more information.

9.5.2 Post-discovery action design considerations

The following conditions apply to post-discovery actions.

- The NFM-P performs no logic validation of the associated scripts.
- Deletion of the associated control script is blocked.
- After a discovery rule runs, the action is not automatically executed on new NEs that are added to the rule. To execute the action on the new NEs, you must rescan the network using the rule.
- An action is performed on an NE only when the Full Resync Status of the NE is Full Resync Done.
- If a device fails discovery, the action is not performed, and the script execution status is Not Executed. Resynchronization of the NE does not invoke the script or change the execution status.
- If the control script execution fails, an alarm is raised.
- After a post-discovery action is performed on an NE:
  - Unmanaging and remanaging the NE has no effect.
  - Deleting and rediscovering the NE causes the action to be performed on the NE.

9.6 Configuring SSH security on devices

9.6.1 About the SSH protocol

The SSH protocol provides secure file transfer and file system access between the NFM-P and managed NEs. SSH version 2, or SSH2, is enabled by default on many devices. SSH2 uses paired public and private encryption keys to perform authentication. After an SSH key pair is generated on an NE, the private key is stored locally, and the public key is used by SSH2 clients. A public key persists in the NFM-P for future SSH2 communication with the NE, and is used to verify that the client is connecting to the correct SSH2 entity. An SSH2 server on an NE identifies itself to a client by sending a message that is signed using the private key of the server. The NFM-P uses the public key of the SSH2 server to authenticate the server identity.

9.6.2 SSH2 host key management

An NE sends an SSH2 host key when the NFM-P first tries to establish an SSH2 connection. The NFM-P automatically accepts the public key fingerprint and stores it locally. The NFM-P uses the local fingerprint copy for authentication during subsequent sessions.

If an NE sends a different host key in a subsequent session, the NFM-P rejects the connection attempt and raises an alarm. If an operator determines that the host key change is valid, for example, because of an NE reboot while host key persistence on the NE is disabled, the operator can manually delete the mismatched host key from the NFM-P and accept the new key. The NFM-P subsequently accepts SSH2 connection requests from the NE.

Note: Nokia recommends that you enable host key persistence on an SSH2 NE to retain the public key fingerprint after NE reboots. If public key persistence is disabled, connection attempts after an NE reboot fail until an NFM-P operator manually deletes the stored key. See 9.34 “To manually accept a mismatched SSH host key” (p. 323) for information about deleting a mismatched host key.

By default, SSH host keys persistence is disabled on the 7210 SAS, 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS.

9.6.3 SSH2 and device CLI sessions

When SSH2 for CLI sessions is enabled in the mediation policy of an SSH2-capable device, SSH2 instead of Telnet is used for each CLI session.

SSH1 is used only in SSH CLI sessions on NEs that do not support SSH2. SSH1 is not supported for communications with GNEs.

9.6.4 SSH2 and secure file transfers

When secure file transfers are specified in the mediation policy of an SSH2-capable device, SCP is used instead of FTP for backups, restores, software upgrades, and statistics collection.

9.7 Mediation and event notification policies

9.7.1 Mediation policies

To discover and manage devices in your network, you must create a mediation security policy to setup the security and communication infrastructure between the NFM-P and each device within the specified discovery rule IP range.

You can also optionally configure other mediation policies to specify other network mediation tasks such as specifying the polling interval for MIB configuration changes, how the NFM-P processes specific events, and how the NFM-P pings devices in the network to ensure they are reachable.

Caution: When configuring a mediation policy, you need use a community string that uses rwa in order for the trap destination is set on the NE. This ensures fault conditions and alarms are reported.

The following table describes the device mediation policies you can configure using the NFM-P.

Table 9-1 NFM-P device mediation policies

Policy type	Purpose	See
Mediation security	Defines the network security model and communication protocols used to discover and manage NEs for all NFM-P user types except LI users. The policy specifies the SNMP, CLI, and file transfer settings, and the credentials for security functions. A discovery rule requires a mediation security policy for: • read access • write access • SNMP trap access	9.17 “To configure device mediation” (p. 301)
MIB entry	Defines the interval at which the NFM-P polls NEs MIBs for configuration changes.	
Management ping	Defines the management IP address of the device that the NFM-P pings periodically to verify reachability. You can specify which of the following interfaces the NFM-P pings: • Out of Band Management Interface Ping • In Band Management Interface Pin • Standby CPM Ping	
Event notification	Defines how the NFM-P processes specific SNMP events from specific NE types and releases. See 9.7.2 “Event notification polices” (p. 285) in this section for more information.	
LI mediation security	Defines the network security model used for LI mirror service creation; LI users do not use a standard mediation security policy. The policy specifies the NE LI user credentials, the authentication and privacy protocol types, and the required passwords. The NFM-P GUI displays an LI mediation security policy only for an LI user. See Chapter 94, “Lawful Intercept” for information about configuring LI.	94.14 “To configure an LI security mediation policy” (p. 3207)

9.7.2 Event notification polices

CAUTION

Service Disruption

The NFM-P uses SNMP notifications to verify network management functions. Before you configure an event notification policy, consult Nokia support to ensure that you do not disable the processing of traps that the NFM-P requires.

To reduce the NFM-P processing load associated with SNMP events, you can configure a policy that specifies which SNMP traps the NFM-P processes. An event notification policy acts as a filter that enables or disables the processing of specific SNMP traps for a specific NE type and release.

An NFM-P operator with an admin or operations scope of command role and write access permission to the policy package can create an event notification policy during mediation configuration. The NFM-P assigns a default event notification policy to an NE during initial device discovery. You can optionally specify which traps to process and which to ignore; processing is enabled for all supported traps in the default policy.

The following conditions apply to event notification policies:

- After an NE upgrade, the NFM-P processes the new traps generated by the upgraded NE. An NFM-P administrator must ensure that the event notification policies that are in effect before an NE upgrade are correctly configured, and must modify them after the upgrade for new traps, as required.
- An NFM-P software upgrades does not affect an existing policy.

SNMP statistics that include the number of ignored traps for an NE are available from the Statistics tab of an NE properties form.

9.8 NE resynchronization

9.8.1 Full and custom NE resynchronization

The Resync contextual menu option for an object specifies that SNMP MIB and CLI information bases for the object are reread to resynchronize them with the NFM-P database. The NFM-P database is updated to reflect the NE information. When you choose the Resync option at the NE level, you can choose to resynchronize all MIBs or perform a customized resynchronization. When you choose to resynchronize all MIBs, the color of the NE object in the navigation tree is yellow during the resynchronization. When you choose a customized resynchronization, you can choose the MIB entries to synchronize. See 9.33 “To partially or fully resynchronize NEs with the NFM-P database” (p. 322) for more information about performing NE resynchronizations.

If SNMP streaming is enabled on an NE, the resynchronization time may be significantly reduced in a high-latency network. See 9.2.3 “SNMP streaming” (p. 279) in “Discovering devices using the NFM-P” (p. 277) for more information.

9.8.2 Initial synchronization

An NFM-P operator must ensure that the discovery of each managed NE completes successfully, and that the NE and NFM-P database information are fully synchronized. The First Full Resync Done attribute on the Managed State tab of the Discovery Manager form indicates whether an NE has undergone one full resynchronization since the NE discovery, most recent NE software upgrade, or an NFM-P system upgrade. If an NE has not undergone one full resynchronization, and the Resync Status of the NE on the Resync Status tab is Resync Failed, the operator must manually resynchronize the NE.

9.8.3 Scheduled NE resynchronization

Some MIBs and device types are set with a scheduled resynchronization value when the device is discovered. Resynchronization is scheduled to capture any device changes that are not event notified or available in SNMP traps. Scheduled polling can be added or changed for MIBs as needed. Scheduled resynchronization is enabled in the Polling tab in the device properties and configured in the Mediation menu. See 8.10 “To configure polling for a 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, VSR, or Wavence SM” (p. 260) and 9.17 “To configure device mediation” (p. 301) for more information.

9.8.4 Resynchronization status

The Resync Status tab of the Discovery Manager form displays the following resynchronization attributes for each managed NE; the attributes are updated after an attempted resynchronization of all MIBs, but not after a customized resynchronization attempt:

- Resync Status
- Last Resync Start Time
- Last Resync End Time
- Scheduled Resync Status
- Last Scheduled Resync Start Time
- Last Scheduled Resync End Time

The Resync Status value is one of the following:

- Full Resync Done—a resynchronization has successfully completed
- Full Resync Failed—a resynchronization attempt has failed

  Full Resync Failed does not indicate that the SNMP agent on the NE is unreachable; if the agent is unreachable, an alarm is raised against the NE and the NE icon color changes to red. When the agent is again reachable, the alarm clears and the NE icon color changes to green.
- In Progress—a resynchronization is in progress
- Not Attempted—no resynchronization has been requested
- Requested—the resynchronization request is queued for processing.

The Scheduled Resync Status value is one of the following:

- In Progress—a resynchronization is in progress
- Not Attempted—no resynchronization has been requested
- Scheduled Resync Done—a resynchronization has successfully completed
- Scheduled Resync Failed—a resynchronization attempt has failed

9.9 Server resource management

9.9.1 Resource grouping

CAUTION

Service Disruption

Only a qualified Nokia support representative can assess or configure NFM-P server resource management. Contact Nokia technical support for more information.

The NFM-P supports the grouping of NEs by network function for increased SNMP mediation efficiency.

Resource allocation is automatically configured during NFM-P installation and network discovery based on the available system resources, and the numbers and types of NEs that are to be managed. The resource group to which an NE belongs is initially determined by the device type. The following are the NFM-P NE resource groups and the default device assignments:

- global—for NFM-P tasks that are not associated with NEs
- default—for NEs that do not belong to a specific resource group
- core—for 7450 ESS, 7750 SR, and 7950 XRS devices
- middle—for 7210 SAS and 7705 SAR devices
- edge—for OmniSwitch and similar devices
- GNE—for other network devices

The NFM-P logs the system load of each resource group at regular intervals. An Nokia support representative can monitor system resource usage, determine whether the current resource allocation is appropriate, and reconfigure resource management, if required.

9.10 Workflow for device discovery

9.10.1 Stages

Preconfigure device for management

1

Configure the Nokia and GNE devices for management; see Chapter 8, “Device commissioning and management” .

Configure SNMP

2

As required, enable SNMPv3 device management; see 9.11 “To enable SNMPv3 management of a device” (p. 291) .

3

As required, enable or disable SNMP streaming on a device for configuration data transfers to the NFM-P; see 9.14 “To enable or disable SNMP streaming on an NE” (p. 299) .

Configure SSH

4

As required, configure SSH for secure device CLI access.

1. Configure SSH servers on devices; see the device documentation for configuration information.
2. Ensure that SSH2 is enabled on devices; see 9.15 “To verify that SSH2 is enabled on a device” (p. 300) .
3. Enable SSH2 host key persistence on devices that support host key persistence; see 9.16 “To enable SSH host key persistence on a device” (p. 301) .

Configure device mediation

5

Specify the network mediation parameters, and optionally configure the following mediation policies for devices in specific discovery rules; see 9.17 “To configure device mediation” (p. 301) :

- MIB entry policies that define how to poll specific device MIBs for changes
- management ping policies that define periodic device connectivity checks
- event notification policies that define how the NFM-P processes SNMP traps

6

As required, assign event notification policies to devices; see 9.18 “To assign an event notification policy to an NE” (p. 305) .

Configure additional management networks

7

If required, configure one or more additional management networks.

1. Define the management network and subnets to associate with non-default management interfaces on main and auxiliary servers; see 9.19 “To configure a management network” (p. 306) .
2. Create additional main-server management interfaces; see 9.20 “To configure an additional management interface on a main server” (p. 307) .
3. Create additional auxiliary-server management interfaces; see 9.21 “To configure an additional management interface on an auxiliary server” (p. 308) .

Configure device discovery

8

If required, create a workflow for post-discovery actions using Workflows.

9

If required, create post-discovery actions; see 9.22 “To configure a post-discovery action” (p. 309) .

Note: You can also create post-discovery actions during discovery rule creation.

10

Create the required device discovery rules and use the rules to discover devices; see 9.23 “To configure a discovery rule” (p. 310) .

11

Manage discovery rules; see 9.24 “To enable, disable, or delete a discovery rule” (p. 313) .

12

As required, manage the post-discovery actions in discovery rules.

a. View the execution status of an associated script; see 9.25 “To view the post-discovery action execution status” (p. 314) .

b. Manually execute a failed script; see 9.26 “To manage a post-discovery action failure on an NE” (p. 315) .

Incidental tasks

13

As required, perform the following.

a. Change the management state of a device; see 9.27 “To manage, suspend, or unmanage a device” (p. 316) .

b. Change from SNMPv2c to SNMPv3 management of a device; see 9.29 “To change from SNMPv2 to SNMPv3 management of a device” (p. 318) .

c. Switch from non-secure to secure mediation; see 9.30 “To switch from non-secure to secure mediation” (p. 319) .

d. Specify which management IP address is saved for a device when the device is set to an unmanaged state; see 9.31 “To specify which management address the NFM-P uses to remanage a device” (p. 321) .

e. Rescan the network using a discovery rule; see 9.32 “To rescan the network for a device according to a discovery rule” (p. 322) .

f. Synchronize the NFM-P information with the device configuration; see 9.33 “To partially or fully resynchronize NEs with the NFM-P database” (p. 322) .

g. Manage SSH2 host key acceptance; see 9.34 “To manually accept a mismatched SSH host key” (p. 323) and 9.35 “To view the SSH2 host keys to identify active and mismatched keys” (p. 324) .

h. View and save collected SNMP MIB information; see 9.36 “To list and save SNMP MIB information” (p. 324) .

i. Delete a device from the NFM-P; see 9.37 “To delete a device from the managed network” (p. 325) .

Procedures for device discovery

9.11 To enable SNMPv3 management of a device

9.11.1 Before you begin

If you are configuring an NE for LI, you must create a second access group. See Chapter 94, "Lawful Intercept" for information about creating an LI user and access group.

SROS 22.2 R1 and later NEs do not support some combinations of authentication and privacy:

- hmac-md5-96 cannot be combined with cfb128-aes-192 or cfb128-aes-256
- hmac-sha1-96 cannot be combined with cfb128-aes-192 or cfb128-aes-256
- hmac-sha2-224 cannot be combined with cfb128-aes-256

9.11.2 Steps

1

Open a CLI session on the device.

2

Enter the following commands in the order shown to create a read-write-notify group for general SNMP mediation on the managed device:

```
configure system security snmp ↲
access group "SNMPv3_group" security-model usm security-level privacy
read "iso" write "iso" notify "iso" ↲
```

where

SNMPv3_group is the name to assign to the new SNMP group

3

If mediation of VPRN objects is required (for configured VPRN or internal VPRN from an Ethernet satellite), enter the following command to create a read-write-notify group for this purpose on the managed device:

```
access group "SNMPv3_group" security-model usm security-level privacy
context vprn prefix read "vprn-view" write "vprn-view" notify "iso" ↲
```

where

SNMPv3_group is the name to assign to the new SNMP group

4

Enter the following command to exit the SNMP group configuration.

```
exit ↲
```

5

Enter the following command to obtain the SNMP engine ID of the device.

show system info ↵

The SNMP engine ID is displayed as SNMP Engine ID.

6

Generate an authentication key and a privacy key.

- An authentication key is used to encrypt a user password.
- A privacy key is used to encrypt the user SNMP packets.

Note: The key authentication method determines the key length.

1. Log in to an NFM-P single-user client, client delegate server, or main server station.

   Note:

   If you log in to a RHEL main or client delegate server station, you must log in as the nsp user.

   If you log in to a single-user client station, you must log in as the user who installed the client, or as a local administrator.

2. Open a console window.
3. On a RHEL station, navigate to the *install_directory*/nms/bin directory, where *install_directory* is one of the following:
   - the NFM-P main server installation location, /opt/nsp/nfmp/server
   - the NFM-P single-user client or client delegate server installation location, typically /opt/nsp/client
4. On a Windows station, navigate to the *install_directory*\nms\bin directory, where *install_directory* is the NFM-P single-user client or client delegate server installation location, typically C:\nsp\client.
5. Enter one of the following to create an authentication key:
   - on a RHEL station:

     ./nmsclient.bash password2key *method password engine_ID* ↵
   - on a Windows station:

     nmsclient.bat password2key *method password engine_ID* ↵

   where

   method is the authentication method, either MD5, SHA, SHA224, SHA256, SHA384 or SHA512

   password is the authentication key password

   engine_ID is the SNMP engine ID obtained in Step 5

   Note: You must enclose a password that contains a special character in single quotation marks; for example:

   password2key *method* 'Mypa$$word'

   Only use the authentication key from the output.
6. Enter the following to create a privacy key.
   - on a RHEL station:

     ./nmsclient.bash password2key *method password engine_ID* ↵

- on a Windows station:

  nmsclient.bat password2key *method password engine_ID* ↵

  where

  method is the authentication method, either MD5, SHA, SHA224, SHA256, SHA384 or SHA512

  password is the privacy key password

  engine_ID is the SNMP Engine ID of the SR, in hexadecimal form with 10-64 hex digits (5-32 bytes)

  Note: You must enclose a password that contains a special character in single quotation marks; for example:

  password2key *method* 'Mypa$$word'

  The list of privacy keys for each privacy method is displayed.

7. Store the generated keys for your applicable authentication and privacy methods.

7

Using the keys generated in Step 6 , create an SNMPv3 user on the managed device.

1. Enter the following sequence of commands at the prompt:

   configure system security user *SNMPv3_user* ↵

   access snmp ↵

   snmp ↵

   authentication *auth_method authentication_key* privacy *priv_method privacy_key* ↵

   group *SNMPv3_group* ↵

   exit all ↵

   where

   SNMPv3_user is the name to assign to the new user

   SNMPv3_group is the name of the new SNMP user group created in Step 2

   auth_method can be:

   hmac-md5-96 hmac-sha1-96 hmac-sha2-224 hmac-sha2-256 hmac-sha2-384 hmac-sha2-512

   authentication_key is the authentication key value generated in Step 6

   priv_method can be:

   cbc-des cfb128-aes-128 cfb128-aes-192 cfb128-aes-256

   privacy_key is the privacy key value generated in Step 6

2. Enter the following to save the configuration changes:

   admin save ↵

   The device is now ready for management using SNMPv3.

3. Close the CLI session.

8

Create an SNMPv3 NE user in the NFM-P.

See the section on NE user and device security in the *NSP System Administrator Guide* for specific information about creating and configuring NE users.

1. Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens.
2. Click Create. The NE User (Create) form opens.
3. Configure the parameters as follows:
   - Enter the *SNMPv3_username* value from Step 7 as the User Name.
   - Enable the SNMP option of the Access parameter.
   - On the SNMPv3 tab:
     — Configure the authentication and privacy parameters. Ensure the NEs support a valid combination of authentication and privacy protocols.
     — Enter the password used to generate the authentication key in Step 6 .
     — Enter the password used to generate the privacy key in Step 6 .

   Note: In the mediation security policy that you create for the device, you must specify the following:
   - SNMPv3 (USM) as the Security Model value
   - the *SNMPv3_user* value from Step 7 as the User Name

   See 9.17 “To configure device mediation” (p. 301) for specific information about creating and configuring mediation security policies.

9

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.12 To enable SNMPv3 management and discover an 1830 VWM device

9.12.1 Purpose

Perform this procedure to enable SNMPv3 management of an 1830 VWM device, configure the mediation policies, and discover the 1830 VWM device.

Note: You can either use the 1830 VWM default SNMPv3 user or configure a new SNMPv3 user in the 1830 VWM device.

9.12.2 Steps

SNMPv3 user configuration in the 1830 VWM device

1

Perform the following to configure a new SNMPv3 user in the 1830 VWM device, if required.

1. Open a CLI session on the device.
2. Enter the following commands in the order shown to create a read-write-notify group for general SNMP mediation on the managed device:

   # config admin snmpusers add "*SNMPv3_user*" admin auth sha "*password*" priv des "*password*"

   # config admin snmpuser edit "*SNMPv3_user*" status enabled

   where

   SNMPv3_user is the name to assign to the new SNMPv3 user

   password is the password of the new SNMPv3 user.
3. Close the CLI session.

SNMPv3 user configuration in the NFM-P

2

Note: See the section on NE user and device security in the *NSP System Administrator Guide* for specific information about creating and configuring NE users.

Perform the following to configure an SNMPv3 user in the NFM-P.

1. Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens.
2. Click Create. The NE User (Create) form opens.
3. Enter the *SNMPv3_username* value from Step 1 or the default user name as the User Name.
4. Check the SNMP option of the Access parameter.
5. Click on the SNMPv3 tab and configure the parameters as follows:
   - Choose SHA as the Authentication Protocol.
   - Choose AES 128 as the Privacy Protocol.
   - Enter the password provided in Step 1 or the default password.
6. Save your changes and close the form.

Mediation policy configuration

3

Note: The 1830 VWM discovery rule requires the SNMPv3 read, write, and security access mediation policies and SNMPv2 trap access mediation policy.

Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.

4

Configure the parameters on the General tab and verify the SNMP trap destination parameters in the SNMP Trap panel.

5

Click on the Mediation Security tab and click Create. The Mediation Policy (Create) form opens.

6

Configure the Policy ID and Displayed Name parameters.

7

Perform the following to configure the SNMPv3 mediation policy.

1. Configure the Security Model parameter to SNMPv3 (USM).
2. Select the required SNMPv3 user in the SNMPv3 panel.

8

Perform the following to configure the SNMPv2 mediation policy.

1. Configure the Security Model parameter to SNMP v2c.
2. Configure the Community String parameter in the SNMP v1/v2c panel.

   The community string configured in the NFM-P is automatically added in the node.

9

Configure the parameters in the SNMP panel. The Table 9-2, “ SNMP timeout” (p. 296) lists the recommended values for Timeout (milliseconds) and Retry parameters.

Table 9-2 SNMP timeout

Timeout (milliseconds)	Retry
60000	1
50000	2

10

Save your changes and close the form.

Discovery of the 1830 VWM device

11

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

12

Click Create to configure a new discovery rule. The Create Discovery Rule step form opens.

13

Configure the required parameters in the Specify General Attributes step and click Next. The Add Rule Elements step form opens.

14

Click Create, configure the required IP parameters, and click OK. The rule element is added to the list.

15

Click Next. The Add Auto Discovery Rule Elements ACL step form opens.

16

Click Create, configure the required IP parameters, and click OK.

17

Click Next and the Configure Mediation Security step form opens.

18

Click Select in each of the following panels to specify the mediation security policies. Choose the SNMPv3 mediation security policy for read, write, and security access mediation policies and choose the SNMPv2 mediation policy for trap access mediation policy.

- Read Access Mediation Policy
- Write Access Mediation Policy
- Trap Access Mediation Policy
- Security Access Mediation Policy

Note: If you do not specify a policy, the default policy is applied.

19

Click Finish to close the Create Discovery Rule form.

20

Apply your changes in the Discovery Manager form. If the Administrative State of the newly created discovery rule is set to Up, the NFM-P scans the network using the discovery rule.

END OF STEPS

9.13 To configure the AIM mediation and discovery for management of the VSR-I

9.13.1 General information

The NFM-P uses SNMPv2 or SNMPv3 to manage the AIM for VSR-I. After the AIM is discovered, the VSR-I IP is automatically added to the AIM discovery rule, so that it is automatically discovered.

Note: AIM does not support auto-registration of the NFM-P IP address for trap forwarding during discovery. You must manually enter the NFM-P IP address during AIM agent configuration at start-up so that traps from the AIM agent are sent to the NFM-P.

Note: You must configure the SNMPv2 or SNMPv3 parameters the same way for AIM and the VSR-I.

9.13.2 Steps

1

Configure an SNMPv2c or SNMPv3 mediation security policy on the NFM-P.

1. Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.
2. Click on the Mediation Security tab.
3. Click Create. The Mediation Policy (Create) form opens.
4. Configure the Displayed Name parameter.
5. Set the Security Model parameter to SNMP v2c or SNMPv3.
6. In the SNMP panel, set the Port parameter to 161.
7. To create an SNMPv2 user on the NFM-P, go to the next step. To create an SNMPv3 user on the NFM-P, go to step 9.
8. **Create an SNMPv2 user on the NFM-P:**
   a. Configure the Community String parameter to XYZ, where XYZ is the SNMP community string configured on the AIM node.
   b. Go to step 10
9. **Create an SNMPv3 user on the NFM-P:**
   a. Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens.
   b. Click Create. The NE User, Global Policy (Create) form opens.
   c. Configure the user name and other required parameters, and choose the SNMP option for the Access parameter.

d. Click on the SNMPv3 tab and provide the authentication and privacy parameters.
e. In the CLI/NETCONF panel, configure the required parameters.

10. Save your changes.

2

Create a discovery rule for the VSR-I on the NFM-P:

a. Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.
b. Click Create. The Create Discovery Rule step form opens.
c. In step 1 (Specify General Attributes), configure the required parameters, then click Next.
d. In step 2 (Add Rule Elements), click Create. The Topology Discovery Rule Element (Create) form opens.
e. Enter the AIM management IP address as the rule element for the discovery rule.
f. Click OK, then click Next.
g. In step 3 of the step form (Configure Mediation Security), click Select and choose the mediation security policy that you created in step 1 for the Read Policy ID, Write Policy ID, and Trap Policy ID parameters.
h. Save your changes.

3

Use the NFM-P client to discover the VSR-I. The NFM-P discovers both AIM and VSR-I in a new topology group under the equipment tree named VSR-a..

END OF STEPS

9.14 To enable or disable SNMP streaming on an NE

9.14.1 Purpose

Perform this procedure to control whether an NE uses SNMP streaming for configuration data transfers to the NFM-P.

9.14.2 Steps

1

Right-click on an NE in the network navigation tree and choose Properties. The NE properties form opens.

2

Click on the Polling tab, then on the Management tab.

3

Configure the Enable SNMP Streaming parameter.

4

Save your changes and close the form.

5

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.15 To verify that SSH2 is enabled on a device

9.15.1 Purpose

Perform this procedure to verify that SSH2 is enabled on a device, and to enable SSH2 on the device, if required.

9.15.2 Steps

1

Open a CLI session on the device.

2

Run the following CLI command to see whether SSH2 is enabled:

```
show system security ssh ↵
```

3

If required, enter the following command at the prompt to enable SSH2:

```
configure system security ssh version 2 ↵
```

Note: A 7705 SAR may become temporarily unreachable after you enable SSH and start the SSH server on the device.

4

Close the CLI session.

5

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.16 To enable SSH host key persistence on a device

9.16.1 Purpose

Perform this procedure to enable the persistence of the server host key on a 7210 SAS, 7450 ESS, 7750 SR, 7705 SAR, or 7950 XRS. For other devices, see the user documentation for information about configuring SSH2 host key persistence.

9.16.2 Steps

1

Open a CLI session on the device.

2

Enter the following at the prompt to disable the SSH server:

```
configure system security ssh server-shutdown ↵
```

3

Enter the following at the prompt to enable host key persistence:

```
configure system security ssh preserve-key ↵
```

4

Enter the following at the prompt to enable the SSH server:

```
configure system security ssh no server-shutdown ↵
```

5

Enter the following at the prompt to verify that the preserve-key function is enabled on the server:

```
show system security ssh ↵
```

6

Close the CLI session.

7

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.17 To configure device mediation

9.17.1 Purpose

Perform this procedure to create a mediation policy.

As part of creation of a mediation policy, you will configure the following communication and security policies for discovery rules:

- MIB entry policies that define how to poll specific device MIBs for changes
- management ping policies that define periodic device connectivity checks
- event notification policies that define how the NFM-P processes SNMP traps

Note: LI mediation has special configuration requirements; see 94.14 “To configure an LI security mediation policy” (p. 3207) for information.

Note: To configure SNMPv3 security, you must preconfigure the NFM-P and managed devices. See 9.11 “To enable SNMPv3 management of a device” (p. 291) for information.

Note: For Wavence SNMPv3 password information, see the Wavence documentation.

9.17.2 Steps

1

Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.

2

Configure the parameters on the General tab.

3

Configure a mediation security policy.

Note: Some NEs require specific mediation security policies. See the NFM-P device-specific documentation for information.

1. Click on the Mediation Security tab.
2. Click Create or choose the default policy and click Properties. The Mediation Policy form (Create | Edit) opens.
3. Configure the general policy and SNMP mediation parameters.

   Note:

   The Community String value must match the community string value on the managed device.
4. If you set the Security Model parameter to SNMPv3 (USM), click Select to choose the required SNMPv3 user. SNMPv3 user creation is described in 9.11 “To enable SNMPv3 management of a device” (p. 291) .

   Note:

   If you specify the use of SNMPv3 in the policy, you must enable SSH2 in the policy.
5. Configure the remaining parameters.

   Note:

If you choose Secure as the File Transfer Type value, you must set the Communication Protocol parameter to SSH2.

To use SSH2, you must enable SSH2 and host key persistence on the SSH server of the device. See 9.6 “Configuring SSH security on devices” (p. 283) for information about configuring SSH2.

If you choose SSH2 as the Communication Protocol, you must specify the mediation policy for read and write access in the discovery rule of the device.

If you choose SSH2 as the Communication Protocol, the User Name and User Password parameters must be the SSH user name and password configured on the device.

See Chapter 8, “Device commissioning and management” for information about enabling FTP access for a device user account.

6. Save your changes and close the form.

4

To configure one or more MIB Entry policies, perform the following steps.

1. Click on the MIB Entry Policies tab and click Search. A list of MIBs appears, organized by the product name of the device that supports the MIB.
2. Choose one or more MIBs from the list and click Properties. The MIB Entry Policy (Edit) form opens.
3. Configure the required parameters in the Configuration panel.

   Caution: Changing the Number of Varbind per PDU value may affect the time required for subsequent NE resynchronizations and degrade NFM-P system performance. Do not configure the parameter without consulting Nokia technical support.
4. Save your changes and close the form.

5

To configure a management ping policy, perform the following steps. A management ping policy specifies how the NFM-P checks the connection to device management IP addresses. Each managed device may provide one or more of the following addresses:

- the system IP address, which is an in-band management interface
- the management IP address, which is an out-of-band management interface
- the IP address of the standby CPM

Note: When the device does not have one or two of the IP addresses, for example, there is no CPM IP address, ensure that you do not enable the schedule on the ping policy assigned to missing interfaces. This allows the assignment of an inactive ping policy during discovery configuration.
The destination interface of the ping is determined during discovery-rule creation. You can also perform an unscheduled ping from the Managed State tab of the Discovery Manager configuration form.
You can view management connection alarms from the NE properties form of the managed device that was pinged. For example, from the Discovery Manager form, click on the Resync Status tab. Choose a device from the list and click Properties. From the NE properties form, click on the Faults tab to view the alarms.

1. Click on the Ping tab.
2. Click Create or choose a ping policy and click Properties. The Management Ping Policy (Create | Edit) form opens.
3. Configure the required parameters.

   Note:

   You must enable scheduling for a ping policy to be active. When scheduling is not enabled, and an assigned managed device is not reachable, management connection alarms may not be raised.
4. Save your changes and close the form.

6

CAUTION

Service Disruption

The NFM-P uses SNMP notifications to verify network management functions. Before you configure an event notification policy, consult Nokia support to ensure that you do not disable the processing of traps that the NFM-P requires.

To configure event notification policies, perform one of the following. Otherwise, go to Step 8 .

a. Select a default trap policy from the list and click Properties. The Event Notification Policy (Edit) form opens.

b. Create an event notification policy:
   1. Click on the Event Notification Policies tab.
   2. Click Create or select a policy and click Properties. The Event Notification Policy (Create | Edit) form opens.
   3. Configure the parameters.
   4. Click OK. The policy is listed on the form.
   5. Apply the changes in the Mediation (Edit) form.
   6. Select the new policy in the list and click Properties. The Event Notification Policy (Edit) form opens.

7

Enable or disable collection of traps in a policy.

Note: You cannot configure the Administrative State parameter in a default event notification policy. A default policy is indicated by a check mark beside the Default Policy indicator on the General tab of the Event Notification Policy (Edit) form.

1. Click on the Notification Policies tab. A list of SNMP traps is displayed.
2. Select an SNMP trap entry in the list and click Properties. The MIB Entry Policy - trappolicy (Edit) form opens.
3. Configure the Administrative State parameter.

4. Save your changes and close the forms.

8

Click OK to save your changes and close the Mediation (Edit) form.

9

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.18 To assign an event notification policy to an NE

CAUTION

Service Disruption

The NFM-P uses SNMP notifications to verify network management functions. Before you assign an event notification policy to a device, consult Nokia support to ensure that you do not disable the processing of traps that the NFM-P requires.

9.18.1 Steps

1

On the equipment tree, right-click on an NE icon and choose Properties. The Network Element (Edit) form opens.

2

Click on the Polling tab and click Clear in the Assigned Event Notification Policy panel. The Select button in the panel is enabled.

3

Select an event notification policy.

4

Save your changes and close the form.

5

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.19 To configure a management network

9.19.1 Purpose

Perform this procedure to create a management network and subnets to segregate the traffic between the non-default management interfaces on the main and auxiliary servers.

9.19.2 Steps

1

Choose Administration→System Information from the NFM-P main menu. The System Information form opens.

2

Click on the Management Networks tab.

3

Click Create. The Management Network (Create) form opens.

4

Configure the required parameters.

5

Click on the Management Interfaces tab.

6

Click Add. The Select Interface for Management Network (Create) form opens.

7

Click Search. The management interfaces are listed.

Note: The default management interfaces are not listed.

8

Select an interface and click OK.

9

In a redundant NFM-P system, you must assign one interface from each main server to a managed network.

To add another interface, repeat Step 6 to Step 8 .

10

Click on the Management Subnets tab.

11

Click Create. The Management Subnet (Create) form opens.

12

Configure the required parameters.

Note: The Subnet Address value that you specify must include a mask that is between /24 and /32 for IPv4 NEs, and between /120 and /128 for IPv6 NEs.
The subnets must not overlap.

13

To create an additional subnet, click Apply to save your changes and repeat Step 12 .
Otherwise, click OK to save your changes and close the form.

14

Click Apply to save your changes. The NFM-P attempts to discover NEs in the subnets using the associated interfaces. The discovered NEs are listed on the Member Network Elements tab.

15

Save your changes and close the forms.

16

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.20 To configure an additional management interface on a main server

Note: In a redundant NFM-P system, you must configure management interfaces for the required networks and subnets on each main server. Otherwise, visibility of the networks and subnets is lost after a server activity switch.

9.20.1 Steps

1

Choose Administration→System Information from the NFM-P main menu. The System Information form opens.

2

Click Properties in the Primary Server panel. The Main Server (Edit) form opens.

3

Click on the Management Interfaces tab.

4

Click Create. The Management Interface (Create) form opens.

5

Configure the required parameters.

6

Click OK. The interface appears in the list.

7

Save your changes and close the forms, as required.

8

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.21 To configure an additional management interface on an auxiliary server

9.21.1 Steps

1

Choose Administration→System Information. The System Information form opens.

2

Click on the Auxiliary Servers tab.

3

Select an auxiliary server and click Properties. The Auxiliary Server (Edit) form opens.

4

Click on the Management Interfaces tab.

5

Click Create. The Management Interface (Create) form opens.

6

Configure the required parameters.

7

Click OK. The interface appears in the list.

8

Save your changes and close the forms, as required.

9

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

End of steps

9.22 To configure a post-discovery action

9.22.1 Purpose

Perform this procedure to create a post-discovery action that you can associate with one or more discovery rules.

Note: You can also create a post-discovery action during discovery rule creation.

9.22.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click on the Post Discovery Actions tab.

3

Click Create or select a post-discovery action and click Properties. The Post Discovery Action (Create | Edit) form opens.

4

Configure the required parameters.

5

Click Select beside the Control Script Name parameter and choose a control script.

6

Click Select beside the Control Script Instance Name parameter and choose a control script instance.

7

Save your changes and close the form.

8

If you are performing this procedure for device discovery, see 9.10 “Workflow for device discovery” (p. 288) for the next steps.

END OF STEPS

9.23 To configure a discovery rule

9.23.1 Purpose

Perform this procedure to configure a discovery rule that contains the mediation, IP information, and policies required to add NEs or EMS to the managed network.

9.23.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click Create to create a new discovery rule. The Create Discovery Rule step form opens.

Note: When you are finished each step, click Next to advance to the next step. You can click Finish at any time.
Depending on the purpose of the discovery rule, not every step in this procedure needs to be completed.

3

Configure the required parameters in the Specify General Attributes step.

Note: When you close the Discovery Manager form, the NFM-P scans the network using the discovery rule, if the Administrative State parameter of the discovery rule is set to Up. The Group Name parameter specifies the equipment group that discovered NEs are added to. If the selected equipment group reaches the maximum element limit, any additional discovered NEs are automatically added to the Discovered NEs group.
The Discovered NEs Group Name cannot be chosen when configuring a Discovery Rule for an EM system.

4

To add rule elements in the Add Rule Elements step, click Create, configure the required IP parameters, and click OK. The rule element is added to the list.

5

To add auto discovery rule elements in the Add Auto Discovery Rule Elements ACL step, click Create, configure the required parameters, and click OK. The rule element is added to the list.

6

In the Configure Mediation Security step, click Select in each of the following panels, as required, to specify the mediation security policies for:

- Read Access Mediation Policy
- Write Access Mediation Policy
- Trap Access Mediation Policy
- Security Access Mediation Policy

Note: If you do not specify a policy, the default policy is applied.
A Wavence MSS-1c discovery rule requires the default SNMPv2 read and write access policies.

7

In the Configure Management Ping Policy step, click Select in each of the following panels, as required, to specify the ping policies for:

- Out Of Band Management Interface Ping
- In Band Management Interface Ping
- Standby CPM Ping

Note: If you do not specify a policy, the default policy is applied.
Management ping policies are created using the Mediation configuration form. These are the policies applied during discovery rule creation. You must apply a ping policy even to interfaces that do not exist. When there is no interface, you must choose a ping policy that has its Schedule Enabled parameter set to disabled.
For example, if there is no standby CPM or out-of-band management IP address, specify a ping policy that has the Schedule Enabled parameter set to disabled for the nonexistent management and standby CPM interfaces. In this example, a ping policy with the Schedule Enabled parameter enabled is required for the in-band interface that does exist.
See 9.17 “To configure device mediation” (p. 301) for information about creating ping policies.

8

In the Configure MIB Statistics Policy step, select a MIB statistics policy, if required.

9

In the Add Discovered Routers to Span(s) step, click Add to specify a span of control for the new NE contained in the discovery rule. If you do not specify a span, the default span is applied.

Note: New devices added to a span from a discovery rule are added explicitly to the span. When a discovered NE group is part of a user-defined span, new devices that are discovered using the discovery rule are automatically added to the span.

10

In the Configure Backup Policy step, select a backup policy, if required. If you do not specify a backup policy, the default policy is applied.

11

If the NFM-P manages wireless NEs in your network, perform the Add NE Self Config Policies step.

12

To configure an EMS in the Add EM Systems step:

1. Click Create. The EM System (Create) form opens.
2. Configure the required parameters in the General tab.
3. Click on the Element Managers tab.
4. Click Create. The Element Manager (Create) form opens.

   Note:

   After discovering an EM system, the system should be opened and its administrative state set to Up in order to enable the connection to the NFM-P. A maximum of four EM Systems may have their administrative state set to Up at any time within the NFM-P.

   A manual resynchronization must be performed in order for alarms from an EM system to be displayed in the NFM-P. This does not apply to the NEtO EM system for Wavence devices; alarms are aggregated automatically provided the Trap Access Mediation policy in the Discovery step form is configured.

   The NFM-P is able to manage 10 000 alarms per EM system. When this limit is reached, a critical alarm is raised and new incoming alarms are deleted. When the alarm counter for the EM System drops to 9000, the critical alarm is cleared and the NFM-P performs an alarm resynchronization on the EM system.
5. Configure the required parameters for host, communication, and identification.
6. Save your changes and return to the Create Discovery Rule form.

13

To specify a post-discovery action in the Add Post Discovery Action step, perform one of the following:

a. Click Select and use the list form that opens to associate an existing post-discovery action with the discovery rule.

b. Create a post-discovery action to associate with the discovery rule. Perform the following steps.
   1. Click Select. The Select Post Discovery Action form opens.
   2. Click Create. The Post Discovery Action (Create) form opens.
   3. Configure the required parameters:
   4. Click Select beside the Control Script Name parameter and use the list form that opens to choose a control script.
   5. Click Select beside the Control Script Instance Name parameter and use the list form that opens to choose a control script instance.
   6. Save your changes and close the form.
   7. Select the new post-discovery action and click OK. The new post-discovery action is associated with the discovery rule.

14

Click Finish to close the Create Discovery Rule form.

15

Apply your changes in the Discovery Manager form. If the Administrative State of the newly created discovery rule is set to Up, the NFM-P scans the network using the discovery rule.

16

See 9.10 “Workflow for device discovery” (p. 288) for other optional steps.

END OF STEPS

9.24 To enable, disable, or delete a discovery rule

9.24.1 Before you begin

When a discovery rule is enabled, the network is scanned according to the discovery rule when the discovery rule is saved or rescanned. The network is also scanned according to the discovery rule as specified by the Discovery Rule Scan Interval parameter in the Mediation form. If your discovery rule is disabled, the network is not scanned as specified by these conditions.

When you delete a discovery rule, only the rule is removed. The NEs discovered using the rule are not removed from the NFM-P, but the Discovery Rule ID for each of the NEs set to 0.

9.24.2 Steps

1

Choose Administration→Discovery Manager from the main menu. The Discovery Manager (Edit) form opens.

2

Select a discovery rule and click Turn Up to enable the rule, Shut Down to disable the rule, or Delete to delete the rule.

3

Close the form.

END OF STEPS

9.25 To view the post-discovery action execution status

9.25.1 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click on the Managed State tab.

3

View the entries in the Script Execution Status column. The possible values are:

- unspecified—post-discovery action not specified in discovery rule
- Pending—post-discovery action execution awaits discovery completion
- Not Executed—control script not executed
- In Progress—control script execution in progress
- Successful—control script execution completed without failure indication
- Failed—control script execution unsuccessful
- Cleared—Failed state manually cleared

4

For each NE that has a Failed entry, perform the following steps to view the script execution result.

1. Click Show Results. The Results form for the NE opens.
2. Click Search. A list of script results is displayed.
3. Select a result in the list and click Properties. The View Result form opens. The form displays information such as the script name, version, run time, and execution status.
4. View the information to determine the required corrective action.
5. If required, perform the following steps to export the result information to a file.
   - Choose File→Export from the View Result form main menu, or click on the Export icon. The Export form opens.

- Navigate to the location in which you want to save the file, and enter a filename in the File Name field.
- Click Export. The result text is saved in the specified file.

Note:

You can also export a script result to a file using the Export button on the Results form for the NE.

5

To delete a result, perform the following steps.

1. Select the result on the Results form and click Delete. A dialog box appears.
2. Click Yes. The result is deleted.

END OF STEPS

9.26 To manage a post-discovery action failure on an NE

9.26.1 Purpose

Perform this procedure to do the following:

- reset the script execution status for an NE
- manually run the post-discovery action script against the NE

This action may be required after the script associated with a post-discovery action fails and the cause of the failure is corrected.

Note: Clearing the status also clears the associated alarm that is raised when post-discovery script execution fails.

9.26.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click on the Managed State tab.

3

Select an NE with an Execution Status of Failed and click Properties. The Node Discovery Control (Edit) form opens.

4

Click on the Post Discovery tab.

5

Click Clear. The Script Execution Status changes from Failed to unspecified.

6

To execute the script:

1. Click Execute Script. The Execute script form opens.
2. Click Execute. The script executes, and a script object is displayed in the navigation tree of the form.
3. Select the script object in the navigation tree.
4. Expand the Status and Script Results panels on the right side of the form to reveal the script execution information.
5. Click Save Result to save the result in the NFM-P database.

7

Close the forms, as required.

END OF STEPS

9.27 To manage, suspend, or unmanage a device

9.27.1 Purpose

CAUTION

Service Disruption

Suspending the management of a device results in a loss of communication with the device.

Unmanaging a device results in a loss of management data for the device. See 9.3.2 "Unmanaging and deleting devices" (p. 280) in "Discovering devices using the NFM-P" (p. 277) for more information.

Unmanaging a device results in the loss of all historical AA statistics for the device.

When you suspend the management of a device, the NFM-P raises a critical alarm against the device.

Perform this procedure to change the management state of a device.

9.27.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.

2

Click on the Managed State tab.

3

Select a device in the list.

4

Perform one of the following:

a. Click Manage to return an unmanaged or suspended device to the managed state.

b. Click Suspend to suspend the management of the device.

c. Click Unmanage to stop the management of the device.

5

Save your changes. The Managed State of the device changes to one of the following:

- managed, if you are remanaging the device
- suspended, if you are suspending management of the device
- not managed, if you are unmanaging the device

6

If Service Bandwidth Management is enabled, perform a system-wide CAC audit to ensure that the bandwidth is properly calculated:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Click on CAC Audit.

END OF STEPS

9.28 To associate a device with a discovery rule

9.28.1 Purpose

After a discovery rule is deleted, any managed devices associated with the discovery rule are not associated with a discovery rule. Perform this procedure to associate one or more managed devices with a discovery rule to ensure that the devices are included in subsequent network administration activities.

9.28.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens with the Discovery rules tab displayed.

2

Select the discovery rule to associate the NE with and click Properties, or use the Create button to create a discovery rule. The Topology Discovery Rule (Edit) form opens.

Note: Ensure that the mediation policy and other policies associated with the discovery rule are appropriate for the devices that you intend to add.

3

Click on the Rule Elements tab.

4

Select an exiting rule element and click Properties, or use the Create button to create a new rule element.

Note: The rule element must be configured to include the device IP address, for example, explicitly or using a subnet specification.

5

Click Assign Node and use the form that opens to select the devices that you want to add to the discovery rule.

6

Click OK. The devices are associated with the discovery rule.

7

Click OK to save the change, if required, and close the form.

8

Close the Discovery Manager (Edit) form.

END OF STEPS

9.29 To change from SNMPv2 to SNMPv3 management of a device

9.29.1 Purpose

Perform this procedure to change the management protocol that the NFM-P uses for a device from SNMPv2 to SNMPv3.

9.29.2 Steps

1

Perform 9.11 “To enable SNMPv3 management of a device” (p. 291) to do the following:

- Configure an SNMPv3 user and group on the device.
- Configure an SNMPv3 user in NFM-P.

2

Perform 9.17 “To configure device mediation” (p. 301) to configure an SNMPv3 mediation policy.

3

Modify the mediation policy of the managed device that you need to change from SNMPv2 to SNMPv3:

1. Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens with the Discovery rules tab displayed.
2. Click on the Managed State tab, choose the managed device that you need to change from SNMPv2 to SNMPv3, and click Properties. The Node Discovery Control (Edit) form opens.
3. Click Select and choose the SNMPv3 mediation policy configured in Step 2, in the read, write, trap, and Security Access Mediation Policy panels.
4. Save your changes and close the forms.

4

Perform 9.23 “To configure a discovery rule” (p. 310) to configure a new discovery rule for the SNMPv3 device.

5

Edit the SNMPv2 discovery rule for the device to exclude the device from discovery.

6

Use a CLI session on the device to remove the following items, which are no longer required:

- SNMPv2 user
- SNMPv2 group
- SNMPv2 trap targets

7

Resynchronize the device with the NFM-P; see 9.33 “To partially or fully resynchronize NEs with the NFM-P database” (p. 322).

END OF STEPS

9.30 To switch from non-secure to secure mediation

9.30.1 Steps

1

Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.

2

Click on the Mediation Security tab. You can create a mediation policy as described in 9.17 “To configure device mediation” (p. 301) , or modify the default mediation policy to use SSH2. Perform one of the following.

a. Click Create to create a new mediation policy. The Mediation Policy (Create) form opens.

b. Choose the default policy from the list, and click Edit. The Mediation Policy (Edit) form opens.

3

Configure the required parameters using the prescribed values:

- SSH2 as the Communication Protocol
- The SSH2 server user name and password as the User Name and User Password
- Secure as the File Transfer Type

  Note:

  When the File Transfer Type is set to Secure, the Communication Protocol, User Name, and User Password parameters must be configured with the SSH2 information.
- The port that the NE uses for SSH2 communication as the SSH2 Server Port

4

Configure the remaining parameters on the form.

5

Save your changes and close the form.

6

Create a new discovery rule that uses the new SSH2 mediation policy as the read-access mediation policy and the write-access mediation policy. See 9.23 “To configure a discovery rule” (p. 310) for information about creating discovery rules.

7

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.

8

Click on the Managed State tab. A list of managed devices appears.

9

Choose the device from the list that you want to use SSH2 for CLI sessions and secure file transfers.

10

Click Properties and click on the Mediation Security tab.

11

Choose the newly created SSH2 mediation policy as the read access mediation policy.

Note: You can also change the mediation policy for a discovery rule by performing Step 11 and choosing a discovery rule from the Discovery Rules tab.

12

Save your changes and close the form.

END OF STEPS

9.31 To specify which management address the NFM-P uses to remanage a device

9.31.1 Before you begin

You can specify which management IP address is saved for a device when the device is set to an unmanaged state. You can configure the NFM-P to save the original management IP address or the most recently used management IP address for a device.

9.31.2 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.

2

Click on the Managed State tab.

3

Choose a device and click Properties. The Node Discovery Control (Edit) form opens.

4

Configure the Use Original Management IP parameter and click OK.

5

Save your changes and close the forms.

END OF STEPS

9.32 To rescan the network for a device according to a discovery rule

9.32.1 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.

2

Select one or more discovery rules and click Rescan.

The NFM-P scans the network as specified by the discovery rules and discovers devices. After a device is discovered, the NFM-P server sets the device in a managed state and adds the NE properties to the NFM-P database.

END OF STEPS

9.33 To partially or fully resynchronize NEs with the NFM-P database

9.33.1 Steps

1

On the equipment tree, select one or more NEs, as required.

Note: If more than one NE is selected, only full resynchronization is available.

2

Choose a resynchronization option.

a. Right-click on one of the selected devices you want to resynchronize with the NFM-P database and choose Resync→Resync All MIBs.

Note: Choosing Resync All MIBs will fully resynchronize all MIBs for the devices, ignoring timestamps.

b. Perform a partial resynchronization.

1. Right-click on one of the selected devices you want to resynchronize with the NFM-P database and choose Resync→Customized Resync. The Resync Site(s) form opens.
2. Select Choose MIB Entries and click Next.
3. Select one or more MIB entries from the list and click Next.
4. Configure the Ignore Timestamps parameter.
5. Click Finish and close the form.

END OF STEPS

9.34 To manually accept a mismatched SSH host key

9.34.1 Purpose

Perform this procedure to manually accept a rejected SSH host key in order to establish a connection to the SSH server on an NE.

Note: Before you accept a mismatched host key, you must verify the validity of the SSH connection.

9.34.2 Steps

1

Open the SSH2 Known Host Key Manager by performing one of the following.

a. Choose Administration→Security→SSH2 Known Host Key Manager from the NFM-P main menu. The SSH2 Known Host Key Manager filter form opens.

b. Perform the following steps.

1. Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form appears.
2. Click on the Managed State tab and click Search. A list of managed devices appears.
3. Select a device in the list and click Properties. The Node Discovery Control (Edit) form opens.
4. Click on the Mediation Security tab.
5. Click SSH2 Known Host Key. The SSH2 Known Host Key Manager filter form opens.

2

In the SSH2 Server Key Status column header, choose Mismatch SSH2 Host Key as a match criterion using the drop-down menu.

3

Click Search. A list of mismatched host keys is displayed.

4

Select the required host key entry.

5

Verify with the device management that the key fingerprint is the host key of the required device.

6

Click Delete to delete the mismatched host key. The mismatched host key is deleted and a connection to the SSH server can be established.

7

Close the forms.

END OF STEPS

9.35 To view the SSH2 host keys to identify active and mismatched keys

9.35.1 Steps

1

Open the SSH2 known host key manager by performing one of the following.

a. To view the active and mismatched SSH host keys on all managed NEs, choose Administration→Security→SSH2 Known Host Key from the NFM-P main menu. The SSH2 Known Host Key Manager form opens.

b. To view the active and mismatched SSH host keys on specific managed NEs:

1. Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form appears.
2. Click on the Managed State tab and click Search. A list of managed devices appears.
3. Select a device and click Properties. The Node Discovery Control (Edit) form opens.
4. Click on the Mediation Security tab.
5. Click SSH2 Known Host Key. The SSH2 Known Host Key Manager form opens.

2

Configure the filter criteria and click Search. A list of active and mismatched host keys appears.

3

Close the forms.

END OF STEPS

9.36 To list and save SNMP MIB information

9.36.1 Purpose

You can list and save SNMP MIB information, which may be of use for purposes such as the following:

- to maintain a record of the SNMP MIBs that the NFM-P supports
- to compare SNMP MIB support between NFM-P releases
- to identify the polling interval for a MIB entry

9.36.2 Steps

1

Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.

2

Click on the MIB Entry Policies tab.

3

View the polling interval for a MIB entry.

1. Select a MIB entry and click Properties. The MIB Entry Policy (Edit) form opens.
2. View the polling interval setting in the Configuration panel.
3. Close the form.

4

Sort the list according to the contents of a column by clicking on the column header.

5

Save the listed information to a file.

1. Right-click on a column header and choose Save to File. The Save As form opens.
2. Use the form to specify the file that is to contain the saved information.
3. Click Save. The information is saved to the specified file.

6

Close the form.

END OF STEPS

9.37 To delete a device from the managed network

CAUTION

Service disruption

Deleting a device results in a loss of management data and completely removes the device from the managed network.

Deleting a device results in the loss of all historical AA statistics for the device.

See "Discovering devices using the NFM-P" (p. 277) for more information.

9.37.1 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens.

2

Click on the Managed State tab and select a device.

3

If the device is managed, click Unmanage and wait for the state to change to Unmanaged. Otherwise, go to Step 4 .

4

Click Delete.

5

Save your changes and close the form. The device is deleted from the NFM-P database.

END OF STEPS

10 Device CLI sessions

Managing device CLI sessions using the NFM-P

10.1 Overview

10.1.1 Functions that require CLI

You can perform most NE management functions using the NFM-P client GUI. Functions such as the following, however, require CLI access to a managed NE:

- validating GUI configuration actions
- configuring items that the GUI cannot, such as LI user access
- troubleshooting using device debug files

The NFM-P client GUI provides CLI access to the managed NEs from the main menu and from NE contextual menus in topology maps and navigation trees.

10.1.2 Security

Scope of command roles and user permissions on the cli package control CLI command access on an NE. The user permissions are set when security is configured on the NE locally or by using NFM-P configuration forms. See the section on NE user and device security in the *NSP System Administrator Guide* for information about setting NE access privileges.

Note: When you use a CLI to change security parameters on an NE, the changes may not be synchronized with the NFM-P and subsequently not displayed in the GUI. For security reasons, the NFM-P cannot retrieve parameters such as passwords from an NE. To ensure that the security parameters are synchronized, you must use the client GUI to change the values.

Note: One CLI login failure from the NFM-P client GUI may generate more than one login failure message in the NE session logs. This is because the CLI window is opened using an internal utility rather than a command shell, and the utility may attempt the login more than once for various reasons, for example, excessive retries or the use of multiple authentication sources.

10.1.3 vi editor support

On a 7450 ESS, 7750 SR, or 7950 XRS, you can use the vi editor to modify local files. The vi editor is available in NFM-P Telnet and SSH sessions. The vi editor supports the standard vi navigation keys as well as the cursor keys (which are often called the arrow keys). See the appropriate device documentation for information about the supported vi commands and functions for a specific device.

10.2 Workflow to use an NFM-P CLI

10.2.1 Purpose

The following workflow describes the sequence of high-level tasks required to open and use a device CLI from the NFM-P client GUI. The following must be true before you can use the NFM-P to open a device CLI using the GUI.

- The NFM-P user account has console access privileges on the managed devices. See the section on NE user and device security in the *NSP System Administrator Guide* for more information about user account privileges.
- The Telnet server on the managed devices is started.
- If required, SSH2 is properly configured on the managed devices; see Chapter 9, “Device discovery” for more information.

10.2.2 Stages

1

Configure the NFM-P CLI preferences. See 10.3 “To configure the NFM-P CLI console preferences” (p. 328) for more information.

2

Open a CLI session and log in to a managed device. See 10.4 “To open and close an NFM-P device CLI session” (p. 329) for more information.

3

Configure the device, view device information, and modify files, as required.

4

Close the CLI session. See 10.3 “To configure the NFM-P CLI console preferences” (p. 328) for more information.

10.3 To configure the NFM-P CLI console preferences

10.3.1 Purpose

Perform this procedure to customize the NFM-P CLI window settings, such as the following:

- the console text style and appearance
- the size of the scrolling buffer
- whether to save the session output to a file

10.3.2 Steps

1

Open a CLI session, as described in 10.4 “To open and close an NFM-P device CLI session” (p. 328) .

2

Right-click in the CLI window and choose Configure. The Terminal Configuration form opens.

3

Configure the required parameters.

Note: When the Send Console To a File parameter is enabled and the Append to file parameter is disabled, the log file is overwritten each time a new CLI session starts.

4

Click OK. The Terminal Configuration form closes.

END OF STEPS

10.4 To open and close an NFM-P device CLI session

10.4.1 Purpose

Perform this procedure to open a Telnet or SSH session on a managed device using the NFM-P client GUI.

10.4.2 Steps

1

Perform one of the following to open a device CLI session.

a. Use the NFM-P navigation tree.
   1. Choose Equipment from the view selector.
   2. Right-click on a device icon and choose NE Sessions→Telnet Session or NE Sessions→SSH Session from the contextual menu.

      The Telnet Session or SSH Session window opens.

      Note: The managed device must be configured for Telnet access. See the appropriate device documentation for information about configuring Telnet access to the device.

      Note: SSH sessions use SSH2 by default on the 7210 SAS, 7450 ESS, 7750 SR, 7950 XRS, and GNEs. Ensure that SSH2 is properly configured on the device and that an SSH2 mediation policy is specified in the device discovery rule. See Chapter 9, “Device discovery” for more information about configuring SSH2 security.

b. Use the NFM-P main menu.

1. Choose one of the following main menu options:
   - Tools→Network Elements→NE Sessions→Telnet Session
   - Tools→Network Elements→NE Sessions→SSH Session

   The Telnet Session or SSH Session window opens.
2. Perform one of the following to specify an NE as the CLI session target.
   - Enter the management IP address of the NE beside the dimmed Connect button.
   - Select an IP address from the drop-down menu.
   - Click Select to search for the NE.

   The Connect button is enabled.
3. Click Connect.

c. Use the Manage Equipment form.
1. Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.
2. Select an NE in the list.
3. Click NE Sessions and choose Telnet Session or SSH Session. The Telnet Session or SSH Session window opens.

d. Use a service topology map.
1. Choose one of the following from the NFM-P main menu:
   - Manage→Service→Services
   - Manage→Service→Composite Services
   - Manage→Service→Mirror Services

   The appropriate management list form opens.
2. Choose a port and click Topology View. The Service Topology map opens.
3. Right-click on an NE in the map and choose NE Sessions→Telnet Session or NE Sessions→SSH Session. The Telnet Session or SSH Session window opens.

e. Use the Physical Topology map.
1. Choose Application→Physical Topology from the NFM-P main menu. The Physical Topology map opens.
2. Right-click on an NE in the map and choose NE Sessions→Telnet Session or NE Sessions→SSH Session. The Telnet Session or SSH Session window opens.

2

Enter the login credentials. You can use the CLI as specified by your user account permissions.

Note: One login failure may generate more than one login failure message in the NE session logs. This is because the CLI window is opened using an internal utility rather than a command shell, and the utility may attempt the login more than once for various reasons, for example, excessive retries or the use of multiple authentication sources.

3

Perform NE management tasks, as required.

4

Click Disconnect to end the CLI session.

Note: When you disconnect from a CLI session, the session window remains open to facilitate the opening of a subsequent CLI session.

5

Perform one of the following.

a. Open a new CLI session:
   1. Perform one of the following to specify an NE as the CLI session target.
      - Enter the management IP address of an NE beside the dimmed Connect button.
      - Select an IP address from the drop-down menu.
      - Click Select to search for an NE.

      The Connect button is enabled.
   2. Click Connect.

b. Close the Telnet Session or SSH Session window.

END OF STEPS

11 Working with network objects

Working with network objects using the NFM-P

11.1 Overview

11.1.1 NFM-P management of network objects

The NFM-P is used to create, configure, and manage a device with the various hierarchical objects required to be part of a network. Equipment such as the routers, which are at the top of the hierarchy, have properties that are configured using the CLI and discovered when the NFM-P discovery process is run. After the device is discovered, it is displayed as an object in the navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about the navigation tree.

Objects in the NFM-P are considered to have parent/child relationships that are contained within a hierarchy. For example, a card in a card slot is the parent object of a daughter card. The behavior of each object is defined using parameters that are specific to the function required. Those parameters can be managed to suit the needs of the service required. Objects are created and managed using the properties forms found in the contextual menus of the equipment view.

The network is the top object in the navigation tree. The network icon in the navigation tree is the parent object of all managed devices. When you expand the network icon, all managed devices are shown as children of the network parent.

The device object is the discovered device at the top of the hierarchy in the navigation tree, directly below the network icon. The following children objects of the router are created automatically in the navigation tree after the device is discovered.

- CCAG
- ISA-Tunnel Group
- ISA-AA Group
- ISA-LNS Group
- ISA-NAT Group
- ISA-Video Group
- LAG
- IGH
- Shelf
- Power Shelf
- Card Slot
- Card Slot A for the CPM and switch fabric
- Card Slot B for the redundant CPM and switch fabric
- Fans
- PCM Trays
- Power Supplies
- CCM
- SFM Slot

The following objects must be created using property forms or create forms from the contextual menus of the equipment view or the ring group view.

- equipment groups
- individual CCAGs with VSM-CCA members
- individual ISA-Tunnel Groups with ISA IPsec MDA members
- individual ISA-AA Groups with ISA-AA MDA or ESA-AA VM members and ISA-AA partitions
- individual ISA-LNS Groups with ISA broadband applications MDA members
- individual ISA-NAT Groups with ISA broadband applications MDA members

- individual ISA-Video Groups with ISA-Video MDA or ESA Video Group members
- individual LAGs with subgroups of LAG member ports
- IGH members
- cards
- ring groups
- daughter cards

  Ports are automatically created when the daughter card is created.
- channels

Configuring an object is accomplished in two steps. First the object must exist or be created, second, the object parameters are modified. See 12.4 “To create an object” (p. 343) to create objects using the navigation tree.

11.1.2 Object deployment status

When you configure an object using the NFM-P GUI or the NFM-P OSS, the NFM-P saves the configuration and attempts to deploy the changes to the network. Although the value is changed on the configuration forms, the configuration change may not be fully deployed to the NE.

The NFM-P allows you to monitor the deployment state for objects. The Deployment tab on all object properties forms notifies you of configuration changes that are not fully deployed to the NE. The parameter that was unsuccessfully deployed is listed, along with the old value and the new value. A deployment icon or warning indicator appears in the following locations when you change the configuration of an object:

- beside all of the parameters on an object properties form for which a deployment is in progress (yellow icon) or has failed (red icon)
- beside the object in the navigation tree (blue icon, regardless of state)
- in info tables in map views (blue icon, regardless of state)
- in the Deployment column on object list forms (blue icon, regardless of state)
- on the Deployment tab (warning indicator)

See Chapter 25, “NE deployment” for information about how to configure a deployment policy, manage deployments, and troubleshoot failed deployments. See 11.5 “To monitor the deployment status of a network object” (p. 336) for information about how to monitor the deployment of an object configuration change.

11.2 Working with equipment group objects

11.2.1 Overview

The equipment group icons in the navigation tree represent logical equipment groups. Initially, an equipment group object is created by choosing Create→Equipment→Group from the NFM-P main menu. A Group (Create) form opens. Additional equipment groups can be created using the Copy button on the equipment group properties form.

The Discovered NEs and the Unmanaged NEs equipment groups are created by default. The Discovered NEs equipment group contains the discovered devices. The Unmanaged NEs

equipment group contains non NFM-P managed NEs that have been added to the network. You can create new equipment groups and move devices to them using the topology maps. The topology map supports equipment groups of up to 250 NEs.

Note: For existing topologies, an NFM-P managed NE is not moved automatically to the Unmanaged NEs equipment group if it is unmanaged by the user.

See Chapter 4, “Topology map management” for information about configuring equipment group objects and managing equipment groups using maps.

11.2.2 Performance considerations

Although there is no limit to the number of equipment groups, Nokia recommends a maximum of 10 000 equipment groups per system for optimal performance. Each equipment group can contain a maximum of 2000 objects. Network elements and immediate child groups are considered objects in an equipment group. A bracketed number is displayed beside the name of each equipment group in the navigation tree. This number indicates the current number of objects in the group.

11.3 Working with physical links

11.3.1 Overview

The NFM-P allows you to create and manage links at the Layer 1 level. The physical links represent the actual physical configuration of network connections between ports. You can view and manage physical links from the physical topology map and the Manage Equipment list form.

When a physical link becomes operationally down, an alarm is raised against the link, and correlated to alarms on the ports on either end of the link.

Although there is no limit to the number of physical links you can have in a system, Nokia recommends a maximum of 10 000 physical links for optimal system performance.

Radio links between Wavence SMs are shown as physical links by the NFM-P, with ports as the endpoint type.

See Chapter 4, “Topology map management” for information about configuring physical links and managing physical links using maps.

11.4 Workflow to manage network objects

11.4.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure network objects. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: Network objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information.

11.4.2 Stages

1

Create network equipment groups as required; see Chapter 4, “Topology map management” .

2

Configure device objects as required; see Chapter 12, “Device object configuration” .

3

Configure logical group objects as required; see Chapter 13, “Logical group object configuration” .

4

Configure shelf, card, and daughter card objects as required; see Chapter 15, “Shelf and card object configuration” .

5

Configure port and channel objects as required; see Chapter 16, “Port and channel object configuration” .

6

Create physical links as required; see Chapter 4, “Topology map management” .

7

Monitor the deployment status of network objects as required; see 11.5 “To monitor the deployment status of a network object” (p. 336) .

8

View the current network resources assigned to network objects such as IP or MAC addresses as required; see 11.6 “To view network resources assigned to network objects” (p. 337) .

11.5 To monitor the deployment status of a network object

11.5.1 Steps

1

Open the properties form for the network object.

You can click on the Deployment tab if a warning indicator appears when you open a properties form.

2

Configure the required parameters and click Apply. An icon appears beside the modified parameter:

- yellow icon—deployment in progress; the icon disappears when the deployment succeeds
- red icon—failed deployment

In addition, a blue deployment icon appears beside the NE object in the navigation tree, in list forms, and in configured info tables in map views.

A warning indicator appears on the Deployment tab.

3

Click on the Deployment tab.

4

In the Name column, verify the name of the parameter for which the deployment was unsuccessful. You can also verify the old value of the parameter before the change was applied, the new parameter value, and the deployer ID.

5

Navigate to the parameter on the configuration form that failed to deploy, if necessary, by selecting an entry in the Attributes list and clicking Show Attribute. The attribute is highlighted for a few seconds on the configuration form.

Note: If you choose an entry in the Attributes list that is located on a hidden tab and click Show Attribute, the hidden tab is temporarily displayed. See 1.24 “To set local tab preferences for configuration forms” (p. 118) on how to configure hidden tabs.

6

Troubleshoot the cause of the failed deployment; see 25.5 “To view and manage failed deployments” (p. 769) .

End of steps

11.6 To view network resources assigned to network objects

11.6.1 Purpose

Perform this procedure as required to view network resources assigned to network objects such as IP or MAC addresses.

11.6.2 Steps

1

Choose Tools→Network Resources→ Network Resources from the NFM-P main menu. The Network resources form opens.

2

Select the appropriate tab for the network resource you need to query.

3

CAUTION

Service Disruption

This operation may take a long time to complete for some network resources such as viewing IP addresses, route distinguishers and route targets, and MAC addresses. Nokia recommends that you use a specific search filter to narrow down the scope of search.

Customize the search fields as required to view the chosen network resource assigned to the network objects.

4

Click Search. For some search queries, for example IP addresses, a dialog box appears.

5

If required, acknowledge the dialog box and click Search. The results of your customized search are displayed.

6

Close the form.

END OF STEPS

12 Device object configuration

Working with device objects

12.1 Overview

12.1.1 Device objects

The device object is the discovered device at the top of the hierarchy in the navigation tree, directly below the network icon.

This chapter contains the procedures to configure device objects using the navigation tree. Device object properties forms, which are used to configure specific parameters for discovered devices, are accessed using objects on the NFM-P navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the navigation tree.

The device icons in the navigation tree represent device objects. See Chapter 6, “Device support” for more information about the supported device types.

Most of the configured properties for this object are inherited from the device. The Properties contextual menu option from the navigation tree allows you to configure or modify the parameters for the object. See the specific procedures in this chapter for the respective tab of the properties form that is to used to configure the device object.

12.1.2 Ring group objects

The NFM-P allows you to create and manage ring groups in the ring group view of the navigation tree. Ring groups are used to group devices logically to represent a typical network topology.

Note: The NFM-P uses VLAN groups instead of ring groups to group OmniSwitch NEs. See Chapter 76, “VLL service management” for more information about VLAN groups.

12.2 Workflow to manage device objects

12.2.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure device objects. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: Device objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

12.2.2 Stages

1

As required, configure the following general NE capabilities:

a. Create an object; see 12.4 “To create an object” (p. 343) .

b. Modify an NE configuration; see 12.5 “To modify NE properties” (p. 343) .

c. Customize the NE property labels and values in the System preferences form, for example, to specify the location and site name that differs from the actual NE site name. See the *NSP System Administrator Guide* for more information about configuring system preferences.

d. Customize the NE property values in the NE form; see 12.6 “To configure NE custom properties” (p. 343) .

e. Create an operational group; see 12.8 “To create an operational group” (p. 345) and 12.9 “To configure a 7210 SAS operational group” (p. 346) .

f. Enable and configure global Cflowd on an NE, and create Cflowd collectors; see 12.10 “To enable and configure global Cflowd sampling on an NE” (p. 347) .

g. Associate a span of control with an NE; see 12.13 “To add a span of control to an NE” (p. 350) .

h. Configure load balancing on an NE; see 12.14 “To configure load balancing” (p. 351) .

i. Configure 802.1x authentication on an NE; see 12.17 “To enable or disable 802.1X” (p. 353) .

j. Enable or disable an exclusive policy editing restriction on one or more NEs or remove any exclusive edit locks that were previously set; see 12.18 “To configure an exclusive policy editing restriction on an NE” (p. 354) .

k. Configure active card alarms on an NE; see 12.19 “To configure active card alarms on an NE” (p. 355) .

l. Configure a TWAMP server on an NE; see 12.20 “To configure a TWAMP server” (p. 355) .

m. Enable LLDP on an NE; see 12.21 “To enable LLDP on an NE” (p. 358) .

n. Configure BFD flap detection on an NE; see 12.22 “To configure the BFD flap detection interval on an NE” (p. 360) .

o. Enable a Q in Q untagged SAP on an NE; see 12.23 “To enable a Q in Q untagged SAP on an NE” (p. 360) .

p. Configure the DHCPv6 advertise message format on an NE; see 12.24 “To configure DHCPv6 Advertise messages on an NE” (p. 361) .

q. Configure Python script protection on an NE; see 12.25 “To configure Python script protection on an NE” (p. 362) .

r. Configure SFLOW on an NE; see 12.29 “To configure sFlow on an NE” (p. 364) .

s. Create a chassis-level PBB configuration; see 12.31 “To create a chassis-level PBB configuration” (p. 367) .

t. Configure serving network information; see 12.32 “To configure serving network information on an NE” (p. 367) .

u. Configure the RADIUS CoA port on an NE; see 12.36 “To configure the RADIUS CoA port on an NE” (p. 369) .

v. Configure data persistence on an NE; see 12.37 “To configure data persistence on an NE” (p. 370) .

w. Configure DNS security extensions on an NE; see 12.38 “To configure DNS security extensions” (p. 371) .

x. Configure the use of vendor-specific ICMP extensions on an NE; see 12.39 “To enable or disable ICMP extensions on a NE” (p. 372) .

2

As required, configure the following functionality on 7210 SAS NEs:

a. Configure operational groups for fault propagation on a 7210 SAS; see 12.9 “To configure a 7210 SAS operational group” (p. 346) .

b. Configure a no-service loopback port on a 7210 SAS; see 12.46 “To configure a no-service loopback port on the 7210 SAS” (p. 378) .

c. Configure two WRED slopes on a 7210 SAS; see 12.48 “To configure two WRED slopes on a 7210 SAS” (p. 379) .

d. Enable frame-based accounting on a 7210 SAS; see 12.49 “To configure frame-based accounting for QoS policies on a 7210 SAS” (p. 380) .

e. Configure the system resource profile on a 7210 SAS or 7250 IXR; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

f. Configure a system resource profile policy for 7210 SAS-R or 7210 SAS-S/Sx VC NEs; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382) .

g. Configure port-based access egress scheduling on 7210 SAS-Mxp or 7210 SAS-R NEs; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384) .

h. Configure buffer management for the 7210 SAS; see 12.54 “To configure buffer management for the 7210 SAS” (p. 385) .

3

As required, configure the system resource profile on a 7250 IXR; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380).

4

As required, configure the following functionality on 7705 SAR NEs:

a. Configure the watermark settings on a 7705 SAR; see 12.59 “To configure watermark settings on a 7705 SAR” (p. 391) .

b. Configure the QoS ingress aggregate rates on a 7705 SAR-M; see 12.60 “To configure QoS

ingress aggregate rates on the 7705 SAR-M/ME" (p. 1818) .

c. Launch an MCT from a Wavence SA NE connected to a 7705 SAR; see 12.61 "To launch an MCT on a Wavence SA NE connected to a 7705 SAR" (p. 392) .

5

As required, configure the following functionality on OmniSwitch NEs:

a. Start and stop a Webview or Secure Webview session on an OmniSwitch; see 12.62 "To start and stop a Webview or Secure Webview session on an OmniSwitch" (p. 393) .

b. Configure the dying gasp alarm on an OmniSwitch; see 12.63 "To configure the dying gasp alarm on an OmniSwitch" (p. 394) .

6

As required, configure ring group objects; see 12.73 "To create a ring group" (p. 401) and 12.74 "To remove a device from a ring group or a ring group" (p. 402) .

12.3 Workflow to configure UNP 802.1x at port level for OmniSwitch devices

12.3.1 Purpose

The following is the sequence of high-level actions required to configure UNP 802.1x at port level.

12.3.2 Stages

1

Manage Omni nodes and add it in VLAN Group; see 75.13 "To create a VLAN group" (p. 2083).

2

Configure an AOS UNP Profile; see 61.6 "To configure an OmniSwitch Ethernet UNP profile" (p. 1818).

3

Configure UNP profile at port level; see 12.70 "To configure an UNP at port-level" (p. 398).

Note: UNP profile can also be configured globally; see 12.69 "To configure global-level UNP" (p. 398)

General device configuration procedures

12.4 To create an object

12.4.1 Steps

1

On the equipment tree, right-click on an empty object and choose Properties, or, when available, Create <*objectname*>. The properties form or the create form, as applicable, opens.

2

Configure the required parameters.

Certain object parameters are available for configuration. Configuring these parameters creates the object, however, after the object is created you may need to edit it using properties forms.

3

Save your changes and close the form.

END OF STEPS

12.5 To modify NE properties

12.5.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Use the form to view or modify the NE parameters, as required.

3

Save your changes and close the form.

END OF STEPS

12.6 To configure NE custom properties

12.6.1 Steps

1

On the equipment or routing tree, right-click on an NE. The Network Element (Edit) form opens.

2

In the Custom Properties panel, configure the user-defined values.

Note: The NE custom properties support the extended character set including multi-byte characters.
The configured labels and values are displayed in the following locations. You can configure the labels and values, if required.

- NE list form: Choose Manage→Equipment→Equipment→Network Element (from the object drop-down menu). The appropriate columns display the values and labels.
- Tree labels: See 2.12 “To customize tree labels” (p. 152) .
  You can also configure Custom NE property values using the XML API. See the *NSP NFM-P XML API Developer Guide*.
  The NFM-P GUI supports localized language display. See Chapter 1, “NFM-P GUI” for more information.

3

Save your changes and close the form.

END OF STEPS

12.7 To enable FIPS-140-2

12.7.1 Purpose

Perform this procedure to enable FIPS-140-2 level 1 support.

The NFM-P returns an SNMP deployment error from the NE for operations that are not supported because of FIPS restrictions.

Consider the following when you enable FIPS-140-2:

- The following algorithms are not accessible for keychains:
  - MD5
  - HMAC-MD5
  - DES
- The following algorithms are not available:
  - MD5
  - DES

Note: FIPS-140-2 level 1 is only supported on the 7705 SAR, Release 8.0 R6 or later. If you enable FIPS-140-2 from the node CLI and perform an NE reboot or software upgrade from Release 8.0 R4 to Release 8.0 R5 in the NFM-P GUI, the 7705 SAR NE will not boot.

12.7.2 Steps

1

On the equipment or routing tree, right-click on an NE. The Network Element (Edit) form opens.

2

Click on the Polling tab. In the General sub-tab, enable the Enable FIPS-140-2 parameter.

3

Save your changes and close the form.

4

Reboot the NE.

END OF STEPS

12.8 To create an operational group

12.8.1 Before you begin

Multiple objects can be bundled in operational groups. If you change the status of the operational group, the status of every object in the group is changed.

12.8.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the Service tab.

3

Click on the Operational Groups tab and then click Create. The Operational Group (Create) form opens.

4

Configure the required parameters.

5

Bind a BFD interface to the operational group.

1. In the BFD Interface Monitored panel, click Create. The Operational Group BFD Interface (Create) form opens.
2. Configure the Interface Type parameter.
3. For service interfaces, click Select in the Service panel. For network interfaces, click Select beside the Interface Name parameter.
4. Choose a service or interface from the displayed list and click OK.

5. Configure the Destination Address parameter.

6

Save your changes and close the forms.

END OF STEPS

12.9 To configure a 7210 SAS operational group

12.9.1 Purpose

You can create operational groups for VLANs on supported 7210 SAS NEs. Operational groups provide fault propagation for ports with null, dot1q, and QinQ encapsulation types.

You can assign an L2 uplink port to an operational group. Access ports on the same device can monitor the group. When the L2 uplink port goes down or loses connectivity, the access ports monitoring the group go operationally down. SAPS associated with the affected ports also switch to operationally down.

Perform this procedure to create or modify an operational group. To associate ports with the group, see 16.24 “To configure Ethernet ports” (p. 599) .

12.9.2 Steps

1

In the equipment view of the navigation tree, right-click on a 7210 SAS NE and choose Properties. The properties form for the NE opens.

2

Click on the Globals tab, then on the Oper Group tab.

3

Click Create, or choose an operational group and click Properties.

4

Configure the required parameters.

You cannot modify the Displayed Name of an operational group after creation.

Note: Oper Group Member and Open Group LAG Member are mutually exclusive, meaning, the Oper Group can be associated only to a Port or LAG, not both. Multiple port/LAG can be associated to a monitor an Oper Group.

5

For existing operational groups, click on the Monitor Oper Group Ports tab to view a list of the ports that are monitoring the group.

6

For existing operational groups, click on the Monitor Oper Group LAGs tab to view a list of the LAGs that are monitoring the group.

7

Save your changes and close the form.

END OF STEPS

12.10 To enable and configure global Cflowd sampling on an NE

12.10.1 Purpose

Cflowd collects statistical data based on traffic flows. A flow is a series of packets in a user session. Cflowd data is used to monitor network usage trends and detect security threats. You can use the NFM-P to configure Cflowd sampling for:

- global traffic that passes through an NE; you must enable Cflowd globally on an NE before you can configure Cflowd collectors. You can create multiple Cflowd collectors on an NE. Each collector is deleted if Cflowd sampling is disabled.
- traffic that passed through an ISA-AA group and/or partitions within the group as well as TCP performance data collection for AA applications and application groups; see 13.9 “To configure Cflowd collectors on an ISA-AA group or partition” (p. 422) .
- unicast and multicast traffic on a routing network interface, IES/VPRN L3 access interface, IES/VPRN tunnel interface, and VPRN network interface; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) , 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249) , 78.28 “To configure an L3 access interface on an IES site” (p. 2472) , 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) , or 79.44 “To configure a network interface on a VPRN site” (p. 2603) .
- traffic that passes through a group interface on an IES or VPRN; see 78.19 “To configure a group interface on an IES” (p. 2449) or 79.37 “To configure a group interface on a VPRN” (p. 2586).
- traffic that is redirected to Cflowd using an ACL IP or IPv6 filter policy; see 51.5 “To configure an ACL IP filter policy” (p. 1671) or 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

You can create up to five Cflowd sampling rate profiles per system. Sampling rate profiles allow you to set different sampling rates on a per interface basis. This allows you to choose the appropriate sampling rate based on the services or type of traffic that is to be received or sent through the associated interface.

12.10.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the CFLOWD tab.

3

Set the CFLOWD State parameter to Enabled.

4

Configure the required parameters in the CFLOWD Attributes panel.

5

Add a collector.

Note: A global NE Cflowd configuration can contain up to eight collectors.

1. Click on the Collector tab and click Create. The Cflowd Collector Configuration (Create) form opens.
2. Configure the required parameters.

   The parameters vary based on the value of the Version parameter.
3. Save your changes and close the form.

6

If you added a collector with the value of the Version parameter set to version 9 or 10, you can create filters that specify which flows are sent to a collector.

1. Select a version 9 or 10 collector and click Create. The Cflowd Collector Configuration (New Instance) (Create) form opens.
2. Click on the Export Filters and Interface List tabs and click Create. The Export Filter Interface List (Create) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

7

As required, create a Cflowd sampling rate profile.

1. Click on the Sample Profile tab. The Cflowd Sample Profile (New Instance) form opens.
2. Configure the parameters.
3. Save your changes and close the form.
4. Repeat these steps for up to five profiles.

8

Save your changes and close the form.

END OF STEPS

12.11 To enable the automatic selection of an RD on an NE

12.11.1 Before you begin

Since an RD must be unique on each PE in the network, you can allocate either a route distinguisher that you manually select or an NE-selected route distinguisher for each service. When you configure an auto-RD on an NE, a Type-1 RD is automatically allocated by the NE based on the community range that you configure.

12.11.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the Service tab.

3

Click on the BGP Groups tab and configure the Enable Auto Route Distinguisher parameter.

4

Configure the IP Address parameter in the Auto RD panel.

5

Configure the parameters in the Community Range panel.

6

Save your changes and close the forms.

END OF STEPS

12.12 To configure a Service MAC list

12.12.1 Purpose

Use this procedure to configure a list of MAC addresses to be excluded from auto learn mac protect (ALMP). ALMP is used along with the Restrict Protected Source, Discard-Frame action to prevent specified MACs from moving across PEs. This is used to prevent loops or MAC spoofing attacks.

The list can be added to the FIB configuration of a VPLS L2 Access Interface or Spoke SDP Binding, to an SHG, or to a PW Template.

12.12.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the Service tab.

3

Click on the Service MAC List tab and click Create. The Service MAC List (Create) form opens.

4

Click Create to add MAC addresses to the list.

5

Add additional MAC addresses as needed and click OK.

6

Save your changes and close the forms.

END OF STEPS

12.13 To add a span of control to an NE

12.13.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Spans tab and click Add. The Select Span(s) - Network Element form opens with a list of available spans.

3

Choose one or more spans to apply to the NE.

4

Save your changes and close the forms.

END OF STEPS

12.14 To configure load balancing

12.14.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab. The Load Balancing tab is displayed.

3

Configure the required load balancing parameters.

To enable system-wide enhanced multicast load balancing to send multicast traffic over a LAG, set the Multicast Enhanced Load Balancing parameter to true.

To improve the LAG spraying of VLL service packets when ECMP and LAG hashing are performed by the same NE, set the Service ID Lag Hashing parameter to true.

For 7250 IXR, Release 19.7 and later, you can specify the CRC polynomials used in ECMP and LAG hashing, by setting the Hash Polynomial values. This is a global setting that affects system-wide ECMP and LAG hashing. The same value cannot be used for both ECMP and LAG. The Hash Polynomial settings use the hash seed result obtained by shifting to optimize the load balancing distribution of packet flows under ECMP and LAG. The number of bits the hash seed result is shifted is specified by the Hash Seed Shift parameter at the card level.

4

Save your changes and close the form.

END OF STEPS

12.15 To configure proxy ARP and proxy node discovery for an NE

12.15.1 Purpose

Use this procedure to create a list of MAC addresses that can be advertised in EVPN for proxy ARP or proxy ND IP entries on VPLS services.

12.15.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then the Service sub-tab, then the Proxy ARP/ND sub-tab.

3

To configure a MAC list entry:

1. Click on the MAC list tab.
2. Click Create. The Proxy ARP/ND MAC List (Create) form opens.
3. Configure the MAC List name parameter.
4. To add a MAC address to the list, click Create, configure the MAC Addr parameter, and click OK.

4

Save the changes and close the forms.

END OF STEPS

12.16 To configure a node discovery profile on an NE

12.16.1 Purpose

Use this procedure to create a profile for node discovery using OSPF Type 10 LSA TLV. When a network element profile is assigned to a VPRN OSPFv2 area site, its information will be advertised via LSA type 10 opaque.

See 79.22 “To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance” (p. 2555) and 28.67 “To create an OSPF area” (p. 963) for information about configuring an OSPFv2 area site on a VPRN.

Note: When a node is discovered using OSPF Type 10 LSA TLV, the network element information will not be automatically added to the routing table (RIB) or forwarding table (FIB). If the NE IP address provided in the profile needs to be visible to the network, you need to configure a loopback interface or physical interface with the same IP address and add it to the OSPFv2 area.

12.16.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then the Node Discovery sub-tab.

3

Click Create. The Node Discovery Profile (Create) form opens.

4

Configure the required parameters.

5

Save the changes and close the forms.

END OF STEPS

12.17 To enable or disable 802.1X

Note: Before you can create an 802.1X policy, 802.1X must be enabled on the NE. Before you can configure 802.1X on an Ethernet access port, 802.1X must be enabled on the NE and an 802.1X policy must be created and distributed to the NE.

Note: To create and distribute 802.1X policies to the NEs that use 802.1X, see Chapter 59, “802.1x policies” .

Note: To configure 802.1X on Ethernet access ports, see 16.24 “To configure Ethernet ports” (p. 599) in Chapter 16, “Port and channel object configuration” .

12.17.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab.

3

Click on the PAE tab and set the Administrative State parameter to Up.

To disable 802.1X on a managed NE, set the Administrative State to Down.

4

Save your changes and close the form.

END OF STEPS

12.18 To configure an exclusive policy editing restriction on an NE

12.18.1 Before you begin

You can enable or disable an exclusive policy editing restriction on one or more NEs. When enabled, this prevents any local policy definitions that reside on the NE from being changed by a global policy update. You can also remove any exclusive edit locks that were previously set for the selected NE(s).

Note: This procedure can only be performed by admin users or users with an assigned policy management scope of command role.

Note: This functionality is not supported on all devices.

12.18.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

Note: You can perform this procedure on multiple NEs if needed. Select the required NEs from the navigation tree, right-click on any one of them, and then select Properties.

2

Click on the Globals tab. The Load Balancing tab is displayed.

3

Click on the Exclusive Edit tab.

4

Perform one of the following:

a. Configure the Use Exclusive Edit parameter to enable or disable an exclusive policy editing restriction for the selected NE(s).

b. Click Reset Exclusive Edit and click Yes to remove any exclusive edit locks that were previously set for the selected NE(s).

Note: The Reset Exclusive Edit option is also available in global routing policies under the Local Definitions tab.

5

Save your changes and close the form.

END OF STEPS

12.19 To configure active card alarms on an NE

12.19.1 Before you begin

You can configure supporting devices to view NE alarms through the CLI.

12.19.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab. The Load Balancing tab is displayed.

3

Click on the Alarm Management tab and configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

12.20 To configure a TWAMP server

12.20.1 Before you begin

You can enable or disable Two-Way Active Measurement Protocol on a server and view session and connection statistics.

Note: You cannot configure a TWAMP server after an NE upgrade until the NE is fully resynchronized.

12.20.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab. The Load Balancing tab is displayed.

3

Click on the OAM tab and scroll to the OAM TWAMP panel.

4

Configure the Administrative Status parameter.

Note: To make changes to TWAMP server parameters, the Administrative Status parameter must be set to Disabled.
If you are configuring a 7705 SAR, Release 6.1 R5 or later, you can modify the TWAMP prefix Description parameter when the Administrative Status parameter is set to Enabled.

5

Configure the required parameters in the Server Configuration panel.

6

If you are configuring a 7705 SAR, the Reflector Configuration panel is displayed.

Expand the Reflector Configuration panel and configure the Reflector Test Session Timeout (Seconds) parameter.

7

Configure the required parameters in the Time stamping Configuration panel.

8

Click on the Prefix tab.

9

Perform one of the following:

a. Click Create to add a new prefix. The Prefix TWAMP Server (Create) form opens, with the General tab displayed.

Note: The NFM-P allows you to create an unlimited number of prefixes. However, an NE can only accept a maximum of 100 prefixes. If more than 100 prefixes are created in the NFM-P, the affected NE will issue a deployment error.

b. Select an existing prefix and click Properties. The Prefix TWAMP Server (Edit) form opens, with the General tab displayed.

10

Configure the required parameters.

11

Click on the Statistics tab.

12

Select either Srv Conns Stats (Assurance) or Prefix Srv Stats (Assurance), as required, from the Select Object Type drop-down menu.

13

Select either Past 4 Hour(s) or No Filter, as required, from the filter selector.

14

Click Statistics Policies and select Statistics Policy from the menu. The Statistics Policy form opens, with the General tab displayed.

15

Configure the required parameters.

16

Click on the Thresholds tab.

17

Perform one of the following, depending on your selection in Step 12 .

a. Configure the Prefix Srv Stats (Assurance) threshold parameters.

b. Configure the Srv Conns Stats (Assurance) threshold parameters.

18

Save your changes and close the forms.

19

To view TWAMP statistics on a server:

1. Perform Step 1 to Step 3 of this procedure.
2. Click on the Prefix tab.
3. Select either Srv Conns Stats (Assurance) or Prefix Srv Stats (Assurance), as required, from the Select Object Type drop-down menu.
4. Select either Past 4 Hour(s) or No Filter, as required, from the filter selector.
5. Click Collect or Collect All. The system compiles the statistics and displays the available records in the table.

   Note:

   The objects listed in the table control the behavior of the TWAMP connections that match the defined IP prefix.
6. Select the required record from the table and click Properties. The Statistics Record is displayed.
7. Save your changes and close the forms.

END OF STEPS

12.21 To enable LLDP on an NE

12.21.1 Before you begin

LLDP is a not a routing protocol, but instead, a neighbor-discovery protocol that allows an NE to advertise its identity and capabilities to other NEs attached to the same physical IEEE 802.1 LAN. As such, it is configured in a different manner than standard routing protocols.

LLDP also permits information that the device discovers about peer devices to be stored. LLDP is only applicable for devices using Ethernet connectivity. To enable LLDP, you must configure the protocol at both the system level and at the port level. This is done using the NFM-P equipment navigation tree.

When LLDP is enabled on a device, it sends and receives LLDP messages on all of the physical interfaces that are enabled for LLDP transmission. These messages are sent periodically to ensure that information is accurate. These messages are stored on the local device for a configurable amount of time, and after this time has expired, the information is discarded.

The NFM-P uses the information stored in the applicable LLDP tables on the node to automatically discover the physical topology in the network. You can use this information to examine the L1/L2 topology and perform appropriate diagnostics and troubleshooting. Care must be taken when configuring the Port ID SubType parameter, as the setting may affect the ability to build the Layer 2 topology map using LLDP. See the NE documentation for more information.

In LLDP, a single LLDP Protocol Data Unit is transmitted in a single Ethernet frame. The basic LLDP PDU consists of a header, followed by a variable number of information elements known as TLVs that each include fields for Type, Length, and Value. Type identifies what kind of information is being sent. Length indicates the length of the information string. Value is the actual information sent. Each LLDP PDU includes three mandatory TLVs followed by optional TLVs.

Mandatory TLVs include:

- Chassis ID: represents the identification of the device transmitting the LLDP frame
- Port ID: represents the identification of the port transmitting the LLDP frame
- TTL: represents the length of time the receive frame shall be valid

Optional TLVs include:

- Port Description: represents the description of the port
- System Name: this is the administratively-assigned name of the device
- System Description: this is a textual description of the device
- System Capabilities: this identifies the capabilities of the device and its function (such as router, switch, repeater, etc.)

LLDP also supports multiple transmission scopes. The destination MAC address in the LLDP PDU determines how a frame is propagated through the network, thereby determining the LLDP message scope. The following table and figure identify a set of destination MAC address and describes the different transmission scopes associated with each address.

Table 12-1 MAC Addresses and transmission scopes

Name	Value	Purpose
Nearest Bridge	01-80-C2-00-00-0E	Propagation constrained to a single physical link
Nearest non-TPMR bridge	01-80-C2-00-00-03	Propagation constrained by all bridges other than TPMR; intended for use within provider bridged networks
Nearest Customer Bridge	01-80-C2-00-00-00	Propagation constrained by customer bridges

Figure 12-1 LLDP Multiple Transmission Scopes

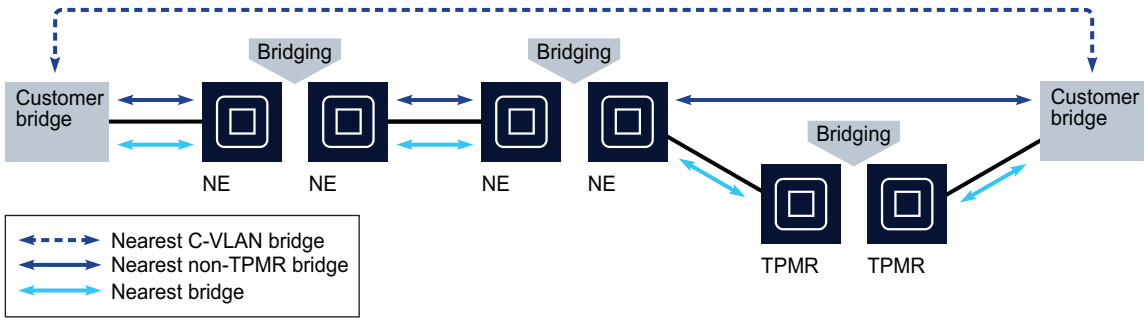

12.21.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Enable LLDP.

1. Click on the Globals tab and then on the LLDP tab.
2. Configure the required parameters.

   Care must be taken when configuring the Port ID SubType parameter, as the setting may affect the ability to build the Layer 2 topology map using LLDP. See the NE documentation for more information.

3

Save your changes and close the forms. See Table 12-2, “LLDP related tasks” (p. 360) for additional procedures related to LLDP information.

Table 12-2 LLDP related tasks

Task	See
Configure LLDP on supported ports to send the ifDesc value in place of the ifAlias value over port TLV. Additionally, IPv6 addresses can be configured on the Transmit MGMT Address tab.	16.24 “To configure Ethernet ports” (p. 599)
Configure LLDP on OmniSwitch Ethernet ports.	16.56 “To configure OmniSwitch Ethernet ports” (p. 647)
Configure LLDP on module cards on a 7705 SAR-M/ME or 7705 SAR-H	15.83 “To configure a module card on a 7705 SAR-M/ME or 7705 SAR-H” (p. 543)

END OF STEPS

12.22 To configure the BFD flap detection interval on an NE

12.22.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the BFD tab.

3

Configure the Operational State Transition Interval (seconds) parameter.

4

Save your changes and close the form.

END OF STEPS

12.23 To enable a Q in Q untagged SAP on an NE

Note: The Enable Q in Q Untagged Sap parameter is not supported on the 7950 XRS as it is enabled by default.

12.23.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Ethernet tab.

3

Enable the Enable Q in Q Untagged Sap parameter.

When you enable the Enable Q in Q Untagged Sap parameter, the NFM-P allows the creation of the following two default SAP types:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- The SAP type *.* functions as a default SAP for double-tagged frames in a Q on Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

4

Save your changes and close the form.

END OF STEPS

12.24 To configure DHCPv6 Advertise messages on an NE

12.24.1 Purpose

Complete this procedure to specify which DHCPv6 client applications on an NE are sent NO_ADDRS_AVAILABLE messages when the server runs out of addresses.

12.24.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Subscriber Management tab.

3

On the DHCPv6 panel, specify one or more DHCPv6 Advertise message format options for the Advertise Message parameter, and configure other required parameters if available.

4

Save your changes and close the form.

END OF STEPS

12.25 To configure Python script protection on an NE

12.25.1 Purpose

Complete this procedure to configure an NE with a protection key, shared with a Python script policy, and with source and destination URLs for a protected Python script. For information on Python script policies, see 58.1 “Python policies and Python script policies” (p. 1807) .

12.25.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Subscriber Management tab.

3

Configure the required parameters on the Python Script Protection panel.

The Protection Key parameter value must match the protection key value of the specified Python script policy. The two values are validated against each other when Python script protection is enabled.

4

Click on the Protect Script button.

5

Save your changes and close the form.

END OF STEPS

12.26 To configure home LAN extension functionality on an NE

12.26.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Subscriber Management tab.

3

Configure the Router Target AS Number parameter on the LAN Extension panel.

4

Save your changes and close the form.

END OF STEPS

12.27 To configure ISA service chaining on an NE

12.27.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Subscriber Management tab.

3

Configure the MAC Prefix parameter on the ISA Service Chaining panel.

4

Save your changes and close the form.

END OF STEPS

12.28 To configure optimized HTTP redirects on an NE

12.28.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Subscriber Management tab.

3

Configure the Optimized Mode parameter on the CPM HTTP Redirect panel.

4

Save your changes and close the form.

END OF STEPS

12.29 To configure sFlow on an NE

12.29.1 Purpose

Perform this procedure to configure an sFlow receiver and counters on a supporting NE. Data for sFlow is collected on ports; for information about enabling sFlow on ports, see 16.24 “To configure Ethernet ports” (p. 599).

12.29.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the SFLOW tab.

3

Configure SFLOW receivers.

1. Select a receiver in the Receiver tab and click Properties. The Receiver (Edit) form opens.
2. Configure the required parameters.
3. Save and close the form.

4

Note: If you are configuring egress counter queues for 7250 IXR, go to Step 5.

Configure egress counter policers and queues

1. Click on the Egress Counter Map tab, then on the Policers tab.
2. Click Create. The SFlow Egress Counter Policer form opens.
3. Configure the required parameters.
4. Save and close the form.
5. Click on the Queues tab and then click Create. The SFlow Egress Counter Queue form opens.
6. Configure the required parameters.
7. Save your changes and close the form.

5

Configure egress counter queues for 7250 IXR

1. Click on the Egress Counter Map tab.
2. In the Queues tab, click Create. The SFlow Egress Counter Queue form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

6

Note: If you are configuring ingress counter policers for 7250 IXR, go to Step 7.

Configure ingress counter policers and queues

1. Click on the Ingress Counter Map tab, then on the Policers tab.
2. Click Create. The SFlow Ingress Counter Policer form opens.
3. Configure the required parameters.
4. Save and close the form.
5. Click on the Queues tab and then click Create. The SFlow Ingress Counter Queue form opens.
6. Configure the required parameters.
7. Save your changes and close the form.

7

Configure ingress counter policers for 7250 IXR

1. Click on the Ingress Counter Map tab.
2. In the Policer tab, click Create. The SFlow Ingress Counter Policer form opens.
3. Configure the required parameters.
4. Save and close the form.

8

Click on the Counter Pollers tab to view counter pollers.

9

Click on the Statistics tab and select Receiver Statistics (sFlow) in the Object Type drop-down to view sFlow statistics.

10

Save your changes and close the form.

END OF STEPS

12.30 To configure ANYsec encryption on an NE

12.30.1 Purpose

Perform this procedure to configure an sFlow receiver and counters on a supporting NE. Data for sFlow is collected on ports; for information about enabling sFlow on ports, see 48.5 “To configure a local connectivity association” (p. 1458).

12.30.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the ANYSec tab.

3

Configure the required general parameters.

4

Click on the Security Termination Policy tab and click Create. The Security Termination Policy (Create) form opens.

1. Configure the required parameters.
2. Save your changes and close the form.

5

Click on the Encryption Group tab and click Create. The Encryption Group (Create) form opens.

1. Select a security termination policy in the Security Termination Policy Name panel.
2. Select a MACsec connectivity association site in the Connectivity Association panel.
3. Configure the other parameters, as required.
4. Click on the Peer tab and click Create. The Peer (Create) form opens.
5. Configure the required parameters.
6. Save your changes and close the forms.

6

Save your changes and close the form.

END OF STEPS

12.31 To create a chassis-level PBB configuration

12.31.1 Purpose

Perform this procedure to provide the PBB MAC name configuration at the NE level.

12.31.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

Configure the required parameters.

4

Click on the MAC Name tab and click Create. The PBB MAC Name (Create) form opens.

5

Configure the required parameters.

6

Save your changes and close the form.

END OF STEPS

12.32 To configure serving network information on an NE

12.32.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Serving Network tab.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

12.33 To configure L2TP on an NE

12.33.1 Purpose

Perform this procedure to configure a non-multi-chassis tunnel ID range on an NE.

12.33.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the L2TP tab.

3

Configure the Non MC Tunnel ID Range Start and End parameters.

The Non MC Tunnel ID Range Start value must not be greater than the End value.

4

Save your changes and close the form.

END OF STEPS

12.34 To configure WLAN GW redundancy on an NE

12.34.1 Purpose

Perform this procedure to configure a virtual chassis identifier on an NE for the purpose of fail-over redundancy in WLAN GW configurations. You must perform this procedure on each NE in the redundancy configuration. The NEs must all use the same virtual chassis ID string or dual homing key.

12.34.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the WLAN GW Sys Config tab.

3

Configure the Virtual Chassis ID and Max Number of GTP Session Held parameters.

4

Save your changes and close the form.

END OF STEPS

12.35 To configure call-trace debug storage on an NE

12.35.1 Purpose

Perform this procedure to configure debug log file storage information for NEs associated with trace profiles.

12.35.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Call-Trace tab.

3

Configure the Maximum Number of Log Files and Primary Compact Flash Card parameters.

4

Save your changes and close the form.

END OF STEPS

12.36 To configure the RADIUS CoA port on an NE

12.36.1 Purpose

Perform this procedure to configure the RADIUS change of authorization port on an NE.

Note: You cannot configure the RADIUS CoA port on the 7950 XRS.

12.36.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the RADIUS tab.

3

Configure the RADIUS CoA Port parameter.

Note: If the routing instance is bound to a RADIUS proxy server and the RADIUS CoA port is set to 1812, all CoA messages are received by the RADIUS proxy server and dropped by port 1812.

4

Save your change and close the form.

END OF STEPS

12.37 To configure data persistence on an NE

12.37.1 Purpose

Use this procedure to configure data persistence on an NE that allows you to store dynamic object information about an application in a file on the NE.

12.37.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Node Persistence tab.

3

Configure the required parameters on each of the following application panels:

- ANCP
- Application Assurance
- DHCP Server
- NAT Port Forwarding

- Python Policy Cache
- Subscriber Management
- DHCP Lease Time Threshold

Note: Configuring the File Location parameter for any of the above applications enables data persistency for that application. Ensure that the specified location has sufficient disk space to handle the volume of dynamic objects that will persist on it. If the system is unable to store persistency data because the device is inaccessible or full, an alarm is raised.

4

Save your changes and close the form.

END OF STEPS

12.38 To configure DNS security extensions

12.38.1 Before you begin

The NFM-P supports configuring DNS security options on 7450 ESS, 7750 SR, and 7950 XRS nodes.

12.38.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the DNS tab.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

12.39 To enable or disable ICMP extensions on a NE

12.39.1 Purpose

Use this procedure to configure the use of vendor-specific ICMP extensions on the 7950 XRS, 7450 ESS, 7705 SAR, and 7750 SR NEs. Network operators can use the information provided by the ICMP extensions to diagnose routing problems. By default, ICMP extensions are disabled on devices.

12.39.2 Steps

1

On the equipment tree, right-click on a NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the ICMP tab.

3

Configure the Vendor-Specific ICMP Extensions parameter.

4

Save your changes and close the form.

END OF STEPS

12.40 To configure a PPPoE Intermediate Agent on an NE

12.40.1 Purpose

Perform this procedure to enable a supporting OmniSwitch NE as an Intermediate Agent (IA) for Point-to-Point Protocol over Ethernet (PPPoE).

12.40.2 Steps

1

On the equipment navigation tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the PPPoE-IA tab.

3

Configure the required parameters.

Configure port parameters

4

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

5

Expand to the port level and click on a port object. The Physical Port (Edit) form opens.

6

Click on the PPPoE-IA tab. The PPPoE-IA configuration form opens.

7

Configure the required parameters.

8

Save your changes and close the forms.

Configure LAG parameters

9

On the equipment navigation tree, expand Network→*NE*→Logical Groups→LAGs→LAG *n*.

10

Right-click on the LAG *n* object and choose Properties. The LAG (Edit) form opens.

11

Click on the PPPoE-IA tab. The PPPoE-IA configuration form opens.

12

Configure the required parameters.

13

Save your changes and close the forms.

END OF STEPS

12.41 To create an FPE

12.41.1 Before you begin

Use this procedure to configure a Forward Path Extension on an NE for traffic pre-processing. FPE is supported on the following NEs:

- SR-7/12/12e (chassis mode D)
- SR-1e/2e/3e
- SR-a4/a8
- ESS-7/12 (chassis mode D) with Mixed-Mode enabled
- XRS-20/16/40

Creation of a PXC or LAG path is required for enabling PwPort or a VXLAN. For a LAG path, two LAGs are required. LAG members should be PXC SubPorts.

When the FPE path is composed of PXC SubPorts (PXC path), the direction b sub-port is the egress sub-port. When the FPE path is composed of a pair of LAGs, the LAG configured as the FPE Xb LAG represents the egress direction, regardless of the direction of the SubPorts in the LAG.

To use the FPE for PW Port FPE on an Epipe service site, the Pw Port parameter must be enabled. To use the FPE for VXLAN tunnel termination, Vxlan termination must be enabled.

Supported breakout ports can be associated to a PXC. The PXC ports can added to either a LAG or an FPE. The LAG with the PXC can be associated to the FPE.

12.41.2 Steps

1

On the equipment tree, right-click on a NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Forward Path Extension tab.

3

Click on the SPD ID sub-tab and configure the required SPD ID parameters.

4

Configure an FPE:

1. Click on the FPE tab.
2. Click Create. The Forwarding Path Extension (Create) form appears.

   If the FPE is to be assigned to a GTP group interface, place a check mark in the Extensions Enabled parameter.
3. In the Path panel, select one or more of the PXC, FPE Xa LAG, and FPE Xb LAG objects.

4. Configure the required parameters.

   If the FPE is to be assigned to a GTP group interface or to a bonded group interface configuration, place a check mark in the Extensions Enabled parameter.

5. Save your changes and close the form.

END OF STEPS

12.42 To configure satellite file transfer

12.42.1 Purpose

Perform this procedure to specify the satellite file transfer protocol to be used for the boot process when transferring boot images and configuration files from the 7x50 host to the Ethernet satellite.

12.42.2 Steps

1

On the equipment tree, right-click on a NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Satellite Features tab and configure the Satellite File Transfer parameter.

3

Save your changes and close the form.

END OF STEPS

12.43 To create a port template

12.43.1 Before you begin

Perform this procedure to create a port template on an NE, Release 15.0 R4 or later, that supports Ethernet satellites. A port template allows the satellite port type to be configured as client or uplink. A port template can only be edited or deleted if it is not associated to an active satellite.

12.43.2 Steps

1

On the equipment tree, right-click on the NE with the Ethernet satellite shelf and choose Properties.

2

Click on the Port Templates tab and click Create. The Satellite Port Template Create form opens.

3

Configure the required parameters.

4

Save your changes and close the form.

5

Click on the Satellite Port Template and click Properties. The Satellite Port Template form opens.

6

Set the Admin State parameter to Down and configure the rest of the required parameters on the General tab.

7

Click on the Ports tab, click on a Satellite Physical Port Id, and click Properties.

8

Set the Port Role parameter.

For 100 GE client ports, 1/1/c3/1 and 1/1/c4/1 must both have the same roles.

Role of none means the port is neither a data port nor uplink satellite port.

For 10 GE uplink port template, you must set the 100 GE ports to role of none, and there is a maximum of 16 sequential 10 GE ports that can be configured as uplink ports.

9

Configure the other parameters, as required.

10

Save your changes and close the form.

11

Associate the port template to a satellite. See 15.98 “To associate a port template on an Ethernet satellite shelf” (p. 560).

END OF STEPS

12.44 To configure the Sender-ID TLV of a CFM PDU for an NE

12.44.1 Before you begin

On supporting NEs, you can specify a user-defined ID in the Sender ID TLV of CFM PDUs. To use this option, the ID-Permission parameter on the MEG Service form for the NE MEG must be set to “chassis”; see 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109).

12.44.2 Steps

1

On the equipment tree, right-click on a device and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

Click on the Globals tab and on the OAM tab.

3

Configure the parameters in the Sender ID panel.

When the Sender ID Type parameter is set to Local, you can configure the Sender ID Name parameter with a user-defined ID to be used in the Sender ID TLV of CFM PDUs.

4

Save your changes and close the form.

End of steps

12.45 To configure the global EVPN proxy ARP and node discovery on an NE

12.45.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

Configure the System EVPN Proxy ARP ND parameter.

4

Save your changes and close the form.

End of steps

12.46 To configure a no-service loopback port on the 7210 SAS

12.46.1 Before you begin

When you configure a port loopback with MAC swap on a 7210 SAS NE, you must configure a no-service port. You can select a physical port or a virtual port. Not all chassis types support virtual no-service ports.

If you select a physical port as a no-service port, the port cannot contain any SAPs.

For virtual no-service ports, the number of available ports and the port names vary depending on the chassis type. Virtual ports are not displayed on the navigation tree.

You cannot select the same no-service port for more than one function.

See the NE documentation for more information.

12.46.2 Steps

1

On the equipment tree, right-click on a 7210 SAS NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then click on the Service tab.

3

Select a loopback no-service port. Perform one of the following:

a. Select a physical loopback no-service port.

b. Select a virtual loopback no-service port and configure all other required parameters.

The availability of virtual ports varies depending on the chassis type and card type.

4

Save your changes and close the forms.

END OF STEPS

12.47 To configure CFM DMM version 1 interoperability on the 7210 SAS

12.47.1 Steps

1

On the equipment tree, right-click on a 7210 SAS NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then click on the CFM tab.

3

Configure the DMM Version 1 Compatibility parameter.

4

Save your changes and close the form.

END OF STEPS

12.48 To configure two WRED slopes on a 7210 SAS

12.48.1 Purpose

The 7210 SAS-D and 7210 SAS-M support three WRED slopes per queue:

- High priority/in-profile TCP
- Low priority/out-of-profile TCP
- Non-TCP

Perform this procedure to specify whether only the high priority and low priority slopes are used, and non-TCP slopes are ignored.

12.48.2 Steps

1

On the equipment tree, right-click on a 7210 SAS NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the QoS tab.

3

Configure the Use WRED Slopes parameter.

4

Save your changes and close the form.

END OF STEPS

12.49 To configure frame-based accounting for QoS policies on a 7210 SAS

12.49.1 Steps

1

On the equipment tree, right-click on a 7210 SAS NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Frame Based Accounting tab.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

12.50 To configure the global system resource profile on a 7210 SAS or 7250 IXR

12.50.1 Purpose

Perform this procedure to allocate the system resources on 7210 SAS or 7250 IXR NEs.

The functional areas for which you can allocate resources vary, depending on the chassis type. See 6.5.13 “System resource profile” (p. 220) and the NE documentation for more information about system resource allocation.

For the 7210 SAS-R, system resources for the G8032 Fast Flood, VPLS Up MEP Resource, and Router functions are configured on the device properties form using this procedure. Other system resources for the 7210 SAS-R are allocated using a policy; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382) .

To configure port-based scheduling on the 7210 SAS-Mxp or 7210 SAS-R, see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384) .

Note: System hardware resources are shared. Do not assign more than the available resources. Functions are disabled if you assign more than the amount of available resources.

Supporting releases of the 7250 IXR allow custom forwarding path resource allocation for resource-intensive functions such as multicast, ACL classification criteria, policers, etc.; see 12.57 “To configure forwarding path options or resource allocation on a 7250 IXR” (p. 388).

12.50.2 Steps

1

On the equipment tree, right-click on a 7210 SAS or 7250 IXR NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then click on the System Resource Profile tab.

3

Configure the required parameters.

Configuration changes for some parameters require a reboot of the NE to take effect; see the NE documentation for more information.

If you are configuring ECMP profiles for a 7250 IXR NE, consider the following:

- There are no default ECMP profiles for 7250 IXR.
- For 7250 IXR with SR OS, ECMP profiles can be modified.
- For 7250 IXR with SR OS, MPLS profiles are supported.

If you are configuring a 7210 SAS NE, consider the following:

- To configure MAC authentication, set the Ingress MAC Authentication Resource and Egress MAC Authentication Resource parameters to 1 in the MAC Authentication panel.
- To enable port-based access-ingress policies on a 7210 SAS-R6 or R12, set the Sap Scale Mode parameter to High.
- The SAP Egress Aggregate Meter parameter is configurable when the Total ACL Resource parameter in the SAP Egress Internal ACL panel is set to a value of 1 or less.
- If you set the SAP Egress Aggregate Meter parameter and then configure egress meters on SAPs, you cannot reset the parameter to a lower value unless you remove all egress meter configurations for all SAPs on the NE.
- In the SAP Ingress Internal ACL panel, the sum of the values for the IPv4 Resource, IPv4/128-bit IPv6 Resource, and 64-bit IPv6 Resource parameters must not be greater than the value specified for the Total ACL Resource parameter.
- In the SAP Egress Internal ACL panel, the sum of the values for the MAC Resource, IPv4 Resource, 128-bit IPv6 Resource, and 64-bit IPv6 Resource parameters must not be greater than the value specified for the Total ACL Resource parameter.
- In the SAP Ingress QoS Resources panel, the sum of the values for the MAC Resource, IPv4 Resource, and IPv4/v6 Resource parameters must not be greater than the value specified for the Total QoS Resource parameter.
- Each 128-bit address consumes two entries of the available entry resources.
- The value for the Up MEP Resource parameter must not be greater than the value for the Total CFM Resource parameter.
- For the Max IP Subnets parameter, IPv4 subnets use one entry each and IPv6 subnets use two entries each.

- You cannot set the Mbs Pool Type parameter to Node when ports on the NE are decommissioned.

4

Save your changes and close the form.

END OF STEPS

12.51 To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC

12.51.1 Purpose

System resource profiles for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx or 7210 SAS-S/Sx VC are configured using the NFM-P policy distribution framework (except for the G8032 Fast Flood, VPLS Up MEP Resource, and Router functions; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380)). See Chapter 49, “Policies overview” for more information about NFM-P policies. You can configure up to 16 system resource profile policies and distribute them to 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx or 7210 SAS-S/Sx VC devices. A selected policy is assigned to the card slot of the device; see 15.71 “To select system resource profile policies for the 7210 SAS-R or 7210 SAS-S/Sx VC” (p. 529) .

To configure port-based scheduling on the 7210 SAS-R using the system resource profile, see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384) .

Note: The Total QoS Resource and MAC Authentication parameters are not supported on 7210 SAS-S/Sx VC.

See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) and the NE documentation for more information about system resource allocation.

12.51.2 Steps

1

Choose Policies→QoS→SROS QoS→7210 System Resource Profile Policy from the NFM-P main menu. The 7210 System Resource Profile Policies form opens.

2

Click Create or choose a policy and click Properties.

3

Configure the required parameters.

When the policy is assigned to a card slot and the SAP Aggregate Meter parameter is modified, you must reboot the card or the node for the modified value to take effect.

To enable Primary VLAN for egress MIPs on supporting 7210 SAS-R or 7210 SAS-S/Sx VC NEs, you must allocate resources to the Bi Directional MIP Egress parameter.

In the SAP Ingress Internal ACL panel, the sum of the values for the MAC Resource, IPv4 Resource, 64-bit IPv6 Resource, and 128-bit IPv4/v6 Resource parameters must not be greater than the value specified for the Total ACL Resource parameter.

In the SAP Egress Internal ACL panel, the sum of the values for the MAC Resource, MAC/IPv4 Resource, 128-bit IPv6 Resource, and 64-bit MAC/IPv6 Resource parameters must not be greater than the value specified for the Total ACL Resource parameter.

In the SAP Ingress QoS Resources panel, the sum of the values for the MAC Resource, MAC/IPv4 Resource, and 64-bit MAC/IPv4/v6 Resource parameters must not be greater than the value specified for the Total QoS Resource parameter.

The sum of the values for the Down MEP Resource and Up MEP Resource parameters must not be greater than the value for the Total CFM Resource parameter.

Each 128-bit address consumes two entries of the available entry resources.

To configure MAC authentication, set the Ingress MAC Authentication Resource and Egress MAC Authentication Resource parameters to 1 in the MAC Authentication panel.

4

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

12.52 To configure chassis MAC address on the 7210 SAS-S/Sx VC

12.52.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then on the Virtual Chassis tab.

3

Configure the required VC Configuration parameters.

The user must specify a range of MAC addresses to use if LAGs are used, with the following conditions:

- The number of MAC addresses range from 1 to 126.
- The first MAC address in the block of MAC addresses is used as the chassis MAC address, and the subsequent MAC addresses are used to assign MAC addresses to LAGs configured in the VC.
- The number of MAC addresses specified must match up with the number of LAGs in use.
- A LAG will be operationally down if no MAC addresses are available for allocation.

- By default, the LAG MAC addresses are allocated from the active CPM MAC address pool (the base MAC address and the number of MAC addresses available in this pool are printed on the label of the chassis).
- The user can use a fixed MAC address for a LAG that is independent of the CPM chassis in use. This simplifies provisioning by removing the need to change the MAC addresses on peer nodes in the network when MC-LAG is in use.

4

Click Ok to save your changes and close the form.

5

Reboot the NE.

Configuration changes require a reboot of the NE to take effect; see the NE documentation for more information. After reboot, the base MAC address is replaced with the configured VC MAC address.

END OF STEPS

12.53 To configure port-based scheduling on the 7210 SAS

12.53.1 Purpose

Using the global system resource profile, you can configure access ports on the 7210 SAS-Mxp or 7210 SAS-R to use either port-based or SAP-based egress scheduling. When port-based scheduling is enabled, the port supports eight egress queues, and all SAPs on the port share the same set of eight queues.

Queues for port-based scheduling are configured in the 7210 port access egress policy assigned to the port; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556) . See 70.15 “Sample QoS configuration on the 7210 SAS” (p. 1954) for more information about configuring scheduling on the 7210 SAS-Mxp and 7210 SAS-R.

SAP-based egress queues are not available when port-based scheduling is enabled. See the NE documentation for more information.

You must enable port-based scheduling before configuring egress meters.

12.53.2 Steps

1

On the equipment tree, right-click on the required NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then click on the System Resource Profile tab.

3

Configure the Port Scheduler Mode (Configured) parameter.

Changes to the Port Scheduler Mode (Configured) parameter require a reboot of the NE to take effect. The Port Scheduler Mode (Active) parameter displays the parameter setting currently in effect on the NE.

4

Save your changes and close the form.

END OF STEPS

12.54 To configure buffer management for the 7210 SAS

12.54.1 Before you begin

For the 7210 SAS-M and 7210 SAS-T, queue buffer resources are by default allocated equally to all ports on the device. You can decommission unused ports and assign their buffer resources to other ports. Decommissioned ports cannot be turned up or used for any function. Configurations deployed to a device with decommissioned ports will not take effect on the decommissioned ports.

Port decommissioning and buffer resource assignment are configured using profiles. You can configure up to 30 buffer management profiles per NE. Changes to profile configuration take effect when the NE is rebooted. The status for a profile indicates whether it is currently active or if it requires a reboot before it takes effect.

For the 7210 SAS-T, you must configure the system resource profile before you configure buffer management; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

See the NE documentation for more information about buffer management.

12.54.2 Steps

1

On the equipment tree, right-click on a 7210 SAS NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Port Decommission tab. A list of buffer management profiles is displayed.

3

Delete one or more buffer management profiles, if required. Choose the profile and click Delete.

4

Configure a buffer management profile:

1. Click Create or choose a profile and click Properties. The Buffer Management form opens.
2. Configure the ID and Description parameters.
3. Click on the From tab, then click Add. The Select form opens.
4. Select one or more ports to decommission and click OK. The Select form closes.

   The selected ports are decommissioned the next time the NE reboots.
5. Click on the To tab, then click Add. The Select form opens.
6. Select one or more ports and click OK. The Select form closes.

   The buffer resources of the decommissioned ports are assigned to the selected ports the next time the NE reboots.
7. Save your changes and close the form.

5

Repeat Step 4 to create additional profiles. You can create up to 30 profiles.

Note: You must reboot the NE for changes to buffer management profiles to take effect. The status for a profile indicates whether it is currently active or if it requires a reboot in order to take effect.

END OF STEPS

12.55 To configure 7210 SAS-R device properties for MVPN

12.55.1 When to use

When you configure 7210 SAS-R sites in a VPRN for multicast VPN using a P2MP LSP, you must configure the device properties appropriately. Also, PIM or IGMP must be enabled on the site.

Perform the following to configure the device properties for a 7210 SAS-R site.

12.55.2 Steps

1

On the equipment tree, right-click on the required 7210 SAS-R NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab.

3

To allocate system resources to enable the P2MP function, click on the System Resource Profile tab and configure the MPLS P2MP Resource parameter.

4

Specify a no-service P2MP bud port. You can use a virtual port or a physical port. Click on the Service tab, then on the General tab, and perform one of the following:

Note: You can only configure a no-service P2MP bud port when LDP is disabled on the site routing instance. You must set the Administrative State to Down on the General tab of the LDP properties form. See 28.51 “To configure global-level LDP” (p. 941).

a. Select the Use Virtual P2MP Bud Port parameter and choose a virtual port. The number of available ports and the port names vary depending on the card type and release. Virtual ports are not displayed on the navigation tree.

b. Select a physical P2MP Bud port. The physical port cannot contain any SAPs.

5

Save your changes and close the form.

To configure MVPN on the 7210 SAS-R site, see 79.28 “To configure an MVPN VRF instance on a VPRN site” (p. 2570).

To enable PIM on the 7210 SAS-R site, see 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003).

To enable IGMP on the 7210 SAS-R site, see 28.104 “To configure an IGMP site on a router” (p. 1017).

END OF STEPS

12.56 To configure IP fragmentation for 7210 SAS

12.56.1 When to use

For 7210 SAS-Mxp, native IP fragmentation for IP packets exceeding the configured MTU is not supported by default. You can configure CPM to extract oversized IPv4 packets from the datapath, fragment them using the system CPU, and insert the fragmented packets back into the datapath.

Note: Even when IP fragmentation is enabled, packets that exceed the configured MTU are dropped if the Do not Fragment (DF) bit is set in the IP header.

Perform the following to configure the IP fragmentation for a 7210 SAS-Mxp site.

12.56.2 Steps

1

On the equipment tree, right-click on the required 7210 SAS-Mxp NE and choose Properties. The Network Element (Edit) form opens.

2

Click the General tab.

3

Configure the Allow CPU Fragmentation parameter.

4

Save your changes and close the form.

END OF STEPS

12.57 To configure forwarding path options or resource allocation on a 7250 IXR

12.57.1 Purpose

Supported releases of the 7250 IXR allow custom forwarding path options such as 802.1x Host Authentication, or resource allocation for resource-intensive functions such as multicast, ACL classification criteria, policers, etc.

A reboot of the NE is required for changes made using this procedure to take effect.

Depending on the 7250 IXR release, the displayed tabs and parameters may vary.

12.57.2 Steps

1

On the equipment tree, right-click on a supporting 7250 IXR NE and choose Properties. The Network Element (Edit) form opens.

2

Click the Globals tab, then click the Forwarding Path tab.

3

Configure the parameters, as required.

4

Click the Options subtab and configure the required parameters.

Note: Forwarding path option resources are shared. When some functions are enabled, others may be disabled. The NFM-P displays an error message when mutually exclusive functions are selected.

Note: After you change the administrative state, enable the Reboot Required check box to reboot the NE and apply the changes.

5

Click the Resource Allocation subtab and configure the required parameters on the following subtabs:

- General
- Filter Ingress IPv4
- Filter Ingress IPv6
- Policer Application Group. To configure the policer scale:
  1. Choose an item in the list and click Properties. The Policer Application Group form opens.
  2. Click on the Policer Application subtab, choose an item, and click Properties. The Application Group (Edit) form opens.
  3. Configure the required parameters.

     When the Override Administrative State and Override Operational State parameters on the General subtab are enabled, you cannot modify the scale values.

After the configuration, the Reboot Required parameter is enabled.

6

Click the LPM subtab and configure the scale options.

7

Save your changes and close the form.

8

Reboot the NE for configuration changes to take effect.

END OF STEPS

12.58 To configure policer allocation

12.58.1 When to use

Perform this procedure to enable policer allocation on 7250 IXR-J2 NE variants to support a higher SAP scale. In this policer allocation model, the 8, 16, and 32 policer allocation model is disabled, and the 2, 4, and 8 policer allocation model is enabled.

12.58.2 Steps

1

On the equipment tree, right-click on a supporting 7250 IXR NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, the Forwarding Path tab, then on the Resource Allocation subtab.

3

The override administrative state must be disabled, and a reboot is required, before you can make any modifications.

1. Click on the General subtab.
2. Disable the Override Administrative State parameter, and enable the Reboot Required check box.
3. Save your changes.

4

Configure the new policer application group, which refers to the 2, 4, 8 policer model; perform 12.57 “To configure forwarding path options or resource allocation on a 7250 IXR” (p. 388), Step 5.

5

Configure the policer application bank:

1. Click on the Policer Application Bank subtab.
2. Click Search to populate the list if required, or choose a policer model, then click Properties. The Application Bank form opens.
3. Configure the required parameters.
4. Save your changes.

6

Configure the application bank group if it refers to the old policer model (8, 16, 32) instead of the new policer model (2, 4, 8). The Override Administrative State must be disabled and a reboot is required before you do this:

1. Perform Step 1 to Step 3 to access the Resource Allocation subtab, then the General subtab, to disable the Override Administrative State parameter, and enable the Reboot Required check box.
2. Save your changes.
3. Perform Step 5 to configure the parameters on the Application Bank form to refer to the new policer model.
4. Perform Step 1 to Step 3 to access the Resource Allocation subtab, then the General subtab, to enable the Override Administrative State parameter, and enable the Reboot Required check box.

7

Create a 7250 IXR SAP access ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544).

8

Perform 12.57 “To configure forwarding path options or resource allocation on a 7250 IXR” (p. 388) to reconfigure the scale options on the LPM subtab, and to change the application bank group to the correct policer model, if required.

9

Create the SAP access ingress policy and assign it to the SAP:

1. Perform 50.28 “To configure a SAP access ingress policy” (p. 1538) to create a SAP access ingress policy.
2. Perform 77.9 “To modify a VPLS” (p. 2259) to associate the SAP access ingress policy to the SAP on the QoS tab.

10

Create a 7250 SROS network ingress policy and assign it to the network interface:

1. Perform 50.43 “To configure a 7250 SROS Network Ingress policy” (p. 1575).
2. Associate the 7250 SROS network ingress policy to the network interface on the Policies tab of the Network Interface – Routing Instance (Edit) form; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information about creating a network interface.

11

Save and close the forms.

END OF STEPS

12.59 To configure watermark settings on a 7705 SAR

12.59.1 Steps

1

On the equipment tree, right-click on a 7705 SAR NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Security tab.

3

Configure the required parameters in the Water Mark Settings (Active Sessions) panel.

4

Save your changes and close the form.

END OF STEPS

12.60 To configure QoS ingress aggregate rates on the 7705 SAR-M/ME

12.60.1 Steps

1

On the equipment tree, right-click on the 7705 SAR-M/ME NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the QoS tab.

3

Configure the Aggregate Rate parameter in the Access Ingress panel and in the Network Ingress panel.

4

Save your changes and close the form.

END OF STEPS

12.61 To launch an MCT on a Wavence SA NE connected to a 7705 SAR

12.61.1 Before you begin

The NFM-P supports an integrated MCT from the port or NE level on the 7705 SAR-8 and the 7705 SAR-18. The MCT is distributed with a NEtO client in the NFM-P. The MCT can be launched from the port that is associated with the MW link in the navigation tree or from the MW Link Member (Edit) form. You can configure radio parameters using the MCT.

Consider the following before you launch the MCT:

- Java version 1.6 is required to launch the MCT.
- The MCT is only supported on a Windows client; however, it is not supported on Windows 7.
- Only one MCT session can be open at a time.

12.61.2 Steps

1

On the equipment tree, right-click on a discovered 7705 SAR-8 or 7705 SAR-18 NE and choose External Element Manager.

2

Choose the corresponding MW member from the list and click Launch MCT.

3

Enter your login credentials and click OK. A dialog box appears.

Note: If the client system does not have the required MCT packages, the packages are downloaded from the NFM-P and saved in a third-party folder.

4

Click Accept. The MCT initialization starts, and the MCT main view screen appears.

Note: If the MCT is not installed on the client system, the MCT installer window opens. After the MCT is installed, the MCT main view screen appears for all subsequent launches, with the MCT version adjusted to the NE version from which the MCT launched.

END OF STEPS

12.62 To start and stop a Webview or Secure Webview session on an OmniSwitch

12.62.1 Steps

1

Choose one of the following to start a Webview or Secure Webview session:

a. On the equipment tree, right-click on a discovered OmniSwitch and choose Launch Webview or Launch Secure Webview. A web browser starts and the Webview or Secure Webview login screen appears.

b. On the equipment tree, right-click on a discovered OmniSwitch and choose Equipment Window. The Equipment Window form opens. Select the Launch Webview or Launch Secure Webview button on the form. A web browser starts and the Webview or Secure Webview login screen appears.

2

Enter your user name and password. After a successful login, the Chassis Management Home Page appears.

3

To stop a Webview or Secure Webview session:

1. Click on the Log Out text at the top of the Webview or Secure Webview window. A Confirm Log Out dialog box appears.
2. Click Log Out to end the Webview or Secure Webview session.

See the appropriate OmniSwitch *Switch Management Guide* for information about configuring and using Webview or Secure Webview.

END OF STEPS

12.63 To configure the dying gasp alarm on an OmniSwitch

12.63.1 Before you begin

Perform the following procedure to configure the dying gasp alarm on supporting OmniSwitch NEs.

Note: To view the dying gasp alarm on the NFM-P, the OmniSwitch NE must be managed using SNMP v2 with the community string set to public.

12.63.2 Steps

1

On the equipment tree, right-click on a supported OmniSwitch icon and choose NE Sessions→Telnet Session. The Telnet Session window opens.

2

Perform Step 5 of 8.8 “To commission an OmniSwitch for NFM-P management” (p. 254) to configure the NE to be managed using SNMP v2 with a community string configured to public.

Note: If the Backup/Primary power supply fails, a major alarm will be raised. This alarm will be sent only to the first two SNMP stations configured on the NE. “Show snmp station” lists the management station address. NFM-P ip should be first or second.
The "Loopback 0" address should not be used for the source IP address field on the NE.

END OF STEPS

12.64 To configure shelf craft port IP address on an 1830 VWM device

12.64.1 Steps

1

Perform one of the following:

a. From the device:
   1. On the equipment tree, right-click on the 1830 VWM device object and choose Properties. The Network Element (Edit) form opens.
   2. Click on the Shelf Craft Port tab. The craft port IP address of all the shelves except the OPS shelf are listed.
   3. Choose the shelf for which you need to configure the craft port IP address and click Properties. The Shelf Craft IP (Edit) form opens.

b. From the shelf:
   1. On the equipment tree, expand Network→1830 VWM.
   2. Right-click on the shelf object (not OPS shelf) and choose Properties. The Shelf (Edit) form opens.
   3. Click on the Shelf Craft Port tab.

2

Configure the parameters in the Shelf Craft IP Details panel.

Note:

- The craft port IP address and the craft port gateway cannot be the same.
- Ensure that the craft port IP address and the craft port gateway are in the same subnet.

3

Save your changes and close the form.

END OF STEPS

12.65 To configure an Auto-ID range for policies

12.65.1 Before you begin

NEs that support next-generation CLI use the policy name as the key for internal system reference. For these NEs, a policy ID is also configurable. If a policy ID is not configured, the system auto-assigns a numerical policy ID. When the policy is distributed to an NE, the auto-assigned ID is generated from a specified range of values.

Perform this procedure to configure a numerical range on the NE for auto-assigned policy IDs.

Before you change an existing range configuration on an NE, you must delete any policies on the NE that have IDs within the existing range.

If you configure a range that comprises a policy ID that already exists on the NE, deployment fails.

12.65.2 Steps

1

On the equipment tree, right-click on a supporting NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the Policies tab.

3

For each policy type, configure the start and end values for the ID range.

4

Save your changes and close the form.

END OF STEPS

12.66 To enable USB support on a 7250 IXR

12.66.1 Purpose

Perform this procedure to enable the USB storage option on all 7250 IXR variants, Release 21.2 R1 and later, except for 7250 IXR-s, where USB support is enabled by default.

12.66.2 Steps

1

On the equipment tree, right-click on a supporting NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the USB Support tab.

3

Set the USB Administrative State parameter to Enabled.

4

Save your changes and close the form.

END OF STEPS

12.67 To globally enable or disable Packet Byte Offset on a 7250 IXR

12.67.1 Before you begin

QoS Packet Byte Offset (PBO) allows scheduler policies to operate based on frame sizes with modified overhead values, to automatically account for different service configurations.

Perform this procedure to globally enable or disable QoS PBO on supporting 7250 IXR NEs. The actual offset values for PBO are configured in 7250 SROS VLAN QoS policies or 7250 SROS Port QoS policies; see 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594) and 50.61 “To configure a 7250 SROS Port QoS policy” (p. 1603).

12.67.2 Steps

1

On the equipment tree, right-click on a supporting NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then the Policies tab.

3

In the QoS Packet Byte Offset panel, configure the Enable Packet Byte Offset parameter.

4

Save your changes and close the form.

END OF STEPS

12.68 To configure GNSS receiver functions on supported IXR and SR NEs

12.68.1 Before you begin

This procedure applies to 7250 IXR-e, 7250 IXR-e2, 7250 IXR-e2c, 7250 IXR-R6, 7250 IXR-R6d, 7250 IXR-R6dl, or FP5-enabled SRs.

Note: For FP5-enabled SRs, the A/gnss and B/gnss management ports are supported only on the 7750 SR-2se. For all other FP5 chassis, only the A/gnss port is supported.

12.68.2 Steps

1

On the equipment tree, right-click on the NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Inventory tab and select Management Port (Physical Equipment).

3

Select a GNSS port and click Properties. The Management Port (Edit) form opens.

4

Click on the GNSS tab to view GNSS location information and to configure the required parameters.

5

Close the form.

END OF STEPS

12.69 To configure global-level UNP

12.69.1 Steps

1

On the equipment tree, right-click on the NE and choose Properties. The Network Element (Edit) form opens.

2

Click the Globals tab, and then the UNP tab.

3

Click the Classification tab. Click Create. The UNP Classification LLDP (Edit) form opens.

4

Configure the required parameters.

5

Save your changes and close the form.

END OF STEPS

12.70 To configure an UNP at port-level

12.70.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Port *n/n/n*.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Click Policies tab, and then the General tab.

4

Configure the UNP Port Type parameter.

5

Click UNP Port tab and configure the required parameters.

6

Save your changes and close the form.

END OF STEPS

12.71 To configure DHCPv6 snooping on an NE

12.71.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the DHCPv6 Snooping tab.

3

Configure the DHCPv6 Snooping parameter.

4

Save your changes and close the form.

END OF STEPS

12.72 To configure the log encryption key

12.72.1 Before you begin

Use this procedure to create a log encryption key. The log encryption key prevents external attackers from obtaining critical network information from the configuration file, such as users, IP addresses, and protocols. This procedure applies to 7250 IXR-6, 7250 IXR-R6, 7250 IXR-10, 7250 IXR-R4, 7250 IXR-e, 7250 IXR-e2, 7250 IXR-e2c, 7250 IXR-ec, 7250 IXR-s, 7750 SR, 7950 XRS, and VSR NEs.

12.72.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Log Configuration tab.

3

Configure the Encryption Key parameter.

4

Click OK and close the form.

END OF STEPS

Ring group configuration procedures

12.73 To create a ring group

12.73.1 Before you begin

You can use the ring group creation function to:

- indicate a VLAN ring topology
- indicate a VPLS ring topology
- configure the properties of the ring group to provide VLAN Internet, BTV (MVR), and L2 VPN (TLS) services
- group devices by geographic region

Consider the following before you create a ring group.

- Ensure that the devices are commissioned, as described in Chapter 8, “Device commissioning and management” .
- Ensure that mediation policies are configured to allow CLI access to the managed devices, as described in Chapter 9, “Device discovery” .
- If you want specific access interfaces to be part of a VLAN service, ensure that the parent device of the interface is a member of the ring group.

12.73.2 Steps

1

Choose Ring Group from the navigation tree view selector. The navigation tree displays the Ring Group view.

2

Right-click on the Network object and choose Create Ring Group. The Ring Group (Create) form opens.

3

Configure the required parameters.

4

Click Apply. If you set the Ring Group Type parameter to VPLS in Step 3 , go to Step 6 .

5

Click on the TLS tab to configure the TLS parameters. The Ethertype and Jumbo Frame parameters are configurable when the Enabled parameter is selected.

6

Perform one of the following to add a device to the ring group.

a. To add a device using the current properties form:

1. Click on the Group Members tab click Create. The Select Network Elements form opens.
2. Choose one or more NEs and click OK. The selected devices are listed on the Ring Group (Create) form.
3. Click OK and then click Yes. The device objects are listed in the navigation tree under the new ring group object.

b. To add a device from the navigation tree:

1. Right-click on the device object in the Equipment view of the navigation tree and choose Add To Ring Group. The Select Group form opens to display a list of ring groups.
2. Choose the new ring group and click OK. The device object is listed under the ring group in the Ring Group view of the navigation tree.

Note: All device types can be added as members of a VLAN ring group. Only devices that support VPLS can be added to a VPLS ring group.

7

To apply a span of control to a ring group:

1. Click on the Spans tab and click Add. The Select Span(s) - Ring Groups form opens with a list of available spans.
2. Choose one or more spans of control to apply to the ring group.
3. Save your changes and close the form.

8

As required, provision VLAN services for the ring group as described in Chapter 75, “VLAN service management” .

END OF STEPS

12.74 To remove a device from a ring group or a ring group

12.74.1 Steps

1

Choose Ring Group from the navigation tree view selector. The navigation tree displays the Ring Group view.

2

Navigate to the required ring group object. The path is Network→Ring Group→Device.

3

To remove a device from the ring group, right-click on the device object under the ring group and choose Remove From Ring Group. The device is removed from the ring group.

4

Removing a ring group deletes the VLAN services that are associated with the ring group.

To remove a ring group:

1. Perform Step 3 to remove all devices from the ring group.
2. Right-click on the Ring Group object and choose Remove Ring Group.
3. Save your changes and close the form.

END OF STEPS

12.75 To configure the global sampling rate on an NE

12.75.1 Purpose

Perform this procedure to configure the global sampling rate from the chassis at the NE level.

12.75.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

Click on the Mirror tab.

4

Configure the Global Sampling Rate parameter.

5

Save your changes and close the form.

END OF STEPS

13 Logical group object configuration

Logical group object configuration using the NFM-P

13.1 Overview

13.1.1 Logical group objects

Logical group objects are children of the device object. They appear below the device object in the navigation tree. The following logical group objects are created automatically in the navigation tree after the device is discovered.

- CCAG
- ISA-AA Group
- ISA-Tunnel Group
- ISA-LNS Group
- ISA-NAT Group
- ISA-Video Group
- LAG
- IGH

Properties forms for logical group objects are accessed using the NFM-P navigation tree.

This chapter contains the procedures to configure logical group objects using the navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the navigation tree.

13.1.2 Working with CCAG objects

CCAGs are navigation tree objects located below device icons. CCAGs are configured manually using the CCAG object navigation tree menu and subsequent forms. For proper CCAG configuration, a VSM-CCA card must be present in the NE. In the navigation tree, a VSM-CCA card can be opened to show its ports. The ports are indicated by VSM Port x/y/z, which shows port type and the ID of the port.

You must configure the following in the CCAG configuration forms:

- General properties such as CCAG ID, Description, CCA Rate Enabled, CCA Rate, Access Adapt QoS, and Administrative State
- CCAG MDA Members which are compatible ports that can belong to a CCAG.

When you create a CCAG, the NFM-P creates virtual paths to be used as interconnections to bind services together. Two unidirectional paths are created: Alpha and Beta. For each path, three virtual ports are created: one for SAP-SAP connections, one for SAP-NET connections, and one for NET-SAP connections. The last two virtual ports are used to bind a service and a network interface together.

A maximum of eight cards can be added to a CCAG, with a maximum of eight CCAGs per NE.

13.1.3 Working with ISA-AA groups

ISA-AA groups provide AA load balancing, scaling, and redundancy on NEs that contain multiple ISA-AA MDAs or ESA VMs. ISA-AA MDAs and ESA VMs cannot be members of the same ISA-AA group.

ISA-AA redundancy protects against card failure and minimizes service disruption during maintenance or protocol signature upgrades.

The NFM-P supports AA group configuration on the 7450 ESS and 7750 SR. An NE can have up to seven ISA-AA groups.

Note: A single-slot chassis does not support AA configuration.

ISA-AA group configuration requires the following:

- group number, which can be set only during group creation
- up to seven primary group members and one optional backup member
- one or more forwarding classes to divert to the AA group

You can use multiple ISA-AA groups to scale the application of AA policies, and can divide an AA group into partitions that are customized for specific VPN services. Each partition has an AA policy, and supports the configuration of custom protocols, applications, and application groups. Within a partition, you can use application service options, or ASOs, to create multiple application QoS policies (AQPs). AA partitions support statistics collection and data reporting functions.

Note: When you create an ISA-AA group or partition, a default AA policy is automatically created for the group or partition. When you delete a group or partition, the default policy is also deleted.

See Chapter 87, “Application assurance” for more information about AA groups, partitions, and policies. See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) for information about creating ISA-AA groups and partitions. See 13.8 “To configure an AA subscriber policy override on an ISA-AA group or partition” (p. 421) and 13.9 “To configure Cflowd collectors on an ISA-AA group or partition” (p. 422) for information about configuring additional ISA-AA group and partition functions.

13.1.4 Working with ISA-Tunnel groups

ISA-tunnel groups provide redundancy for IPsec tunneling and encryption functions when multiple ISA-IPSEC MDAs or ESA VMs are installed on an NE. ISA-IPSEC redundancy protects against card failure and minimizes service interruption during maintenance or protocol signature upgrades. You can configure up to 16 ISA-Tunnel groups on one NE.

An ISA-tunnel group can operate in one of the following modes:

- primary/backup active member (ISA-IPSEC MDA only)
- multiple active members (MDAs or ESA VMs)

The primary/backup configuration provides fault tolerance in the event of an MDA failure. An ISA-tunnel group in a primary/backup configuration can contain a maximum of two MDAs.

ISA-IPSEC MDAs and ESA VMs cannot be members of the same ISA-Tunnel group.

Using multiple active members provides load balancing in addition to fault tolerance. This mode of operation is supported on a 7750 SR in chassis mode D or higher, and on a 7450 ESS in mixed mode. You can include up to 16 members in an ISA-tunnel group when the use of multiple active members is enabled.

You can configure an ISA-tunnel group on a 7750 SR or 7450 ESS in mixed mode to act only as an IKEv2 responder during a MC IPsec switchover. This is the automatic behavior for dynamic tunnels, but not for static tunnels. See Chapter 34, “IPsec” for more information about MC IPsec.

13.1.5 Working with ISA-LNS groups

The NFM-P supports the creation and configuration of ISA-LNS groups. ISA-LNS groups provide LNS PPP session termination on 7750 SR NEs. You can assign an ISA-LNS group to a tunnel group profile or a tunnel profile. When an operational L2TP tunnel is established, peers that are associated with the ISA-LNS group are automatically created. Session traffic is automatically balanced across the available active ISA broadband application MDAs in the group. You can add ISA broadband application MDAs to an ISA-LNS group, and up to four ISA-LNS groups on each NE. See 13.12 “To configure an ISA-LNS group” (p. 426) for information about how to create and configure an ISA-LNS group.

You must configure the following on the ISA-LNS groups configuration form:

- General properties, such as Group Number, Description, and Administrative State
- ISA-LNS Group members

Note: You can configure an ISA broadband application MDA only on an IOM3-XP module in a 7750 SR -7 or 7750 SR-12, Release 20.10 and earlier.

13.1.6 Working with ISA-NAT groups

An ISA-NAT group provides a redundant NAT function for routing instances using ISA Broadband Applications MDAs or ESA VMs. See Chapter 30, “NAT” for information about ISA-NAT group configuration.

13.1.7 Working with ISA-Video groups

ISA-Video groups are created to provide packet buffering and packet processing in support of the IPTV video features.

When configured in the router, video ISAs are logically grouped into video groups. An ISA-Video group allows more than one video ISA to be treated as a single logical entity for a specific application, where the system performs a load-balancing function when it assigns tasks to a member of the group.

ISA-Video groups provide a redundancy mechanism to guard against hardware failure within a group. ISA-Video groups pool the processing capacity of all the group members and increase the application throughput because of the increased packet processing capability of the group.

You must configure the following on the ISA-Video groups configuration form:

- General parameters
- ISA-Video Group members
- ESA Video Group members

An ISA-Video group supports ISA or ESA-VM video group members, and not a mix of group members under the same ISA-Video group.

You can add up to four MDAs to an ISA-Video group, and up to four ISA-Video groups on each NE. All members of an ISA-Video group are primary members.

13.1.8 Working with WLAN GW groups

A WLAN GW group is used to represent multiple hardware adapters as a single entity, allowing for warm redundancy between multiple WLAN GW MDAs or ESAs. Only one WLAN GW group can be created per NE. This group can be numbered 1 to 4, but this number must be exclusive and cannot be in use by an ISA-NAT group.

The IOM is the child of the WLAN GW group. The required IOMs should be added before the WLAN GW group is turned administratively up.

WLAN GW members are created automatically by the NE when the WLAN GW group is administratively up, and deleted automatically when the administrative state is down. The number of members is equal to the active IOM limit.

You can globally list WLAN GW groups from the Manage→ ISA Functions→ ISA-WLAN main menu, or from from Manage→ Equipment→Equipment→ISA-WLAN GW Group (WLAN Gateway).

WLAN GW tunnels

WLAN GW tunnels are created dynamically by the NE when the first UE attempts to connect. The tunnel is used to backhaul traffic from the access point or residential gateway to the WLAN GW. The NE creates tunnels on a per-access point or per-residential gateway basis. WLAN GW tunnels terminate on the WLAN-GW MDA or ESA.

WLAN GW tunnels are viewable on the configuration form for base routing instances, VPRN routing instances, group interfaces, and WLAN GW groups. The WLAN GW tunnels tab lists the WLAN GW tunnels belonging to a given routing instance. User equipment information is viewable from the properties form for individual WLAN GW tunnels. GTP session information (including bearer information) is viewable from the properties form for individual user equipment items. You can globally list WLAN GW tunnels from the Manage→ ISA Functions→ ISA-WLAN main menu.

13.1.9 Working with IGH objects

IGHs are navigation tree objects located below the device icon. IGHs are configured manually using the configuration forms available when you choose Create IGH from the IGH object navigation tree contextual menu.

You create an IGH to group together IP links and POS links so that if a configured number of links go out of service for any reason, the remaining links in the IGH go out of service too. This causes the routing protocols to re-converge to switch from the primary path to an alternate path.

The following requirements and restrictions apply to IGHs.

- IGHs are supported only on network links.
- IGHs are supported only on SONET/TDM interfaces with PPP auto encapsulation.
- A port or channel needs to be bound to a router interface after the member is added to an IGH.
- You can assign a port to only one IGH.

13.1.10 Working with LAG objects

LAGs are navigation tree objects located below a device icon which aggregate multiple network connections in parallel to increase throughput beyond what a single connection can sustain, and to provide redundancy in case one of the links fails.

The following minimum configuration is required to enable LACP.

- Enable LACP at either end of the LAG group.
- Set one end of the LAG group as LACP active.

You must configure the following in the LAG configuration forms:

- General properties, such as LAG description, configured address, encapsulation type, and administrative state
- Link aggregation group parameters, such as port threshold, port threshold action, and dynamic cost
- LACP parameters, such as LACP mode, LACP transmit interval, actor administration key, LACP transmission standby, and LACP selection criteria
- LAG members, which are the compatible ports that can belong to a LAG

Because all ports can have their own MAC address, when ports are part of a LAG, the LAG must have an MAC address.

The port configuration of the first port added to the LAG is used to compare with subsequently added ports. If a discrepancy is found with a newly added port, that port is not added to the LAG.

The maximum number of ports you can add to or remove from a LAG is 64. The number of ports depends on the NE type; see the NE documentation for more information. All ports added to a LAG must have the same parameter settings.

Only ports belonging to one LAG subgroup are considered eligible members of a LAG and can be selected as active links.

A LAG can be configured with weighted per-link hashing, which allows a more balanced distribution of subscribers and services across LAG links. For example, significant differences in subscriber or SAP bandwidth requirements could lead to unbalanced traffic on LAG egress. You can configure each service or subscriber on a LAG with one of three unique classes and a weight value to distribute services and subscribers across LAG links.

A LAG can be configured with LAG adaptive load balancing to resolve a traffic imbalance between LAG member ports on the FP-based SR family of NEs. Adaptive load balancing monitors the traffic utilization of each LAG member port and identifies whether the traffic needs to be shifted toward certain ports based on a tolerance value. If the tolerance value is exceeded, the system identifies the best way to shift traffic to resolve this imbalance. Adaptive load balancing is mutually exclusive with per-link-hashing.

On supporting devices, you can enable LACP tunneling to transparently forward LACP packets. LACP tunneling is configured on Ethernet ports; see 16.24 “To configure Ethernet ports” (p. 599).

See 6.13 “OmniSwitch” (p. 230) for OmniSwitch-specific LAG information.

See the *NSP Wavence Device Support Guide* for Wavence-specific LAG information.

LAG link mapping profiles

You can use a LAG link mapping profile to assign a specific LAG egress port to a SAP or interface, and control how egress traffic is handled if the specified port fails. You can then configure the specified port to perform admission control and QoS contract with the knowledge that all traffic for the SAP or interface will egress on that port. You can also specify a secondary port to handle traffic in the event that the primary port fails. See 13.21 “To create a LAG link mapping profile” (p. 443) for more information about creating LAG link mapping profiles.

LAG utilization TCAs

You can enable TCA on a LAG object and determine whether the LAG utilization value exceeds a TCA policy threshold. The following nodes support enabling TCA on LAG objects:

- 7450 ESS
- 7750 SR
- 7950 XRS
- Wavence SM

You can associate a TCA policy with a LAG object from the TCA policy configuration form. See 19.5 “To apply a TCA policy to objects using the object properties forms” (p. 711) for information about applying an existing TCA policy to an object from the object properties form. The NFM-P compares the utilization value at each statistics collection to the threshold value in the TCA policy and raises an alarm when the LAG utilization value crosses the specified limit. You need to enable performance statistics collection on the LAG object to configure a TCA policy. See Chapter 19, “TCA” for more information about TCA.

L1 LAGs support collection of scheduled and on-demand statistics, and plotting of real-time and historical statistics. L2 and Ethernet LAGs support collection of scheduled statistics and plotting of historical statistics.

BFD on Ethernet LAGs

On supporting NEs, you can configure an Ethernet LAG to use BFD to speed up the detection of link failures. When BFD is enabled on a LAG, micro-BFD sessions are automatically created for each link in the group.

For information about configuring BFD on a LAG, see 13.17 “To modify a LAG” (p. 435) . For information about viewing micro-BFD sessions on a LAG, see 13.22 “To view micro-BFD sessions on a LAG” (p. 445) .

13.2 Workflow to manage logical group objects

13.2.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure logical group objects. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: Logical group objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

13.2.2 Stages

1

As required, configure CCAG objects. See 13.4 “To configure a CCAG” (p. 414) for more information.

2

As required, create and configure ISA-AA groups and ISA-AA partitions. See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) for more information. Additionally, perform the following, as required:

a. Configure AA subscriber statistics collection. See 13.6 “To configure AA subscriber statistics collection on an ISA-AA group or partition” (p. 418) for more information.

b. Configure AA special study objects. See 13.7 “To configure special study objects on an ISA-AA group or partition” (p. 419) for more information.

c. Configure a subscriber policy override on an ISA-AA group or ISA-AA partition. See 13.8 “To configure an AA subscriber policy override on an ISA-AA group or partition” (p. 421) for more information.

d. Configure Cflowd collectors on an ISA-AA group and/or ISA-AA partition. See 13.9 “To configure Cflowd collectors on an ISA-AA group or partition” (p. 422) for more information.

3

As required, configure ISA-Tunnel groups. See 13.10 “To configure an ISA-tunnel group” (p. 424) for more information.

4

As required, configure ISA-Tunnel member-pool groups. See 13.11 “To configure an ISA-tunnel member-pool group” (p. 425) for more information.

5

As required, configure ISA-LNS groups. See 13.12 “To configure an ISA-LNS group” (p. 426) for more information.

6

As required, create and configure ISA-NAT groups. See 30.3 “To configure an ISA-NAT group” (p. 1080) for more information. Additionally, perform the following, as required:

a. Configure statistics on an ISA-NAT group. See 30.15 “To configure statistics on an ISA-NAT group” (p. 1096) for more information.

b. Start or stop a NAT address-pool drain operation. See 30.13 “To start or stop a NAT address-pool drain operation” (p. 1094) for more information.

7

As required, configure ISA-Video groups. See 13.13 “To configure an ISA-Video group” (p. 427) for more information.

8

As required, configure WLAN GW groups. See 13.14 “To configure a WLAN GW group” (p. 428) for more information.

9

As required, configure IGH objects. See 13.15 “To create an IGH and add members” (p. 430) for more information.

10

As required, configure LAG objects:

a. Create and configure a LAG. See 13.16 “To create a LAG” (p. 431) and 13.17 “To modify a LAG” (p. 435) for more information.

b. Create and configure a LAG on an OmniSwitch. See 13.19 “To configure an OmniSwitch LAG” (p. 440) for more information.

c. Create and configure a dynamic LAG member on an OmniSwitch. See 13.20 “To configure OmniSwitch dynamic LAG members” (p. 442) for more information.

d. Create a LAG Link Mapping Profile. See 13.21 “To create a LAG link mapping profile” (p. 443) for more information.

13.3 Workflow to configure weighted per-link hashing on a LAG

13.3.1 Purpose

The following workflow describes the sequence of high-level tasks required to configure weighted per-link hashing on a LAG. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: LAG objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

13.3.2 Stages

1

Configure a LAG object with weighted per-link hashing. See 13.16 “To create a LAG” (p. 431) or 13.17 “To modify a LAG” (p. 435) for more information. Configure the required service sites as LAG members.

2

If required, configure LAG weighted per-link hashing on a subscriber profile. See 64.4 “To configure a subscriber profile” (p. 1840) for more information.

3

As required, configure SAPs for weighted per-link hashing:

- Epipe and Ipipe L2 access interface. See 76.41 “To configure LAG per-link hashing on a VLL Epipe or Ipipe L2 access interface” (p. 2180) for more information.
- IES L3 access interface. See 78.29 “To configure LAG per-link hashing on an IES L3 access interface” (p. 2474) for more information.
- IES group interface SAP. See 78.22 “To configure LAG per-link hashing on an IES group interface SAP” (p. 2462) for more information.
- VPLS L2 access interface (B-VPLS, I-VPLS or M-VPLS). See 77.68 “To configure LAG per-link hashing on a VPLS L2 access interface” (p. 2339) for more information.
- VPRN L3 access interface. See 79.86 “To configure LAG per-link hashing on a VPRN L3 access interface” (p. 2660) for more information.
- VPRN group interface SAP. See 79.40 “To configure LAG per-link hashing on a VPRN group interface SAP” (p. 2599) for more information.

Logical group object configuration procedures

13.4 To configure a CCAG

13.4.1 Steps

1

On the Equipment tree, expand *NE*→Logical Groups→CCAGs.

2

Right-click on the CCAGs icon and choose Create CCAG, or right-click an existing CCAG object and choose Properties. The CCAG (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Click on the CCAG MDA Members tab to configure CCAG MDA members.

6

Click Create or select an existing CCAG MDA member and click Properties. The CCAG MDA Member (Create|Edit) form opens.

1. Select an MDA to add to the CCAG.
2. Save and close the form.

7

Click on the CCAG Paths tab to modify the internal paths for the CCAG. The Alpha and Beta paths are displayed.

1. Select a path and click Properties. The CCAG Internal Path (Edit) form opens.
2. Configure the required parameters.
3. Click on the Path Cross Connects tab to configure cross-connects for the CCAG internal path, as outlined in Step 8 .
4. Save your changes and close the form.

8

Click on the Path Cross Connects tab to modify the path cross-connects associated with the CCAG.

1. Select a cross-connect entry and click Properties. The Cross Connect (Edit) form opens.

2. Configure the required parameters.
3. Click on the Policies tab to specify non-default ingress and egress queue policy settings.
4. To specify an egress slope policy other than the default, click Clear in the Egress Slope Policy panel and select an alternate policy.
5. Configure the Egress Reserved CBS parameter.
6. To specify an ingress slope policy other than the default, click Clear in the Ingress Slope Policy panel and select an alternate policy.
7. Configure the Ingress Reserved CBS parameter.

9

Save your changes and close the forms.

END OF STEPS

13.5 To configure an ISA-AA group and ISA-AA partitions

13.5.1 Purpose

Perform this procedure to create an ISA-AA group that contains one or more partitions. ISA-AA partitions provide ISA-AA group functions that are limited to the partition scope.

Note: When you create an ISA-AA group or partition, a default AA policy is automatically created for the group or partition. When you delete a group or partition, the default policy is also deleted.

13.5.2 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups.

2

Right-click on the ISA-AA Groups icon and choose Create ISA-AA Group, or right-click an existing ISA-AA group object and choose Properties. The ISA-AA Group (Create|Edit) form opens.

3

Configure the required parameters.

If the ISA-AA group is to be added to an ISA-WLAN GW, you must set the Subscriber Scale parameter to Lightweight Internet.

If you are configuring an existing AA group and the associated AA group policy has SAP dependencies, you cannot modify the Partitions parameter until you remove the SAP dependencies on the group policy.

The following parameters are configurable when the Subscriber Scale parameter is set to something other than Lightweight Internet:

- Number of IPv4 Transit Prefix Entries
- Number of Remote IPv4 Transit Prefix Entries
- Number of IPv6 Transit Prefix Entries
- Number of Remote IPv6 Transit Prefix Entries

4

Click Apply.

5

Click on the AA Group Members tab to specify one or more AA group members for the ISA-AA group.

Note: The available group member type depends on the Subscriber Scale parameter: each subscriber scale model is applicable to either ISA or ESA only.

- To add an ISA-AA MDA as an AA group member, see Step 6.
- To add an ESA VM as an AA group member, see Step 7.

6

Click on the AA Group Members tab to specify one or more ISA-AA MDAs for the ISA-AA group.

1. Click Create or select an existing AA group member entry and click Properties. The AA Group Member (Create|Edit) form opens.
2. Configure the Role parameter.

   If the ISA-AA group is to be included for Distributed Subscriber Management in an ISA-WLAN GW group, you must set the Role parameter to Primary for each group member.
3. Select a member MDA.

   After a reboot, an NE rebalances the AA processing load. Because of this, the MDA slot number assigned to an AA subscriber or AA SAP on the AA Subscriber or AA SAP Summary form is displayed as Unassigned if the AA subscribers in the configuration file cannot be accommodated.
4. Save your changes and close the form.

Proceed to Step 8.

7

Click on the AA Group Members tab to specify one or more ESA VMs for the ISA-AA group.

1. Click Create or select an existing AA group member entry and click Properties. The AA Group Member (Create|Edit) form opens.
2. Select a member VM.

   The VM type of the selected member VM must be Application Assurance.
3. Save your changes and close the form.

Proceed to Step 8.

8

Click on the ISA-AA Group Diverted FCs tab to specify one or more traffic forwarding classes to divert to the ISA-AA group.

1. Click Create or select an existing ISA-AA group diverted forwarding class entry and click Properties.The ISA-AA Group Diverted FC (Create|Edit) form opens.
2. Configure the Forwarding Class Name parameter.
3. Save your changes and close the form.

9

Click on the QoS tab to configure QoS policy parameters for the ISA-AA group.

1. Configure the required parameters on the Egress From-Subscriber tab.

   You must deselect the Default check box before you can use a slider to specify a Reserved CBS or Buffer Utilization High Water Mark value.

   The value of the Buffer Utilization High Water Mark parameter must be higher than the value of the Buffer Utilization Low Water Mark parameter.
2. Select a slope policy in the Slope Policy panel.
3. Select a network queue policy in the Network Queue Policy panel.
4. Select a port scheduler policy in the Port Scheduler Policy panel.
5. Click on the Egress To-Subscriber tab and repeat Step 9 1 to 4 .

10

You can configure the following objects on the ISA-AA group if the Partitions parameter was disabled in Step 3 .

- AA subscriber statistics (13.6 “To configure AA subscriber statistics collection on an ISA-AA group or partition” (p. 418))
- AA special study objects (13.7 “To configure special study objects on an ISA-AA group or partition” (p. 419))
- AA subscriber policy override (13.8 “To configure an AA subscriber policy override on an ISA-AA group or partition” (p. 421))

11

To configure Cflowd on the ISA-AA group, click on the CFLOWD tab and perform 13.9 “To configure Cflowd collectors on an ISA-AA group or partition” (p. 422).

12

If the Partitions parameter on the General tab is enabled in Step 3 , click on the ISA-AA Partitions tab to configure AA partitions for the group.

1. Click Create or select an existing partition entry and click Properties. The ISA-AA Group Partition (Create|Edit) form opens.
2. Configure the required parameters.
3. Click on the CFLOWD tab to configure Cflowd on the partition (See 13.9 “To configure Cflowd collectors on an ISA-AA group or partition” (p. 422) .

4. The following objects are configurable on ISA-AA group partitions:
   - AA subscriber statistics (13.6 “To configure AA subscriber statistics collection on an ISA-AA group or partition” (p. 417))
   - AA special study objects (13.7 “To configure special study objects on an ISA-AA group or partition” (p. 419))
   - AA subscriber policy override (13.8 “To configure an AA subscriber policy override on an ISA-AA group or partition” (p. 421))

13

Save your changes and close the forms.

END OF STEPS

13.6 To configure AA subscriber statistics collection on an ISA-AA group or partition

13.6.1 Before you begin

If partitions are disabled on an ISA-AA group, subscriber statistics objects are configured directly on the group. If partitions are enabled on an ISA-AA group, subscriber statistics objects are configured on the individual partitions within the group.

13.6.2 Steps

1

If you are performing this procedure as part of the ISA-AA group or partition configuration described in 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) , go to Step 4 .

2

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups→ISA-AA Group *n*.

3

Right-click on the ISA-AA Group *n* icon and choose Properties. The ISA-AA Group (Edit) form opens.

4

If you are configuring the subscriber statistics objects on an ISA-AA partition, perform the following steps.

1. Click on the ISA-AA Partitions tab.
2. Select a partition in the list and click Properties. The ISA-AA Group Partition (Edit) form opens.

5

Click on the AA Subscriber Stats Objects tab.

1. Click Add and choose one of the following options:
   - Application
   - Application Group
   - Charging Group
   - System Protocol
   - Custom Protocol

   The Select *object_type* form opens and lists the available objects of the chosen type.
2. Select an object and click OK. The Select *object_type* form closes and the AA Subscriber Stats Object Config form opens.
3. Configure the parameters in the Export Method Selection panel.

6

Save your changes and close the forms.

END OF STEPS

13.7 To configure special study objects on an ISA-AA group or partition

13.7.1 Before you begin

If partitions are disabled on an ISA-AA group, special study objects are configured directly on the group. If partitions are enabled on an ISA-AA group, special study objects are configured on the individual partitions within the group.

13.7.2 Steps

1

If you are performing this procedure as part of the ISA-AA group or partition configuration described in 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) , go to Step 4 .

2

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups.

3

Right-click on the ISA-AA Group *n* icon and choose Properties. The ISA-AA Group (Edit) form opens.

4

If you are configuring the special study objects on an ISA-AA partition, perform the following steps.

1. Click on the ISA-AA Partitions tab.
2. Select a partition in the list and click Properties. The ISA-AA Group Partition (Edit) form opens.

3. Click on the AA Special Study tab.

5

To configure special study subscribers on the ISA-AA group or partition, click on the Subscribers sub-tab.

1. Click Create or select an existing special study subscriber and click Properties. The AA Special Study Subscriber Config (Create|Edit) form opens.
2. Configure the AA Subscriber Name and AA Stats Type parameters.
3. Save your changes and close the form.

6

To configure special study subscribers on the ISA-AA group or partition, click on the SAPs sub-tab.

Click on one or more of the following buttons and use the list form that opens to select one or more objects, as required:

- Add Per-SAP Protocols—to add a SAP for per-protocol monitoring
- Add Per-SAP Applications—to add a SAP for per-application monitoring

7

To configure special study spoke SDP bindings on the ISA-AA group or partition, click on the Spoke SDP Bindings sub-tab.

Click on one or more of the following buttons and use the list form that opens to add one or more objects, as required:

- Add Per-Spoke SDP Binding Protocols—to add a SAP for per-protocol monitoring
- Add Per-Spoke SDP Binding Applications—to add a SAP for per-application monitoring

8

To configure special study transit subscribers on the group or partition, click on the Transit Subscribers sub-tab.

1. Click Create or select an existing transit subscriber entry and click Properties. The AA Special Study Transit Subscriber Config (Create|Edit) form opens.
2. Select a transit subscriber.
3. Configure the AA Stats Type parameter.

9

To configure special study ESM subscribers on the ISA-AA group or partition, click on the ESM Subscriber Host sub-tab.

1. Click Create. The AA Special Study ESM Subscriber Host Config (Create) form opens.
2. Configure the AA ESM Subscriber Host Name and AA Stats Type parameters.
3. Save your changes and close the form.

10

Save your changes and close the forms.

END OF STEPS

13.8 To configure an AA subscriber policy override on an ISA-AA group or partition

13.8.1 Before you begin

If partitions are disabled on an ISA-AA group, subscriber policy overrides are configured directly on the group. If partitions are enabled on an ISA-AA group, subscriber policy overrides are configured on the individual partitions within the group.

The AA subscriber policy override is rejected if the subscriber does not have an application profile assigned. You can add an AA subscriber policy override to a SAP or Spoke SDP, but not to an ESM subscriber.

13.8.2 Steps

1

If you are performing this procedure as part of the ISA-AA group or partition configuration described in 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) , go to Step 4 .

2

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups→ISA-AA Group *n*.

3

Right-click on the ISA-AA Group *n* icon and choose Properties. The ISA-AA Group (Edit) form opens.

4

If you are configuring the subscriber policy override on an ISA-AA partition, perform the following steps.

1. Click on the ISA-AA Partitions tab.
2. Select a partition in the list and click Properties. The ISA-AA Group Partition (Edit) form opens.

5

Click on the AA Subscriber Policy Overrides tab.

Note: An AA subscriber policy override is rejected if the subscriber does not have an assigned application profile.

6

Click Create or select an existing policy override entry and click Properties. The AA Subscriber Policy Override (Create|Edit) form opens.

7

Configure the AA Subscriber Type parameter.

8

Perform one of the following, based on the AA Subscriber Type parameter setting:

a. If it is SAP, select a SAP subscriber in the SAP Subscriber panel.

b. If it is Spoke SDP Binding, select a spoke SDP binding subscriber in the Spoke SDP Binding Subscriber panel.

c. If it is Transit, select a transit subscriber in the Transit Subscriber panel.

9

Click on the ASO Characteristics tab.

1. Click Create or select an existing policy override ASO characteristic entry and click Properties. The AA Subscriber Policy Override ASO Characteristic (Create|Edit) form opens.
2. Select an override ASO characteristic.
3. Select an override ASO characteristic value.

10

Save your changes and close the forms.

END OF STEPS

13.9 To configure Cflowd collectors on an ISA-AA group or partition

13.9.1 Before you begin

You can configure Cflowd collectors on an ISA-AA group, and on partitions within the group. The options differ, depending on whether you configure a group or partition.

13.9.2 Steps

1

If you are performing this procedure as part of the ISA-AA group or partition configuration described in 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) , go to Step 4 .

2

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups→ISA-AA Group *n*.

3

Right-click on the ISA-AA Group *n* icon and choose Properties. The ISA-AA Group (Edit) form opens.

If you are configuring Cflowd on an ISA-AA partition, go to Step 7 .

4

Click on the CFLOWD tab and configure the required parameters on the General tab.

5

Click on the Collector tab.

1. Click Create or select an existing collector entry and click Properties. The Cflowd Collector (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form

Note:

An ISA-AA group can contain a maximum of two Cflowd collectors.

6

Go to Step 11 .

7

Click on the ISA-AA Partitions tab.

8

Select a partition in the list and click Properties. The ISA-AA Group Partition (Edit) form opens.

9

Click on the CFLOWD tab.

10

To add an AA application or application group for Cflowd performance monitoring, click on the Performance tab. Otherwise, go to Step 13 .

11

To add an AA application for Cflowd performance monitoring:

1. Click Add Applications Comprehensive, or Add Applications TCP. The Select Applications form opens.
2. Select one or more applications and click OK.

12

To add an AA application group for Cflowd performance monitoring:

1. Click Add Application Groups Comprehensive, or Add Application Groups TCP. The Select Application Groups form opens.
2. Select one or more application groups and click OK.

13

To set the Cflowd sampling administrative states, click on the States tab and configure the required parameters.

14

Configure direct Cflowd export, if required.

1. Click on the Direct Export tab.
2. Configure the Direct Export VLAN parameter.
3. Click on the Collectors tab.
4. Click Create. The CFLOWD Direct Export Collector (Create) form opens.

   Note:

   You can create only one collector.
5. Configure the parameters.
6. Click Create. The CFLOWD Direct Export Collector Address (Create) form opens.

   Note:

   You can add a maximum of two host address entries.
7. Configure the parameters.
8. Click OK. The CFLOWD Direct Export Collector Address (Create) form closes.
9. Repeat Step 14 6 to 8 to create an additional host address, if required.

   You can view the statistics for a collector host on the Statistics tab of the collector host properties form.

15

Save your changes and close the forms.

END OF STEPS

13.10 To configure an ISA-tunnel group

13.10.1 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-Tunnel Groups.

2

Right-click on the ISA-Tunnel Groups icon and choose Create ISA-Tunnel Group, or right-click an existing ISA-tunnel group object and choose Properties. The ISA-Tunnel Group (Create|Edit) form opens.

3

Configure the required parameters.

4

To associate a member-pool group to an ISA-tunnel group, configure the parameters under Member Pool Association panel.

5

Click Apply.

6

Click on the ISA Tunnel Group Members tab.

1. Click Create or select an existing tunnel group member entry and click Properties. The ISA Group Member (Create|Edit) form opens.
2. Select a daughter card.
3. Save your changes and close the form.
4. Repeat 1 to 3 to configure the backup ISA-Tunnel group member.

7

Save your changes and close the forms.

END OF STEPS

13.11 To configure an ISA-tunnel member-pool group

13.11.1 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-Tunnel Member-Pool Groups.

2

Right-click on the ISA-Tunnel Member-Pool Groups icon and choose Create ISA-Tunnel Member-Pool Group, or right-click an existing ISA-tunnel member-pool group object and choose Properties. The ISA-Tunnel Member-Pool Group (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Click on the ISA Tunnel Member-Pool Group Members tab.

1. Click Create or select an existing tunnel member-pool group member entry and click Properties. The ISA Member-Pool Group Member (Create|Edit) form opens.
2. Select a daughter card.
3. Save your changes and close the form.
4. Repeat 1 to 3 to configure the backup ISA-Tunnel member-pool group member.

6

Click on the ESA Member-Pool Group Members tab.

1. Click Create or select an existing member-pool group member entry and click Properties. The ESA Member-Pool Group Member (Create|Edit) form opens.
2. Select a VM.
3. Save your changes and close the form.
4. Repeat 1 to 3 to configure the backup ESA member-pool group member.

7

Save your changes and close the forms.

END OF STEPS

13.12 To configure an ISA-LNS group

13.12.1 Before you begin

You can configure an IES or VPRN group interface to terminate LNS PPP sessions. See Chapter 78, “IES management” for information about configuring an IES group interface. See Chapter 79, “VPRN service management” for information about configuring a VPRN group interface.

13.12.2 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-LNS Groups.

2

Right-click on the ISA-LNS Groups icon and choose Create ISA-LNS Group, or right-click an existing ISA-LNS group object and choose Properties. The ISA-LNS Group (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a port policy.

5

Click Apply.

6

Click on the ISA-LNS Group Members tab.

1. Click Create or select an existing LNS group member entry and click Properties. The ISA-LNS Group Member (Create|Edit) form opens.

   Note:

   You can configure up to six ISA-LNS group members. For each group member, you must choose an ISA Broadband Applications MDA. An ISA Broadband Applications MDA can be configured only on an IOM3-XP module on a 7750 SR, Release 20.10 and earlier.
2. Select the LNS group member MDA.
3. Repeat 1 and 2 to configure additional LNS group members, if required.

7

Save your changes and close the forms.

END OF STEPS

13.13 To configure an ISA-Video group

13.13.1 Before you begin

An ISA-Video group supports ISA or ESA-VM video group members, and not a mix of group members under the same ISA-Video group.

13.13.2 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→ISA-Video Groups.

2

Right-click on the ISA-Video Groups icon and choose Create ISA-Video Group, or right-click an existing ISA-video group object and choose Properties. The ISA-Video Group (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

To specify ISA video group members:

1. Click on the Video Group Members tab.
2. Click Create or select an existing video group member entry and click Properties. The Video Group Member (Create|Edit) form opens.
3. Select the video group member card.
4. Repeat 2 and 3 to configure additional group members, if required.

6

To specify ESA-VM video group members:

1. Click on the ESA Video Group Members tab.
2. Click Create or select an ESA existing video group member entry and click Properties. The ESA Video Group Member (Create|Edit) form opens.
3. Select the video group member card.
4. Repeat 2 and 3 to configure additional group members, if required.

7

Save your changes and close the forms.

END OF STEPS

13.14 To configure a WLAN GW group

13.14.1 Steps

1

On the Equipment tree, expand Network→*NE*→Logical Groups→ISA-WLAN GW Groups.

2

Right-click on the ISA-WLAN GW Groups icon and choose Create ISA-WLAN GW Group, or right-click an existing ISA-WLAN GW Group object and choose Properties. The ISA-WLAN GW Groups (Create|Edit) form opens.

3

Configure the required parameters.

The LSN Support parameter in the NAT panel is only configurable when the ISA WLAN GW

group is in the Shut Down state.

4

Select a port policy.

5

Select a tunnel port policy.

6

Select a NAT RADIUS accounting policy in the NAT panel, if required.

7

Select an ISA-AA group in the Distributed Subscriber Management panel, if required.

8

Click Apply. The ISA-WLAN GW Group form refreshes with additional tabs.

9

Click on the ISA-IOMs tab.

1. Click Create or select an existing IOM entry and click Properties. The ISA-WLAN GW Group IOM (Create|Edit) form opens.
2. Select the ISA-WLAN GW group member IOM card.
3. Repeat 1 and 2 to configure additional LNS group members, if required.

10

Click on the Watermarks tab to configure high/low watermark entries for the ISA-WLAN GW group.

1. Click Create or select an existing watermark entry and click Properties. The WLAN-GW Group Watermark (Create|Edit) form opens.
2. Specify the entity type to monitor.
3. Configure the High Watermark and Low Watermark percentage for the entity type.

11

Save your changes and close the forms.

END OF STEPS

13.15 To create an IGH and add members

13.15.1 Steps

1

On the Equipment tree, expand *[NE]*→Logical Groups→IGHs.

2

Right-click on the IGHs icon and choose Create IGH. The Create IGH form opens with the Define General Properties step displayed.

3

Configure the required parameters.

4

Click Next. The Configure IGH Members step displays.

5

Click Create. The Create IGH Member form opens with the Select Ports step displayed.

Alternatively, you can create IGH members anytime by right-clicking on an existing IGH object and choosing Create IGH Members from the contextual menu.

6

Choose a port from the list and click Finish.

7

Repeat Step 5 and Step 6 to add more members, if required.

8

Click Finish.

9

Click Close. The Create IGH form closes.

END OF STEPS

LAG configuration procedures

13.16 To create a LAG

13.16.1 Before you begin

This procedure applies to various variants of the 7210 SAS, 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS. See 13.20 “To configure OmniSwitch dynamic LAG members” (p. 442) for information about configuring a LAG on an OmniSwitch. See the *NSP Wavence Device Support Guide* for information about configuring a LAG on a Wavence NE.

13.16.2 Steps

1

On the equipment navigation tree, expand *NE*→Logical Groups.

2

Right-click on LAGs and choose Create LAG. The Create LAG wizard opens.

Configure general properties

3

Configure the required parameters.

If you are creating a LAG for HSMDA, the Port Type parameter must be set to HSMDA. The Encapsulation Type must be either Dot1Q or QinQ. Setting Mode to network creates a network LAG, and setting it to access creates an access LAG.

If you are creating a LAG containing ports on an XMDA card on a 7705 SAR-18, the Encapsulation Type parameter must be either Dot1 Q or Null and the Mode must be set to Access.

The Weighted parameter is configurable only if the Per-Link Hashing parameter is enabled. The Auto Rebalance parameter is only configurable only if the Weighted parameter is enabled.

Consider the following when you configure the Enable DEI parameter on a hybrid LAG.

- When the Enable DEI parameter is selected, the Enable DEI parameter for the ports in the LAG must also be selected. See 16.24 “To configure Ethernet ports” (p. 599) .
- Network policies assigned to the ports in the LAG must be configured to use DEI as the profile. The Default FC Profile parameter and the Profile parameter on the Ingress Dot1p tab must be set to the DEI option. See 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .

4

If you are configuring Ethernet ports for MWA 1+0 on a 7705 SAR-8 or 7705 SAR-18:

1. Set the Mode parameter to Network and the Encapsulation Type parameter to Dot1Q.
2. Set the MWA port Auto-negotiation parameter to Limited.

3. Create an MWA 1+0 link and add the port to the link.
4. Add the MWA 1+0 port to the LAG.

You can configure up to four MW links per LAG group, with a mix of MW and non-MW links with the same port speed settings. The LAG interface must be configured on the LAG group ID, not on the MW link. 1+1 is not supported, the redundancy is for L2 LAG only.

5

If you are configuring a LAG on a 7210 SAS NE that supports split horizon groups, you can select an SHG in the Split Horizon Group panel. See 16.46 “To create a 7210 SAS SHG” (p. 638) for information about creating a 7210 SAS SHG.

Note: A LAG cannot be added to or deleted from an SHG if it has a SAP configured on it. A LAG cannot belong to more than one SHG; you must remove all SHGs from the LAG before you can add the LAG to a different SHG.

Configure LAG parameters

6

Click Next. The Configure LAG Parameters step displays.

7

Configure the required parameters.

You must set the Port Threshold parameter to 0 when the Standby Signalling parameter is set to Power Off in Step 12 .

8

Perform one of the following:

a. If you set the Mode parameter in Step 3 to Access or Hybrid, go to Step 9 .

b. If you set the Mode parameter in Step 3 to Network, or if you did not configure the Mode parameter, go to Step 11 .

Configure access parameters

9

Click Next. The Configure Access Parameters step displays.

10

Configure the required parameters.

If you are creating a LAG for HSMDA, the QoS Adaptation parameter must be set to Link, otherwise the LAG will not be created.

Configure LACP parameters

11

Click Next. The Configure LACP form opens.

12

Configure the required parameters.

The LACP Mode, LACP Transmit Interval, Actor Administration Key, LACP System ID, and LACP System Priority parameters are configurable when the LACP Enabled parameter is enabled.

When the Standby Signalling parameter is set to Power Off, you must set the Active Sub-Group Selection Criteria parameter to Best Port.

Configure LAG members

13

Click Next. The Configure LAG Members step displays.

14

Click Create. The Only Show Compatible Ports wizard opens.

15

Configure the Show Only Compatible Ports and Class parameters.

16

Click Next. The Select Ports step displays.

17

Choose compatible ports to construct the LAG:

1. Choose one or more ports. The maximum number of ports you can add to a LAG is 64. The number or ports you can choose depends on the NE type; see the NE documentation for more information.

   Note:

   If there are no compatible ports to choose from and you want to edit some of the existing ports to be compatible, disable the Show Only Compatible Ports parameter in Step 15 .

   For ports to appear as compatible ports, ensure no SAPs or services are configured on them.

   You must use the same load balance algorithm for the ports associated with the LAG.

   The dot1x configuration modification is allowed only on the primary member port. These changes will be propagated to other member ports. Any new port added to LAG must have the same dot1x configuration as primary port.

All ports added to a LAG must have the Enable DEI parameter set to the same configuration.

Add subports of the same type, otherwise a validation error appears.

For HSMDAv2 ports in a LAG, you must set the Auto-negotiate parameter to False or Limited. See 16.24 “To configure Ethernet ports” (p. 599) for more information on configuring Ethernet ports.

If you are creating a LAG containing ports on an XMDA card in the 7705 SAR-18, then the following conditions apply to the ports:

- The Encapsulation Type must be set to either Dot1Q or Null.
- The Mode parameter must be set to Access.
- The Auto-negotiate parameter must be set to False or Limited.

2. Click Next. The Specify the Member Properties step displays.
3. Configure the parameters.
4. If required, perform Step 42 of 16.24 “To configure Ethernet ports” (p. 599) to add Egress Secondary Shapers to a port member of an HSMDA LAG.

   The add/delete/modify functions for Egress Secondary Shapers apply only to the primary port.

   There can be a total of 64 egress secondary shapers per MDA, including the default shaper.

   The Egress Secondary Shapers on the primary port of the LAG are propagated to each non-primary port member of the LAG, with a maximum 8 ports per LAG. These propagated shapers cannot be deleted.

   For adding any new port to the LAG, the user must configure each Egress Secondary Shaper as per the primary port, otherwise the addition of the port is blocked.

   When adding a port member to a LAG, the primary port and each selected port (maximum 8 ports) secondary shapers must have the same name and rate.

   Ports with Egress Secondary Shapers cannot be added to a configured LAG which do not match the Egress Secondary Shapers of the primary port.

   Shapers are not removed from the port when the port is removed from a LAG.

   The SAP reference to the Egress Secondary Shapers within a LAG is supported. Only the primary port of the SAP on the LAG is used to determine the list of Egress Secondary Shapers for the SAP.

   Note:

   When port members are removed from the LAG, care must be taken to ensure that the last port member is not removed if the LAG is referenced to a SAP. Shapers must exist on a LAG to be referenced on a SAP. See Chapter 70, “Service management and QoS” to Chapter 85, “Composite service management” for more information.

18

Click Finish.

Configure BFD parameters

19

Click Next. The Configure BFD step displays.

20

Select the IPv4 or IPv6 entry and click Properties. The BFD Configuration on LAG form appears.

21

Configure the required parameters.

22

Save your changes and close the form.

23

Click Finish.

Configure Operational Group parameters

24

Click Next. The Configure Operational Group step displays.

25

Configure the required parameters.

26

Save your changes and close the form.

27

Click Finish. Save your changes and close the wizard.

END OF STEPS

13.17 To modify a LAG

13.17.1 Before you begin

This procedure applies to 7210 SAS, 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS NEs. See 13.20 “To configure OmniSwitch dynamic LAG members” (p. 442) for information about configuring a LAG on an OmniSwitch. See the *NSP Wavence Device Support Guide* for information about configuring a LAG on a Wavence NE.

13.17.2 Steps

1

On the equipment navigation tree, expand Network→*NE*→Logical Groups→LAGs→LAG *n*.

2

Right-click on the LAG *n* object and choose Properties. The LAG (Edit) form opens.

Configure general parameters

3

Configure the required parameters.

The Weighted parameter is configurable only if the Per-Link Hashing parameter is enabled, and the Auto Rebalance parameter is only configurable if the Weighted parameter is enabled.

Consider the following when you configure the Enable DEI parameter on a hybrid LAG.

- When the Enable DEI parameter is selected, the Enable DEI parameter for the ports in the LAG must also be selected. See 16.24 “To configure Ethernet ports” (p. 599) .
- Network policies assigned to the ports in the LAG must be configured to use DEI as the profile. The Default FC Profile parameter and the Profile parameter on the Ingress Dot1p tab must be set to the DEI option. See 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) and 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .

4

If you are configuring a LAG on a 7210 SAS NE that supports split horizon groups, you can select an SHG in the Split Horizon Group panel. See 16.46 “To create a 7210 SAS SHG” (p. 638) for information about creating a 7210 SAS SHG.

Note: A LAG cannot be added to or deleted from an SHG if it has a SAP configured on it. A LAG cannot belong to more than one SHG; you must remove all SHGs from the LAG before you can add the LAG to a different SHG.

5

If you are configuring a LAG on a 7850 VSG, you can select a port profile and an operational group in the OperGroup and ServiceProfle panel.

Configure other LAG parameters

6

Click on the Link Aggregation Group tab and configure the required parameters.

You must set the Port Threshold parameter to zero when the Standby Signalling parameter is set to Power Off in Step 9 .

7

Click on the States tab and configure the Administrative State parameter.

8

If you are configuring a LAG in Access or Hybrid mode, click on the Access tab and configure the required parameters.

If you are configuring a LAG for HSMDA, the QoS Adaptation parameter must be set to Link.

9

Click on the LACP tab and configure the required parameters.

The LACP Mode, LACP Transmit Interval, and Actor Administration Key parameters are configurable when the LACP Enabled parameter is enabled.

When the Standby Signalling parameter is set to Power Off, you must set the Active Sub-Group Selection Criteria parameter to Best Port

10

Click on the LAG Members tab.

You can also add LAG members by right-clicking on an existing LAG object and choosing Create LAG Members from the contextual menu.

To add LAG members, perform Step 14 to Step 18 of 13.16 “To create a LAG” (p. 431) .

11

Perform one of the following steps:

a. If the LAG that you are configuring is in Network mode, the Network Interfaces tab is displayed. Click Create. The Create Network Interface form opens. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information on creating a network interface. Otherwise go to Step 17 .

b. If the LAG that you are configuring is in access mode or hybrid mode with dot1q or QinQ encapsulation, or in network mode with dot1q encapsulation, the MEPs tab is displayed. Go to Step 12 .

12

Add a facility MEP to the LAG:

Note: Only one facility MEP can be configured on a LAG.

1. Click on the MEPs tab and click Create. The MEP (Create) form opens.
2. Select a global MEG.
3. Configure the required parameters.

   Note:

   You cannot enable the collection of LMM statistics on a facility MEP if a service MEP is already configured to collect the LMM statistics information.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM

interval parameter is set to 10 ms or 100 ms.

To enable fault propagation on a LAG facility MEP, the Facility Fault Notify parameter must be enabled.

4. If the MD for the MEP has a Name Type of None, the Y.1731 Tests tab is configurable; click on the Y.1731 Tests tab and configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

5. Save your changes and close the form.

13

To add a tunnel facility MEP to the LAG, perform the following steps.

The LAG that you are configuring must be in Access mode or Hybrid mode and must have an encapsulation type of QinQ.

1. Click on the MEPs tab.
2. Click Create. The MEP (Create) form opens.
3. Select a global MEG.
4. Configure the required parameters.

   Note:

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

   To enable the fault propagation on a tunnel facility MEP, Facility Fault Notify must be enabled.

   To link the tunnel facility status to the SAP, Facility VLAN ID must match the outer encapsulation value of the SAP.

   To accept the fault notification from a facility MEP on the site where the SAP is located and on the SAP, Tunnel Fault Management must be set to Accept.

5. If the MD for the MEP has a Name Type of None the Y.1731 Tests tab is configurable; click on the Y.1731 Tests tab and configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

6. Save your changes and close the form.

14

Configure Ethernet-CFM redundancy on a facility MEP.

You must enable Ethernet CFM redundancy globally on a device before it can function on child objects such as LAGs. See Chapter 12, “Device object configuration” for information about device object configuration.

The tunnel MEP must be part of an MC-LAG with Ethernet CFM redundancy configured.

1. Click on the MEPs tab.
2. Choose the facility MEP for which you need to configure Ethernet-CFM redundancy and click Properties. The MEP (Edit) form opens.

3. On the Facility MEPS panel, enable the Create Redundant MEP parameter.
4. Click OK to confirm. A redundant tunnel MEP with the same ID and Operational MAC address is created on the peer LAG. The redundant MEP has the same state as the MC-LAG in which that MEP resides.

15

Configure BFD parameters:

1. Select the IPv4 or IPv6 entry and click Properties. The BFD Configuration on LAG form appears.
2. Configure the required parameters.
3. Save your changes and close the form.

16

Configure Operational Group parameters:

1. Select the Oper Group and Monitor Oper Group value.
2. Save your changes.

17

Save your changes and close the form.

END OF STEPS

13.18 To manually re-balance LAG ports

13.18.1 Purpose

Perform this procedure to force a manual subscriber load balance function on a LAG that is configured for weighted per-link hashing, or to configure LAG adaptive load balancing to resolve a traffic imbalance between LAG member ports on FP-based SR family of NEs. Adaptive load balancing is mutually exclusive with per-link-hashing.

13.18.2 Steps

1

On the equipment navigation tree, expand Network→*NE*→Logical Groups→LAGs→LAG *n*.

2

Right-click on the LAG *n* object and choose Properties. The LAG (Edit) form opens.

3

To force a manual subscriber load balance, perform the following:

1. Enable the Per-link Hashing parameter.
2. Click the Load Balance button. The LAG Load Balance form opens.

3. Configure the Class ID parameter.
4. Click OK to perform the load balance function.

4

To configure LAG adaptive load balancing, perform the following:

1. Enable the Load Balance parameter.
2. Configure the required parameters.

5

Close the form.

END OF STEPS

13.19 To configure an OmniSwitch LAG

Note: You can configure MVRP fixed ports/LAGs, 802.1 Q ports/LAGs, aggregate ports/LAGs, and VLAN Stacking Network ports/LAGs.

When a port is assigned to a LAG member, you cannot modify the port properties using the NFM-P.

13.19.1 Steps

1

On the equipment navigation tree, expand *NE*→Logical Groups.

2

Right-click on LAGs and choose Create LAG. The Create LAG wizard opens.

Configure general parameters

3

Configure the required parameters.

The Automatic VLAN Binding parameter can only be configured for network LAGs and Stacked VLANs.

Configure LAG parameters

4

Click Next. The Create LAG - Configure LAG Parameters step displays.

5

Configure the required parameters.

6

Perform one of the following:

a. If you set the Type parameter to Static, go to Step 10 .

b. If you set the Type parameter to Dynamic, go to Step 7 .

Configure LACP parameters

7

Click Next. The Create LAG - Configure LACP step displays.

8

Configure the required parameters.

9

Go to Step 15 .

See 13.20 “To configure OmniSwitch dynamic LAG members” (p. 442) for more information about adding members to a dynamic LAG.

Configure LAG members

10

Click Next. The Create LAG - Configure LAG Members step displays.

11

Click Create to add ports to the LAG. The Create LAG Member - Only show compatible ports wizard opens.

12

Configure the Show Only Compatible Ports and Class parameters.

13

Click Next. The Create LAG Member - Select Ports step displays.

14

Choose up to the maximum number of ports you require from the list of ports to create the LAG. You can select more than one port at a time from the list.

Note: If there are no compatible ports to choose from and you want to edit some of the existing ports to be compatible, then disable the Show Only Compatible Ports parameter. Click Next. Choose ports to edit from the list and click Properties.
Mobile, UNI, and NNI ports cannot be members of a LAG.
Only network ports can be members of a network LAG and only access ports can be members of an access LAG.

15

Click Finish. Save your changes and close the forms.

END OF STEPS

13.20 To configure OmniSwitch dynamic LAG members

Note: LAG members are not added during the creation of the dynamic LAG, as they are with static LAGs. When a new dynamic LAG member is created it is placed into the Unassigned Dynamic LAG Members object. The new dynamic LAG members remain under the Unassigned Dynamic LAG Members object until the system identifies a LAG with matching properties. When an unassigned LAG member joins a LAG, the member is removed from the Unassigned Dynamic LAG Members object and added to the appropriate LAG object.

When a port is assigned to a LAG member, you cannot modify the port properties using the NFM-P.

13.20.1 Steps

1

On the equipment navigation tree, expand Network→*NE*→Logical Groups.

2

Right-click on the Unassigned Dynamic LAG Members object and choose Create LAG Members. The Create LAG Member wizard opens.

3

Configure the Show Only Compatible Ports and Class parameters.

4

Click Next. The Create LAG Member - Select Ports step displays.

5

Choose up to the maximum number of ports you require from the list of ports to create the LAG. You can select more than one port at a time from the list.

Note: If there are no compatible ports to choose from and you want to edit some of the existing ports to be compatible, then disable the Show Only Compatible Ports parameter. Click Next. Choose ports to edit from the list and click Properties.
For all ports in an LAG, you must disable auto-negotiation, configure the same speed, and set the ports to full duplex. Mobile ports cannot be part of a LAG.

6

Configure the required parameters.

7

Click Finish. Save your changes and close the form.

END OF STEPS

13.21 To create a LAG link mapping profile

Note: You can only use LAG link mapping profiles on an NE that is in chassis mode D or higher.

13.21.1 Steps

1

On the equipment navigation tree, expand Network→*NE*→Logical Groups→LAGs→LAG *n*.

2

Right-click on the LAG *n* object and choose Properties. The LAG (Edit) form opens.

The following table describes the required parameter values that must be set in order to use LAG link mapping profiles.

Table 13-1 LAG link mapping prerequisites

Tab	Parameter	Required values
General	Port Type	Access or Hybrid
Access	QoS Adaptation	Link
LAG Members	Sub-Group ID	All LAG members must be in the same sub-group.

3

Click on the Link Mapping Profiles tab and click Create. The Link Mapping Profile (New Instance) form opens.

Configure general profile parameters

4

Configure the required parameters on the General tab.

Create mapping profile links

5

Click on the Mapping Profile Links tab and click Create. The Link Mapping Profile Link (New Instance) form opens.

6

Click Select and specify a port. The selected port will be used for egress traffic for interfaces assigned to this LAG link mapping profile.

7

Configure the Type parameter.

8

Save your changes and close the form.

9

Configure a secondary port by repeating Step 5 to Step 8 , if required.

Assign the link mapping profile to an interface or MSAP policy, if required

10

Assign the LAG link mapping profile to a network interface, access interface, or SAP, if required:

1. Click on the Network Interface, L2 Access Interface, L3 Access Interface, or Service Access Points tab. A list of interfaces that terminate on the LAG appears.
2. Select an interface from the list and click Properties. The Interface (Edit) form appears.
3. Click on the Port tab.
4. Select a LAG link mapping profile.
5. Save your changes and close the form.

11

Assign the LAG link mapping profile to an MSAP policy, if required:

1. From the NFM-P main menu, select Policies→Residential Subscriber. The Residential Subscriber Policies form opens.
2. Select an MSAP policy from the list and click Properties. The MSAP Policy (Edit) form opens.
3. Click on the Local Definitions tab.
4. Select the local definition for the site where the LAG link mapping profile is configured, and click Properties. The MSAP Policy (Local Policy) form opens.
5. If the Distribution Mode parameter is set to Sync With Global, click Switch Mode.
6. Set the LAG Link Map Profile parameter to the ID of the LAG link mapping profile you need to assign to the MSAP policy.
7. Save your changes and close the forms.

END OF STEPS

13.22 To view micro-BFD sessions on a LAG

13.22.1 Steps

1

On the equipment navigation tree, navigate to the LAG object that you want to configure. The path is Network→*NE*→Logical Groups→LAGs→LAG *n*.

2

Right-click on the LAG *n* icon and choose Properties. The LAG (Edit) form opens.

3

Click on the BFD tab.

4

Select the IPv4 or IPv6 entry and click Properties. The BFD Configuration on LAG form opens.

5

Click on the Micro BFD Session tab.

6

Select a BFD session and click Properties to view information about the selected session.

END OF STEPS

14 ESA object configuration

Configuring ESA objects using the NFM-P

14.1 Overview

14.1.1 ESA objects in the NFM-P navigation tree

When the NFM-P discovers applicable 7750 SR chassis, Release 20.5 R1 and later, the ESAs group object appears automatically in the navigation tree. The ESA operates as a resource server that provides packet buffering and processing, and is connected to the 7750 SR using standard SR interface ports.

The ESAs group objects are children of 7750 SR device objects. They appear above logical group objects in the navigation tree. ESA objects are children of ESAs group objects, and as such appear below the ESAs group object in the navigation tree. The ESA hosts up to four Virtual Machine (VM) instances for multiservice processing. VM objects are created in the navigation tree as children of the ESA object.

See the *NSP NFM-P Network Element Compatibility Guide* for information about which SR devices support ESAs.

14.1.2 Working with ESA objects

An ESA hosts up to four VM instances for multiservice processing.

An ESA VM may use up to four host ports. Each physical port can be used to bring up each VM within an ESA. With four host ports, four VMs can be operationally up. ESA VMs can be different types or the same type.

ESA VMs are configured as a desired type and size (number of cores and amount of memory).

Multiple ESAs may be configured per IOM and per system as needed for scale. See the node documentation for more information.

14.2 Workflow to manage ESA objects

14.2.1 Purpose

The following workflow describes the sequence of high-level tasks required to create and configure ESA objects. This workflow assumes that the physical 7750 SR devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: ESAs group objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

14.2.2 Stages

1

Create ESAs. See 14.3 “To create an ESA” (p. 449).

2

As required, create VMs. See 14.5 “To create a virtual machine on an ESA” (p. 450) .

3

As required, configure the ESAs. See 14.4 “To configure an ESA” (p. 449) .

4

As required, configure the VMs. See 14.6 “To configure a virtual machine on an ESA” (p. 451) .

5

As required, create and configure ISA-AA groups and ISA-AA partitions. See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) .

6

As required, create and configure ISA-NAT groups. See 30.3 “To configure an ISA-NAT group” (p. 1080).

7

As required, create and configure ISA Video groups. See 13.13 “To configure an ISA-Video group” (p. 427).

Procedures for ESA configuration

14.3 To create an ESA

14.3.1 Purpose

Perform this procedure to create an ESA on a supported 7750 SR system.

14.3.2 Steps

1

On the equipment tree, expand Network→*7750 SR*.

2

Right-click on the ESAs group object and choose Create ESA. The Extended Services Appliance (Create) form opens.

3

Configure the required parameters.

4

Associate a host port.

1. Click Create. The Host Port (Create) form opens.
2. Click Select. The Select Host Port form opens.
3. Select a port and click OK. The Select form closes.
4. Click OK. The Host Port (Create) form closes.

For applicable 7750 SR chassis, Release 21.2 R1 and later, you can associate two host ports.

5

Save your changes and close the form.

END OF STEPS

14.4 To configure an ESA

14.4.1 Purpose

Perform this procedure to configure an ESA.

14.4.2 Steps

1

On the equipment tree, expand Network→*7750 SR*→ESAs group.

2

Right-click on an ESA object and choose Properties. The Extended Services Appliance (Edit) form opens with the General tab displayed.

3

Configure the parameters, as required.

4

Click on the Virtual Machine tab to create, modify, or delete one or more virtual machines, as required.

5

Click on the Deployment tab to monitor the deployment status of the ESA, as required.

6

Click on the Faults tab to view alarm information, as required..

7

Save your changes and close the form.

END OF STEPS

14.5 To create a virtual machine on an ESA

14.5.1 Purpose

Perform this procedure to create an ESA VM.

14.5.2 Steps

1

On the equipment tree, expand Network→*7750 SR*→ESAs.

2

Right-click on an ESA object and choose Create VM. The Virtual Machine Extended (Create) form opens.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

14.6 To configure a virtual machine on an ESA

14.6.1 Purpose

Perform this procedure to configure an ESA VM.

14.6.2 Steps

1

On the equipment tree, expand Network→*7750 SR*→ESAs→ESA→VM.

2

Right-click on a VM object and choose Properties. The Virtual Machine (Edit) form opens with the General tab displayed.

3

Right-click on a VM and perform one of the following:

a. If the VM type has not yet been configured, choose Configure VM.

b. If the VM type has been configured, choose Properties.

The Virtual Machine (Edit) form opens with the General tab displayed.

4

Configure the parameters, as required.

5

Click on the VM Port tab to view VM port details, as required.

6

Click on the NAT VAAPs tab to view any associations to NAT groups, as required.

7

Click on the Statistics tab to view and collect statistics, as required.

8

Click on the TCA tab to apply a TCA policy to the VM, as required.

9

Click on the Deployment tab to monitor the deployment status of the VM, as required.

10

Click on the Faults tab to view alarm information, as required.

11

Save your changes and close the form.

END OF STEPS

14.7 To view virtual ports on an ESA

14.7.1 Before you begin

Perform this procedure to view ESA virtual port details.

14.7.2 Steps

1

On the equipment tree, expand Network→*7750 SR*→ESAs→ESA→VM.

2

Right-click on a VM port object and choose Properties. The VM Port (View) form opens with the General tab displayed.

3

Click on the Statistics tab to view performance statistical information about the virtual port, as required.

4

Click on the Deployment tab to monitor the deployment status of the virtual port, as required.

5

Click on the Faults tab to view alarm information, as required.

6

Save your changes and close the form.

END OF STEPS

15 Shelf and card object configuration

Configuring shelf objects using the NFM-P

15.1 Overview

15.1.1 Shelf objects in the NFM-P navigation tree

In the equipment navigation tree, shelf objects are children of device objects. They appear below logical group objects in the navigation tree. Card slot objects are children of shelf objects, and as such appear below the shelf object in the navigation tree. Daughter card slot objects are created in the navigation tree as children of the card slot object, after the card slot object is configured. APS bundle objects appear below the shelf object and Bundle objects appear below daughter card objects that are configured as channelized MDAs. See Chapter 38, “APS” for more information about APS bundles.

Shelf objects represent the hardware that is configured on a shelf. When you choose the shelf object in the navigation tree and click on Properties in the contextual menu, you can view the states and conditions of the shelf including:

- general information
- fan tray state and speed
- power supply tray statuses
- LED statuses
- card slots
- hardware environment information
- CCM properties
- PoE
- timing
- statistics
- dry contacts
- faults
- port segregation
- software control module
- software bank information
- cross connects

The Display tab displays a graphical representation of the device’s shelf and its equipment components, such as the empty card slots and the cards that are installed on the device. You can double-click on an object under this tab to open its Properties configuration form. Right-click on the object and you have full access to the contextual menus for the object and any child objects, for example the ports of a card (dynamic graphical representation only).

15.1.2 Chassis modes

The chassis mode of a device indicates the minimum IOM or IMM card type that is initialized by the device and determines the scaling numbers and features that are available to the system. See the appropriate device Release Notice for scaling information.

Note: Chassis modes are not supported on the 7950 XRS, and are not required for NEs using SR OS, Release 15.0 or later. After upgrading an NE to SR OS Release 15.0 or later, the chassis mode is displayed as D and cannot be modified.

15.1.3 Timing synchronization

The NFM-P supports the configuration of timing synchronization on supporting NEs. Configurations are typically performed on the shelf properties form. Additional configurations may be required on the NE properties form, for example, to configure specific timing protocols. See 15.6 “Workflow to manage shelf objects” (p. 457).

The NFM-P supports the following timing synchronization options. See the NE documentation for information about synchronization support for a specific NE.

- Synchronous Ethernet (SynchE), based on ITU-T G.8261, G8262, G8264, G8265, and G8275
- IEEE 1588 PTP clock
- BITS

See 15.20 “To configure timing synchronization” (p. 477) for more information about configuring timing synchronization.

Note: NTP and SNTP are also available for compatible devices but are configured separately from timing synchronization. See 15.34 “To configure NTP on supported devices” (p. 490) and 15.32 “To configure SNTP on a 7705 SAR” (p. 489).

15.1.4 IEEE 1588 PTP clock

The NFM-P supports IEEE 1588 PTP clocks and the ITU-T G.8265.1, G.8275.1, or G.8275.2 PTP profile, for packet-based timing synchronization on supporting NEs.

The IEC/IEEE 61850.9.3 PTP clock profile supports recovery of frequency as well as time/phase. PTP transport is based on Ethernet encapsulation with multicast addressing. Instead of the delay-request/response mechanism, the IEC/IEEE 61850.9.3 profile uses the peer delay mechanism.

The C37.238-2017 profile is an extension of IEC/IEEE 61850.9.3 profile with support of additional TLVs.

The IEC/IEEE 61850.9.3 and the C37.238-2017 profiles are supported on a 7705 SAR with a CSM clock ID.

The NFM-P supports IEC/IEEE 61850.9.3 and C37.238-2017 profiles on 7210 SAS-DXP 16 and 24 port variants.

Depending on the device, NEs can be configured with an ordinary master, ordinary slave, or boundary clock. The 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M/ME, 7705 SAR-W, and 7705 SAR-Wx can also be configured with a transparent clock.

The following satellites on 7450 ESS, 7750 SR and 7950 XRS nodes, Release 16.0 R5 and later support transparent clock functionality:

- 48-port 1-Gig Ethernet SFP satellite
- 64-port 10-Gig SFPP Ethernet satellite with 4 100-Gig CFP up-links
- 48-port 1-Gig Ethernet SFP satellite, LCS variant
- 64-port 10-Gig SFPP Ethernet satellite with 4 100-Gig QSFP28 up-links

For some NEs, PTP can operate in two modes. In PTP pure mode, PTP is selected as both the timing reference and frequency source. In PTP hybrid mode, PTP is selected as the timing reference, and SSU (for example, SyncE or BITS) is selected as the frequency source.

15.1.5 Synchronization management

You can use the NFM-P to manage synchronization domains and assign IP path monitors to PTP peers. See Chapter 46, “Synchronization management” for more information.

15.2 SCADA on the 7705 SAR

15.2.1 Overview

SCADA systems are used to monitor and maintain networks from remote monitoring locations. SCADA uses a controlling unit/RTU architecture with a single controlling unit that supports multiple RTUs.

The ISC supports software applications that specifically meet the requirements of TDM-based SCADA systems. The ISC is supported on the 7705 SAR-8 and the 7705 SAR-18.

The SCADA MDDB and PCM multidrop bridge applications feature similar architecture and functionality, with the main exception being that the MDDB application uses a serial RS-232 interface, while the PCM multidrop bridge application uses an E&M analog interface.

15.2.2 Multi-drop data bridge

The MDDB application provides a centralized digital bridging functionality that allows a SCADA bridge to be configured between a controlling unit and remote RTUs. The bridge allows a single data message stream to be broadcast from a controlling unit to multiple RTUs and allows a single RTU to communicate back to the controlling unit.

In a SCADA network, the 7705 SAR provides the communications infrastructure to connect the central controlling units to multiple RTUs at remote locations, where the controlling units and RTUs communicate over serial RS-232 links (synchronous or asynchronous). The 7705 SAR-8 or 7705 SAR-18 located at the controlling unit site contains the ISC, which provides the MDDB bridge functionality and acts as the MDDB controlling unit. Remote 7705 SAR nodes connected to RTUs are referred to as MDDB RTUs. See 6.8 “7705 SAR” (p. 222) for information about the ISC. See 15.36 “To configure SCADA on a 7705 SAR” (p. 495) for more information.

The remote nodes are connected to the SCADA bridge over an IP/MPLS network. The ISC supports up to 16 SCADA bridges. Each bridge supports 32 branches. Two branches (branch 1 and branch 2) are dedicated connections to the SCADA controlling units; the other 30 branches connect to the RTUs.

Note: Larger bridges can be built by cascading individual bridges internally within a single ISC and using the controlling unit output from one bridge as the RTU input to another bridge. Larger bridges can be cascaded across multiple ISCs by using an RS-232 link.

15.2.3 PCM multi-drop data bridge

The PCM multidrop bridge application provides multidrop bridging for SCADA systems that use 4-wire analog modems to connect remote RTUs to a controlling unit. Incoming analog signals from

the controlling unit are converted to PCM (Mu-Law or A-Law) for transport between a remote RTU and the controlling unit. The ISC broadcasts the controlling unit stream to all remote RTUs. Only the addressed remote unit will respond to the broadcast and the response must be transported through the bridge back to the controlling unit via an E&M interface. If the network RTUs support two SCADA systems over the same interface by separating them into high-frequency and low-frequency bands, the PCM multidrop bridge always selects the two loudest branches to be passed through the bridge for communication with the controlling unit.

15.2.4 Redundancy

Most SCADA systems have redundant controlling units where both monitor all traffic but only one controlling unit transmits at a time. The 7705 SAR SCADA system supports duplicated links - one to each controlling unit. Each bridge has two branches that connect to the controlling unit. The bridge supports two controlling unit links with only one being active. If redundancy is required, the second link can be manually configured and activated by forcing a switchover. Data from the bridge to the controlling unit is broadcast to both controlling units. The switchover allows only one controlling unit to access the bridge at a time. The 7705 SAR does not support automatic switchover capability.

15.2.5 Branch squelch

A condition may occur where a single RTU continues to send data to the controlling unit after the normal response period has expired. This condition locks up the bridge so that no other RTU can transmit data back to the controlling unit. To resolve this condition, the squelch command can be enabled on a bridge (it is disabled by default) by configuring a timeout period that, once expired, raises an alarm and triggers the squelching function. Squelching blocks the errant RTU so that other RTUs can continue to use the bridge. The squelch reset command is used to put the bridge back into a normal state. See 15.36 “To configure SCADA on a 7705 SAR” (p. 495) for more information about configuring a squelch reset.

15.3 Power management configuration

15.3.1 Overview

The NFM-P supports configuring power management zones on NEs with power supplies that support configuration, and power priority values on MDA cards installed on those nodes. Each MDA and APEQ provisioned on a node is associated with a power management zone. During startup, MDAs with higher power priority values are started first. You can view information and collect statistics about power availability and shedding using the NFM-P.

See 15.47 “To configure a power management zone” (p. 510) for information about configuring a power management zone. See 15.78 “To configure an MDA” (p. 536) for information about configuring power priority values on an MDA. See the *NSP NFM-P Statistics Management Guide* for information about gathering power statistics using the NFM-P.

Detailed information about power usage is available on the power management zone, shelf, card slot, and daughter card slot configuration forms. Use the Hardware Environment tab to view information about power usage.

15.4 Reboot hold

15.4.1 Functional description

The Reboot Hold option in a contextual menu instructs the selected standby SF or CPM module to gracefully shut down. The module remains shut down until instructed to reboot by another operation, such as the Reboot option.

15.5 Manual chassis reboot

15.5.1 Overview

There are two fields at the Shelf->General->Chassis Reboot level which indicate if and why a manual chassis reboot is required:

- “System Reboot Required”: Indicates if a configuration change has been made that requires an operator-driven reboot to fully activate.
- “Reboot Reason”: If a reboot is required, the reboot field checkbox will be checked with the respective reason for why the node requires a reboot.

To update these fields, a manual NE resync should be executed because the node does not send trap separately for each read-only property. The user can also update these fields through executing the clear command of “clear system reboot-required” on the NE; however, that NE does not send trap to NFM-P.

15.6 Workflow to manage shelf objects

15.6.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure shelf objects and shelf object children. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: Shelf objects and shelf object children can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

15.6.2 Stages

1

As required, perform the following procedures to configure shelf objects.

a. Configure the chassis mode of a device. See 15.13 “To configure the device chassis mode” (p. 471) .

b. Configure a VWM shelf for a 7210 SAS device. See 15.14 “To configure a VWM shelf for a 7210 SAS” (p. 472) .

c. Configure dry contact sensor monitoring. See 15.15 “To configure dry contact sensors” (p. 473) .

d. Configure the IMM card type on a 7210 SAS-R. See 15.16 “To configure the IMM card type on a 7210 SAS-R” (p. 474) .

e. Configure switch fabric multicast ingress replication rates. See 15.17 “To configure switch fabric multicast ingress replication rates” (p. 475) .

f. Configure IMPM overrides. See 15.18 “To configure IMPM overrides” (p. 475) .

g. Enable mixed mode on a 7450 ESS or 7750 SR device. See 15.19 “To enable mixed mode” (p. 476) .

h. Configure an IEEE 1588 PTP port on a 7705 SAR. See 15.30 “To configure an IEEE 1588 PTP port on a 7705 SAR” (p. 487) .

i. Configure system time on a 7705 SAR. See 15.31 “To configure system time on a 7705 SAR” (p. 488) .

j. Configure SNTP on a 7705 SAR. See 15.32 “To configure SNTP on a 7705 SAR” (p. 489) .

k. Configure NTP on supported device types. See 15.34 “To configure NTP on supported devices” (p. 490) .

l. Configure the IEEE 1588 PTP clock on supported devices. See 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478).

m. Configure the IEEE 1588 PTP peer on supported devices. See 15.22 “To configure the IEEE 1588 PTP peer of a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 480).

n. Configure the IEEE 1588 PTP port on supported devices. See 15.23 “To configure IEEE 1588 PTP ports on a 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS” (p. 481).

o. Manage OmniSwitch running configuration. See 15.41 “To manage an OmniSwitch running configuration” (p. 505) .

p. Configure OmniSwitch health monitoring. See 15.42 “To configure OmniSwitch health monitoring” (p. 506) .

q. Configure a CCM on a 7950 XRS-20. See 15.43 “To configure a CCM on a 7950 XRS-20” (p. 507) .

r. Configure power supply trays. See 15.46 “To configure a power supply tray” (p. 509) .

s. Configure power management zones. See 15.47 “To configure a power management zone” (p. 510) , 15.49 “To provision an APEQ” (p. 511) , and 15.50 “To configure a variable power supply APEQ” (p. 511) .

2

Assign a card type to a chassis slot. See 15.56 “To assign a card type” (p. 517) .

3

As required, configure the card and card slot objects as follows:

a. Configure an egress WRED queue control on an IOM 3 or IMM forwarding plane. See

15.62 “To configure egress WRED queue control on an XCM, IOM 3 or IMM forwarding plane” (p. 521) .

b. Configure IMPM on a 2 × XP MDA IOM 3 or IMM forwarding plane. See 15.64 “To configure IMPM on an XCM, 2 x XP MDA IOM 3, or IMM forwarding plane” (p. 522) .

c. Enable named pool mode. See 15.67 “To enable named pool mode” (p. 526) .

d. Configure OmniSwitch stacks. See 15.72 “To configure OmniSwitch stacks” (p. 530) .

e. Configure an OmniSwitch CPU temperature threshold. See 15.73 “To configure an OmniSwitch CPU temperature threshold” (p. 530) .

f. Configure a CPM on a 7950 XRS. See 15.74 “To configure a CPM” (p. 531) .

4

As required, configure daughter card objects as follows:

a. Configure an MDA. See 15.78 “To configure an MDA” (p. 536) .

b. Configure IMPM on an MDA. See 15.79 “To configure IMPM on an MDA” (p. 539) .

c. Specify an event action for a 7x50 MDA. See 15.81 “To specify an event action for a 7x50 MDA” (p. 541) .

d. View operational multicast channel properties on an MDA. See 15.82 “To view the operational multicast channel properties of an MDA” (p. 541) .

e. Configure a module card on the 7705 SAR-M/ME. See 15.83 “To configure a module card on a 7705 SAR-M/ME or 7705 SAR-H” (p. 543) .

5

As required, configure shelf objects that require daughter cards as follows:

a. Configure timing synchronization. See 15.20 “To configure timing synchronization” (p. 477) .

b. Configure an IEEE 1588 clock on a 7705 SAR. See 15.26 “To configure an IEEE 1588 PTP clock on a 7705 SAR” (p. 483) .

c. Configure SCADA on a 7705 SAR. See 15.36 “To configure SCADA on a 7705 SAR” (p. 495) .

d. Configure a 7705 SAR microwave link. See 15.38 “To configure a 7705 SAR MW link” (p. 499) .

e. Configure a 7705 SAR MW link member. See 15.39 “To configure a 7705 SAR MW link member” (p. 501) .

f. Configure auxiliary alarm definitions on a 7705 SAR. See 15.40 “To configure 7705 SAR auxiliary alarm definitions” (p. 504) .

6

As required, create or modify bundle objects as follows:

a. FR group bundle. See 15.86 “To create an FR group bundle” (p. 548) and 15.87 “To modify an FR group bundle” (p. 549) .

b. IMA group bundle. See 15.88 “To create an IMA group bundle” (p. 550) and 15.89 “To modify an IMA group bundle” (p. 551) .

c. MLPPP bundle. See 15.90 “To create an MLPPP bundle” (p. 551) and 15.91 “To modify an MLPPP bundle” (p. 553) .

d. MLPPP bundle for multiclass service transmission. See 15.92 “To configure an MLPPP bundle for multiclass service transmission” (p. 554) .

e. MLPPP bundle as a network interface on a channelized ASAP MDA. See 15.93 “To configure an MLPPP bundle as a network interface on a channelized ASAP MDA” (p. 555) .

Working with card and card slot objects

15.7 Overview

15.7.1 Functional description

When you expand a shelf object in the equipment navigation tree, the card slots of the shelf are displayed as child objects. A card slot object is shown as empty when no card is provisioned.

Card slots A and B are reserved for CPM or CSM cards. When you expand a CPM or CSM card object, the processor and flash memory objects are displayed.

Note: Some 7210 SAS and 7250 IXR devices are equipped with a USB flash memory. In the equipment tree and on Alarm Info forms, the NFM-P displays only Flash Memory. The display does not distinguish between compact flash memory and USB flash memory.

Card slot objects for the OS 6250, OS 6400, OS 6450, OS 6465, OS 6850, OS 6850E, OS 6855, OS 6860, OS 6860E, OS 6860N, and OS 6865 appear automatically when a physical device exists. This occurs because the entire device is represented by a card slot object and no daughter cards are associated with these devices.

When the NFM-P discovers a 7210 SAS or 7705 SAR-F, the integrated IOM and daughter card objects appear automatically in the navigation tree.

The 7250 IXR-R6 uses CPIOM cards that provide both CPM and IOM functionality. For information about the NFM-P equipment tree display of card slots and daughter card slots on the 7250 IXR-R6, see 6.6.1 “Overview” (p. 220).

The 7750 SR-s series of NEs allow for configuration of an IOM-s in a slot position of an XCM-s. This module is an Xiom-s. You can configure either XMA or Xiom-s on XCM-s, but not both. See 15.57 “To configure an Xiom-s card slot” (p. 518).

When the NFM-P discovers a supported PAC FP3 IMM in a card slot, the daughter cards and ports appear automatically in the navigation tree. Daughter cards and ports that are discovered in this way cannot be manually deleted.

Choose Configure Card from the contextual menu of the object and assign a supported card type for the slot. The Assigned Card Type parameter lists the card types that can be assigned to the card slots.

Some card types can be pre-provisioned in a slot before the card is installed in the chassis. A card and daughter card must be provisioned before a port can be configured.

When a card is first configured, the administrative state can be down. The resource is not operationally up until the card is equipped and the administrative state is up. A card can only be provisioned in a slot that is vacant, and no other card can be provisioned (configured) for that specific slot.

The Remove Card option in a contextual menu deletes the card from the slot when the slot and everything contained in the slot is changed to administratively down.

To reconfigure a slot position, delete the card currently in the slot and configure the new card type added to the slot. A card can only be provisioned in a slot when the card type is allowed in the slot.

Note: You can also reconfigure a slot position by changing the chassis mode. The chassis mode determines the minimum card requirements. See 15.1.1 “Shelf objects in the NFM-P navigation tree” (p. 453) for more information about chassis modes.

For example, if a 2 x 10-Gig MDA IOM 2 card is installed in the chassis that is running in chassis mode B, the card behaves as a 2 x 10-Gig MDA IOM Card, B. You can upgrade the chassis mode to C to make the IOM card behave as a 2 x 10-Gig MDA IOM 2 card.

You can configure forwarding plane properties on the properties forms for IOM 3, IMM, and XCM cards. IOM 3 and IMM cards support one forwarding plane, and XCM cards support two forwarding planes — one for each XMA card slot.

Note: You cannot configure a forwarding plane on an XCM card until you have provisioned the corresponding XMA card.

You cannot remove a provisioned XMA card when any of the following is true on the corresponding XCM forwarding plane:

- Egress WRED queue control is administratively enabled
- Access ingress queue groups are configured
- Network ingress queue groups are configured
- The Dynamic Enforcement Policer Pool Size parameter is configured with a nonzero value

15.7.2 PChip FCS Errors for cards

PChip FCS Error counts can be viewed from the IO Card tab of a Card Slot form. The Card Complex Statistics can be viewed from the Statistics tab of a Card Slot form. When frames of data are being corrupted during transmission, or duplex disparities between switch and the end device causes these errors, the NFM-P raises an alarm.

Note: The PChip alarm details, such as the alarms and statistics, cannot be differentiated whether raised against an MDA or IOM.

15.8 Card provisioning and chassis modes

15.8.1 Overview

The behavior of the installed card depends on the chassis mode that has been configured on the device. The following table describes the behavior of the provisioned (configured) card types and installed cards in each chassis mode.

Table 15-1 Behavior of installed IOM cards

Provisioned card type	Installed card type	Behavior
2 x 10-Gig MDA IOM	2 x 10-Gig MDA IOM	The IOM card appears and behaves at the scaling limits of the 2 x 10-Gig MDA IOM.
	2 x 10-Gig MDA IOM Card, B	The IOM card appears and behaves as a 2 x 10-Gig MDA IOM. The 2 x 10-Gig MDA IOM Card, B is fully backward-compatible with the 10-Gig MDA IOM.
	2 x 10-Gig MDA IOM 2	The IOM card appears and behaves as a 2 x 10-Gig MDA IOM. The 2 x 10-Gig MDA IOM 2 is fully backward-compatible with the 10-Gig MDA IOM.
	2 × XP MDA IOM 3 [1]	The IOM card appears and behaves as a 2 x 10-Gig MDA IOM. The 2 x XP MDA IOM is fully backward-compatible with the 10-Gig MDA IOM.
2 x 10-Gig MDA IOM Card, B	2 x 10-Gig MDA IOM	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM Card, B	The IOM card appears and behaves at the scaling limits of the 2 x 10-Gig MDA IOM Card, B.
	2 x 10-Gig MDA IOM 2	The IOM card appears and behaves as a 2 x 10-Gig MDA IOM Card, B. The 2 x 10-Gig MDA IOM 2 is fully backward-compatible with the 10-Gig MDA IOM Card, B.
	2 × XP MDA IOM 3 [1]	The IOM card appears and behaves as a 2 x 10-Gig MDA IOM Card, B. The 2 x XP MDA IOM 3 is fully backward-compatible with the 10-Gig MDA IOM Card, B.
2 x 10-Gig MDA IOM 2	2 x 10-Gig MDA IOM	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM Card, B	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM 2	The behavior of the card depends on the operational chassis mode: • 2 x 10-Gig MDA IOM in chassis mode A • 2 x 10-Gig MDA IOM Card, B in chassis mode B • 2 x 10-Gig MDA IOM 2 in chassis mode C and D
	2 × XP MDA IOM 3 [1]	The behavior of the card depends on the operational chassis mode: • 2 x 10-Gig MDA IOM in chassis mode A • 2 x 10-Gig MDA IOM Card, B in chassis mode B • 2 x 10-Gig MDA IOM 2 in chassis mode C and D

Table 15-1 Behavior of installed IOM cards (continued)

Provisioned card type	Installed card type	Behavior
2 × XP MDA IOM 3 card [1]	2 x 10-Gig MDA IOM	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM Card, B	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM 2	The IOM card fails because the minimum card requirements are not met.
	2 × XP MDA IOM 3 [1]	The behavior of the card depends on the operational chassis mode: • 2 x 10-Gig MDA IOM in chassis mode A • 2 x 10-Gig MDA IOM Card, B in chassis mode B • 2 x 10-Gig MDA IOM 2 in chassis mode C • 2 x XP MDA IOM 3 in chassis mode D. With mixed mode enabled for 7750 SR, only the IPv6 features of the IOM3 card are enabled, but no other IOM3 features are available.
2 × XP MDA IOM 3, B card [1]	2 x 10-Gig MDA IOM	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM Card, B	The IOM card fails because the minimum card requirements are not met.
	2 x 10-Gig MDA IOM 2	The IOM card fails because the minimum card requirements are not met.
	2 × XP MDA IOM 3 [1]	The behavior of the card depends on the operational chassis mode: • 2 x 10-Gig MDA IOM in chassis mode A • 2 x 10-Gig MDA IOM Card, B in chassis mode B • 2 x 10-Gig MDA IOM 2 in chassis mode C • 2 x XP MDA IOM 3 in chassis mode D. With mixed mode enabled for 7750 SR, only the IPv6 features of the IOM3 card are enabled, but no other IOM3 features are available.

Notes:

1. Also applies to any IMM

Working with daughter card objects

15.9 Overview

15.9.1 Daughter card objects

The NFM-P navigation tree displays daughter cards, or MDAs, as child objects of card objects.

An NFM-P license has a specific equipment management capacity. MDAs fall into two licensing categories: premium and standard. Premium MDAs consume more license capacity than standard MDAs. In general, non-high-performance MDA cards are considered standard. All IMM, HSMDA, XMDA, and XP MDA cards are premium.

Each IMM card is associated with a platform, and counts towards the licensed IMMs for that platform regardless of where it is installed. For example, a 7750 SR IMM installed on a 7450 ESS is counted as a premium 7750 SR MDA for the purposes of licensing. Enabling or disabling mixed mode does not change the associated platform type of a card.

FP4 cards on 7x50, Release 16.0 R1 and later, support license levels under card and MDA. The Licensing Details parameters appear on the Card Slot and Daughter Card Slot forms. You can set the Licensed Assigned Level that a card or daughter card slot can accept.

15.9.2 DCO daughter cards

CFP2-DCO daughter cards support DWDM channels, coherent optics, and optical transport channel unit (OTU). The 400 Gig QSFP-DD daughter cards and x2-1000g-wdm XMA support DWDM and coherent optics only. And other daughter cards support DWDM only. See the hardware documentation for more information about daughter card and XMA optical support.

On the NFM-P, DWDM channel and coherent optics are supported under connector ports and OTU is supported under breakout ports. See 16.23 “To configure connector ports and breakout ports” (p. 597) for more information about enabling coherent optics.

15.9.3 IMM daughter cards

The IMM card integrates IOM 3 and high bandwidth MDA functionality on a single card that fits into existing IOM slots. When you configure the IMM the integrated MDAs are automatically configured. The NFM-P equipment tree displays the IMM with two daughter cards, each having half of the total number of ports supported by the IMM.

Each daughter card object contains a number of ports that are specific to the type of service required. The port objects are created automatically under the daughter card but they must be configured based on the function served by the port; for example, as an access interface for a VPRN service.

You can associate policies to daughter cards. Network buffer policies are used to create and edit QoS buffer pool resources on egress network ports, channels, and ingress ports. Ingress and egress network ports and channels have a dedicated buffer pool for egress queuing. The ingress and egress network traffic is handled by a buffer pool at the ingress and a buffer pool at the egress.

You can also configure multicast path management on an IMM daughter card. See Chapter 52, “Multicast policies” for more information.

15.9.4 MDA modes for the 7705 SAR

The NFM-P allows you to configure the MDA mode of some MDA types on the 7705 SAR. The MDA mode on a channelized ASAP determines the capabilities that are available to the MDA. The MDA mode on an XMDA determines the number of ports that are available to the MDA. See the MDA Mode parameter description for more information.

Working with bundle objects

15.10 Overview

15.10.1 Bundle objects

Bundle objects can be configured as multilink PPP, IMA group, or FR group bundles. Multilink PPP, IMA group, and FR group bundles appear as Bundle objects under the daughter card object if the daughter card has been configured as a channelized MDA.

APS Bundle objects appear in the navigation tree as children of the shelf object. APS Bundle objects are configured to protect multilink PPP and IMA bundles on channelized ASAP MDAs. See Chapter 38, “APS” for more information on APS and APS bundles.

15.10.2 Multilink PPP bundles

You can create MLPPP bundles to:

- bundle DS0 channels together to be used as a SAP
- provide a mechanism to distribute data across multiple links to achieve higher bandwidth

The following general rules apply to multilink bundles.

- Table 15-2, “Multilink bundle configuration maximum” (p. 467) lists the maximum number of multilink bundles that can be created on each MDA.
- DS0 channels can be aggregated on a single MDA only.
- Up to 8 members can be added to a bundle.
- The Encap Type parameter of a member cannot be set to FR.
- A DS0 channel can host only one SAP.
- If a channel is being used as a SAP, it cannot be added to a bundle.

Table 15-2 Multilink bundle configuration maximum

MDA	Bundle maximum
4 x Channelized OC3/OC12 ASAP SFP MDA	32
All MDAs other than Channelized ASAP MDAs	56
• 1 x OC12 Deep Channel • 4 x OC3 Deep Channel • 12 x DS3/E3 Deep Channel • 4 x DS3/E3 Deep Channel	56
4 x Channelized DS3/E3 ASAP MDA	112
• 4 x Channelized OC3 ASAP • 1 x Channelized OC12 ASAP • 12 x Channelized DS3/E3 ASAP	256

See 15.90 “To create an MLPPP bundle” (p. 551) to create an MLPPP bundle.

15.10.3 IMA group bundles

IMA group bundles combine ATM-encapsulated DS0 channel groups into a single ATM interface. The following general rules apply to IMA group bundles.

- IMA group bundles can only be created on:
  - channelized ASAP MDAs with SDH or SONET framing the 7750 SR
  - channelized DS1/E1 ASAP daughter cards on the 7705 SAR
- The encapsulation type of the DS0 group channel must be ATM.
- The Clock Source parameter of the DS1 channel must be set to Node-Timed.
- Table 15-3, “IMA group bundle configuration maximums” (p. 467) lists the maximum number of IMA group bundles that can be created on each MDA.
- IMA group members must be on the same MDA.

Table 15-3 IMA group bundle configuration maximums

Device	MDA	Bundle maximum
7750 SR	All MDAs other than channelized ASAP MDAs	56
7750 SR	4 x Channelized DS3/E3 ASAP MDA	112
7750 SR	• 4 x Channelized OC3 ASAP • 1 x Channelized OC12 ASAP • 12 x Channelized DS3/E3 ASAP	256
7705 SAR	16 x Channelized DS1/E1 ASAP	8
7705 SAR	2 x Channelized OC3/STM1 ASAP	16

See 15.88 “To create an IMA group bundle” (p. 550) to create an IMA group bundle.

15.10.4 FR group bundles

FR group bundles are used to fragment lower priority DLCI frames to minimize jitter and delay on higher priority DLCI frames. The following general rules apply to FR group bundles.

- FR group bundles can only be created on channelized ASAP MDAs.
- FRF12 cannot be enabled on a SAP that is using an FR group bundle.
- DS0 channel groups must have all their timeslots selected and have their encapsulation type set to FR to be configured as bundle members.

Note: You can configure a maximum of 128 FR group bundles on an MDA.

See 15.86 “To create an FR group bundle” (p. 548) to create an FR group bundle.

Working with extension shelf objects

15.11 Extension shelves

15.11.1 Overview

An extension shelf is a virtual object that is modeled in the NFM-P as a shelf object. It can represent a variety of equipment connected to the NE by cables connected to uplink ports. This section describes the types of extension shelves supported by the NFM-P.

15.11.2 Ethernet satellite shelf

An Ethernet satellite shelf is a breakout box attached to a router that provides equipment support for additional gigabit Ethernet ports. These breakout boxes are modeled as virtual shelves on the host NE in the NFM-P. See 15.94 “To create a satellite shelf” (p. 557) for information about creating a satellite shelf on an NE. Each breakout box is connected to an NE through uplink ports; see 15.95 “To configure satellite shelf uplink port topology” (p. 557) for information about configuring the uplink.

You can specify the satellite file transfer protocol to be used for the boot process when transferring boot images and configuration files from the 7x50 host to the Ethernet satellite. See 12.42 “To configure satellite file transfer” (p. 375) for more information.

15.11.3 TDM satellite shelf

A TDM satellite shelf is a separate TDM chassis attached to a router that provides additional TDM ports. See 15.94 “To create a satellite shelf” (p. 557) for information about creating a satellite shelf on an NE. Each chassis is connected to an NE through uplink ports; see 15.95 “To configure satellite shelf uplink port topology” (p. 557) for information about configuring the uplink.

Software upgrade

You can upgrade the software on an optical extension shelf using the NFM-P. To perform an upgrade, you must associate the shelf with a software repository, then perform a reboot upgrade. See 15.101 “To perform a software upgrade on an extension shelf” (p. 562) for information about upgrading system software on a TDM satellite shelf.

Procedures for shelf object configuration

15.12 To configure an 1830 VWM shelf

15.12.1 Before you begin

The 1830 VWM manages passive devices SFD-96, SFD-44, BMU-P, and AMP as shelves. The following procedure may be used to configure passive devices as well as shelves.

15.12.2 Steps

1

On the equipment tree, expand Network→1830–VWM-OSU NE.

2

Right-click on the 1830–VWM-OSU NE object and choose Configure Shelf. The Shelf (Create) form opens.

3

Configure the parameters in the Shelf Details panel.

The following shelf types and passive devices can be configured:

1830 VWM-ITP4LCA	1830 VWM-PMU9UC	1830 VWM-SMMAI
1830 VWM-ITP4LCB	1830 VWM-PMU9LC	1830 VWM-SMMAO
1830 VWM-ITP4UCA	1830 VWM-PMU9UCM	1830 VWM-TLU9
1830 VWM-ITP4UCB	1830 VWM-PMU9LCM	1830 VWM-TLU9M
1830 VWM-OPS	1830 VWM-SFD44	1830 VWM-SFD44B
1830 VWM-SFD96	1830 VWM-BMUP	1830 VWM-PMUD21A
1830 VWM-PMUD21B	1830 VWM-PMUD21C	1830 VWM-PMUD21D
1830 VWM-SPLIT2	1830 VWM-OSCCPLR	1830 VWM-OAUNIDIR
1830 VWM-SFD10 and its variants A, B, C, and D	1830 VWM-SFD2 and its variants A, B, C, D, E, F, G , H, I, L, M, N, O, P, Q, and R	1830 VWM-SFD4 and its variants A, B, C, D, E, F, G, and H
1830 VWM-SFD8 and it variants A, B, C, and D	1830 VWM-SFC1 and its variants A, B, C, D, E, F, G, and H	1830 VWM-SFC2 and its variants AB, CD, EF and GH
1830 VWM-SFC4AD and SFC4EH	1830 VWM-SFC8U and SFC8L	1830 VWM-SARO2 and its variants A, B, C, D, E, F, G, and H

1830 VWM-SARO4 and its variants A, B, C, and D	1830 VWM-SARO8A and SARO8B	1830 VWM-LR4BOC
1830 VWM-TLU200	-	-

4

Save your changes and close the form.

END OF STEPS

15.13 To configure the device chassis mode

15.13.1 Before you begin

The device chassis mode must be set to indicate the minimum IOM or IMM card type that is initialized by the NE. This also determines other features available to the system. The behavior of the installed card depends on the chassis mode configured on the device. See 15.1.1 “Shelf objects in the NFM-P navigation tree” (p. 453) and “Working with card and card slot objects” (p. 461) for more information about chassis modes and card types.

15.13.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Configure the parameters:

- Administrative Mode

  If you downgrade the administrative chassis mode, the device must be rebooted to change the operational chassis mode.

- Force Mode

  Forcing a chassis mode change does not change the operational chassis mode unless there is a compatible card type equipped on the device. See 15.1.1 “Shelf objects in the NFM-P navigation tree” (p. 453) for more information about the minimum card type that must be installed for each chassis mode.

4

Save your changes and close the form.

END OF STEPS

15.14 To configure a VWM shelf for a 7210 SAS

15.14.1 Before you begin

The 1830 VWM is an add-on optical module that provides WDM capability to 7210 SAS devices. Support for 1830 VWM shelves varies depending on the 7210 SAS chassis type. See the NE documentation for more information.

15.14.2 Steps

1

On the equipment tree, expand Network→*NE*→VWM Shelves.

2

Right-click on VWM Shelves and choose Configure VWM Shelf. The VWM Shelf (Create) form opens.

3

Configure the required parameters.

Configure the Shelf ID parameter to match the hardware shelf ID value that is preset on the rotary dial of the 1830 VWM module. If the Shelf ID parameter does not match the hardware shelf ID, the correct manufacturer details will not appear on the VWM Shelf form.

Some chassis types support a value of zero for the Shelf ID parameter. When a value of zero is configured, the VWM Type parameter is set to DWDM-A (DWDM-active) by default and is not configurable.

4

On the equipment tree, expand the VWM Shelves object to the VWM Card Slot level.

5

Right-click on a VWM Card Slot object and choose Configure Card. The Configure VWM Card Slot form opens.

6

Configure the Assigned VWM Card Type parameter.

The Assigned VWM Card Type parameter provides a way to check if the equipped card type is acceptable. An alarm is raised when the equipped VWM card type in the module does not match the Assigned VWM Card Type.

7

Save your change and close the form.

8

Repeat Step 5 to Step 7 for each additional card slot that you need to configure.

9

You can view information about VWM controllers, such as Status and manufacturer details, by opening the VWM controller properties form. To open the form, right click on a VWM controller object in the navigation tree and choose Properties.

The Status parameter indicates whether a controller is operational or standing by. The Status parameter is available only for DWDM and DWDM-A shelf types.

10

Save your changes and close the form.

END OF STEPS

15.15 To configure dry contact sensors

15.15.1 Before you begin

You can enable dry contact sensors on the console port or alarm interface port of supported NEs. When an external alarm condition is present, the NE sends an alarm with a user-configured message to the NFM-P. See the NE documentation for more information about dry contact sensors.

15.15.2 Steps

1

If you are configuring a 7210 SAS-D or 7210 SAS-K, perform the following steps. Otherwise, go to Step 2 .

1. On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.
2. Enable the Console Alarm Input parameter.

   Note:

   The console port is not available for other purposes when the Console Alarm Input parameter is enabled.
3. Save your change and close the form.

2

On the equipment tree, expand Network→*NE*→Shelf.

3

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

4

Click on the External Alarms tab. A list of dry contacts is displayed.

5

Choose one of the dry contacts and click Properties. The Dry Contact (Edit) form opens.

6

Configure the required parameters and click OK. The Dry Contact (Edit) form closes.

7

Repeat Step 5 and Step 6 as required to configure additional dry contacts.

Note: You can configure only one dry contact for the 7210 SAS-D and 7210 SAS-K.

8

Save your changes and close the form.

END OF STEPS

15.16 To configure the IMM card type on a 7210 SAS-R

15.16.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Configure the Allow IMM Family Type parameter.

The Allow IMM Family Type parameter requires an NE reboot to take effect. The Operational IMM Family Type parameter displays the card type currently in effect for the shelf.

When None is enabled for the Allow IMM Family Type parameter, the default IMM family type for the NE is applied. The default for the 7210 SAS-R6 is KT1. For the 7210 SAS-R12, the default is KT2.

4

Save your changes and close the form.

END OF STEPS

15.17 To configure switch fabric multicast ingress replication rates

15.17.1 Purpose

Perform this procedure to configure the total multicast replication rates for primary and secondary ingress paths for each switch fabric multicast plane.

15.17.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the Multicast tab.

4

Configure the required parameters.

5

Save your changes and close the form.

END OF STEPS

15.18 To configure IMPM overrides

15.18.1 Before you begin

See 15.64 “To configure IMPM on an XCM, 2 x XP MDA IOM 3, or IMM forwarding plane” (p. 522) to configure IMPM on an IOM or XCM. See Chapter 52, “Multicast policies” for more information on multicast path management.

15.18.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the Multicast tab.

4

Configure the required parameters in the MMRP Ingress Multicast Path Management Override panel.

Note: The Operational Mode parameter is not configurable. If the Administrative Mode parameter is changed, the value of the Operational Mode parameter changes to indicate the same as the Administrative Mode, after the system is rebooted.

5

Save your changes and close the form.

END OF STEPS

15.19 To enable mixed mode

15.19.1 Before you begin

Enabling mixed mode only applies to 7750 SR devices.

Mixed-mode is the default behavior for all 7450 ESS chassis, which allows the 7450 ESS chassis to support all 7750 SR functionality.

7750 SR-7 and 7750 SR-12 devices also support mixed mode, which allows you to enable IPv6 in chassis mode A or B, without having to upgrade an entire chassis. Slots with network ports, and ports that require IPv6 support, can then be selectively upgraded for IOM3-XPs (SROS Release 20.10 and earlier) or IMMs.

Note: You cannot enable mixed mode on a 7750 SR-7 or 7750 SR-12 that has a provisioned ISA TMS MDA.

15.19.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Enable the Mixed Mode State on Chassis Enabled parameter.

4

Save your change and close the form.

END OF STEPS

15.20 To configure timing synchronization

15.20.1 Purpose

The following procedure describes the sequence of tasks required to configure timing synchronization. This procedure assumes that the shelf and its daughter cards have been configured.

You must enable the Synchronous Ethernet parameter on the daughter card in order to configure the Timing Reference One and Timing Reference Two panels for SyncE. See 15.78 “To configure an MDA” (p. 536) .

You must configure an IEEE 1588 PTP clock on the device when you configure PTP as a timing reference; see 15.26 “To configure an IEEE 1588 PTP clock on a 7705 SAR” (p. 483) or 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478).

15.20.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the Timing tab and configure the required parameters.

When you select a port as a timing reference for SyncE, the Port Clock Mode parameter for the port must be set to Manual Slave. See 16.24 “To configure Ethernet ports” (p. 599) . Other restrictions may also apply for SyncE port selection, depending on the device. See the NE documentation for more information.

You can enable the Administrative State parameter in the PTP panel only when the Frequency Source and Oper Frequency Source parameters are both set to Ptp on the IEEE 1588 PTP Clock form; see 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478) .

For the 7210 SAS-D ETR, 7210 SAS-M, and 7210 SAS-X, PTP and SyncE are mutually exclusive. Consider the following when configuring timing synchronization on these 7210 SAS NEs:

- You can only choose PTP as an option for First Timing Reference Input when the Administrative State parameter in the PTP panel is set to Up.
- You cannot configure the Second Timing Reference Input parameter when the First Timing Reference Input parameter is set to PTP.
- You can only set the Administrative State parameter in the PTP panel to Up when the Administrative State parameter in both the Timing Reference One and Timing Reference Two panels is set to Down.

- You can only set the Administrative State parameter in the Timing Reference One and Timing Reference Two panels to Up when the Administrative State parameter in the PTP panel is set to Down.

Note: The displayed parameters vary depending on the NE type, release, and the settings of other parameters.

4

Save your changes and close the form.

END OF STEPS

15.21 To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR

15.21.1 Purpose

Perform this procedure to configure the IEEE 1588 PTP clock on a device that is configured as an IEEE 1588 PTP client. See 15.20 “To configure timing synchronization” (p. 477) for information about how to configure a device as a PTP client.

Depending on the PTP profile you are configuring, you may need to configure PTP peers or PTP ports; see 15.22 “To configure the IEEE 1588 PTP peer of a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 480) and 15.23 “To configure IEEE 1588 PTP ports on a 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS” (p. 481).

15.21.2 Steps

1

In the navigation tree equipment view, right-click on the required device and choose Properties. The Network Element (Edit) form opens.

2

Click on the IEEE PTP Clock tab.

3

Select an IEEE 1588 PTP Clock from the list and click Properties. The IEEE 1588 PTP Clock (Edit) form opens.

4

Configure the required parameters.

At least one IEEE PTP peer or port must be created before you can set the Admin State parameter to Enabled. See 15.22 “To configure the IEEE 1588 PTP peer of a 7210 SAS, 7250

IXR, 7450 ESS, or 7750 SR" (p. 481) and 15.23 "To configure IEEE 1588 PTP ports on a 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS" (p. 481).

When the Admin State parameter is set to Enabled, you cannot change the following parameters:

- Clock Type
- PTP Profile
- Domain
- Clock Announce Rx Timeout
- Clock Announce Interval
- Log sync interval

The PTSF Unusable Admin State parameter is only configurable when the PTP Profile is set to ITU-T G.8275.1.

For the 7210 SAS-D ETR, 7210 SAS-K, 7210 SAS-M, 7210 SAS-R, 7210 SAS-S/Sx VC, 7210 SAS-T, 7210 SAS-T ETR, and 7210 SAS-X: when the Clock Type parameter is set to Boundary, you must set the PTP Profile parameter to IEEE-1588-2008.

If you are configuring an ITU-T G.8275.1 profile on a supporting 7210 SAS NE, you must choose Boundary Clock for the Clock Type parameter and ITU-T G.8275.1 for the PTP Profile parameter.

If you are configuring an ITU-T G.8275.1 profile on a 7250 IXR, you must choose Boundary or Slave for the Clock Type parameter and ITU-T G.8275.1 for the PTP Profile parameter.

The ITU-T G.8265.1 profile only supports master slave clocks and PTP peers. For the 7250 IXR-e big/small, the ITU-T G.8265.1 profile only supports ordinary slave clocks and PTP peers.

Changes to the PTP Profile and Frequency Source parameters may require an NE reboot to take effect; see the NE documentation for more information. The currently effective frequency source is displayed as the Oper Frequency Source.

To make SSU the effective Oper Frequency Source, the timing reference must be qualified and the Status must be set to Master Locked on the Timing tab of the shelf properties form; see 15.20 "To configure timing synchronization" (p. 477). Also, the Admin State parameter on the General tab of the ptp clock tab must be set to enabled.

Note: The displayed parameters vary depending on the NE type, release, and the settings of other parameters.

5

Save and close the forms.

END OF STEPS

15.22 To configure the IEEE 1588 PTP peer of a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR

15.22.1 Purpose

Perform this procedure to configure IEEE 1588 PTP peers for a device that has the IEEE 1588 PTP clock configured. See 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478) for information about how to configure the IEEE 1588 PTP clock on a device.

For 7250 IXR-e CPM GNSS with the ITU-T G.8265.1 profile, PTP peer creation is only supported for the ordinary slave clock profile.

15.22.2 Steps

1

In the navigation tree equipment view, right-click on the required device and choose Properties. The Network Element (Edit) form opens.

2

Click on the IEEE PTP Clock tab.

3

Choose an entry from the list and click Properties. The IEEE 1588 PTP Clock (Edit) form opens.

4

Click on the IEEE PTP Peers tab.

5

Click Create, or choose an IEEE 1588 PTP peer from the list and click Properties. The IEEE 1588 PTP Peer (Create|Edit) form opens.

The number of peers that you can create varies depending on the NE. When the limit is reached, the Create button is not available.

Alternatively, you can create IEEE PTP peers by right-clicking on the Clock object in the navigation tree and choosing Create IEEE PTP Peer.

6

Configure the required parameters.

7

Save and close the forms.

END OF STEPS

15.23 To configure IEEE 1588 PTP ports on a 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS

15.23.1 Purpose

Perform this procedure to configure IEEE 1588 PTP ports for a supporting device that has the IEEE 1588 PTP clock and the ITU-T G.8275.1, ITU-T G.8275.2, or IEEE1588-2008 profile configured. See 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478) for information about how to configure the IEEE 1588 PTP clock on a device.

You can configure multiple IEEE PTP ports on the same device.

15.23.2 Steps

1

In the navigation tree equipment view, right-click on the required device and choose Properties. The Network Element (Edit) form opens.

2

Click on the IEEE PTP Clock tab.

3

Choose an entry from the list and click Properties. The IEEE 1588 PTP Clock (Edit) form opens.

4

Click on the IEEE PTP Port tab.

5

Click Create, or choose an IEEE 1588 PTP port from the list and click Properties. The IEEE 1588 PTP Port (Create|Edit) form opens.

The number of ports that you can create varies depending on the NE. When the limit is reached, the Create button is not available.

Alternatively, you can create IEEE PTP ports by right-clicking on the Clock object in the navigation tree and choosing Create IEEE PTP Port.

6

Configure the required parameters.

When you configure a 7210 SAS-R NE as a Boundary clock, you must disable the Master Only parameter on at least one of the configured IEEE PTP ports.

7

If an NE with embedded GNSS receiver is configured as a PTP boundary with G8275.2 profile, the GNSS input is modeled as a virtual PTP port and the GNSS Port tab appears. The GNSS virtual ports are A/gnss or B/gnss.

Choose a virtual PTP port from the list and click Properties to display the IEEE 1588 GNSS Port (View) form.

8

Save and close the forms.

END OF STEPS

15.24 To configure alternate profiles under IEEE PTP Clock on a 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS

15.24.1 Before you begin

This procedure is not supported on 7250 IXR-s, 7250 IXR-e, 7250 IXR-e2, 7250 IXR-e2c, and 7250 IXR-ec.

An alternate profile can only be configured if the primary profile is either G.8275.1 or G.8275.2, and the clock type is either Master or Boundary. The initially-defined profile functions as a primary profile and up to six alternate profiles can be configured and assigned to desired PTP ports. If the PTP port is created without choosing any alternate profiles, then it will use the primary profile.

15.24.2 Steps

1

In the navigation tree equipment view, right-click on the required device and choose Properties. The Network Element (Edit) form opens.

2

Click on the IEEE PTP Clock tab.

3

Select an IEEE 1588 PTP Clock from the list and click Properties. The IEEE 1588 PTP Clock (Edit) form opens.

4

Click Create in the PTP Clock Alternate Profiles panel.

The PTP Clock Alternate Profile Inter Working (Create) form opens.

5

Configure the required parameters.

6

Save and close the forms.

END OF STEPS

15.25 To associate an alternate profile to an IEEE PTP Port on a 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS

15.25.1 Before you begin

This procedure is not supported on 7250 IXR-s, 7250 IXR-e, 7250 IXR-e2, 7250 IXR-e2c, and 7250 IXR-ec.

15.25.2 Steps

1

In the navigation tree equipment view, right-click on the required device and choose Properties. The Network Element (Edit) form opens.

2

Click on the IEEE PTP Clock tab.

3

Choose an entry from the list and click Properties. The IEEE 1588 PTP Clock (Edit) form opens.

4

Click on the IEEE PTP Port tab.

5

Click Create or choose a port entry and click Properties. The IEEE 1588 PTP Port (Create) form opens.

6

On the Alternate Profile panel, configure an alternate profile to be associated with the port.

7

Save and close the forms.

END OF STEPS

15.26 To configure an IEEE 1588 PTP clock on a 7705 SAR

15.26.1 Before you begin

You can only configure the IEEE 1588 PTP clock on a 7705 SAR that is configured as an IEEE 1588 PTP client. See 15.20 “To configure timing synchronization” (p. 477) for information about how to configure the NE as a PTP client.

15.26.2 Steps

1

On the equipment tree, expand Network→*NE*.

2

Right-click on the NE and choose Properties. The Network Element (Edit) form opens.

3

Click on the IEEE PTP Clock tab and click Create. The IEEE 1588 PTP Clock (Create) form opens.

4

Configure the required parameters.

If you set the PTP Profile to ITU-T G.8275.1 or ITU-T G.8275.2, the Clock Local Priority parameter is configurable.

You can configure the IEC/IEEE 61850.9.3 and the C37.238-2017 PTP clock only on a 7705 SAR with a CSM clock ID. Master clock is only supported for 7705 SAR devices with integrated GNSS. When you configure a 7705 SAR with the IEC/IEEE 61850.9.3 or the C37.238-2017 profile, the default settings are automatically applied.

5

Select a clock MDA for the PTP clock.

6

Select a source interface for the PTP clock.

7

If you are configuring a 7705 SAR-8 or 7705 SAR-18, you can view the properties of the local clocks, active or standby, in the Local Clock-Active/Standby panel.

8

If you set the Clock Type parameter to Boundary or Ordinary, Slave, you can view the Parent clock information in the Parent Clock panel.

9

Click Apply to save your changes.

10

To turn up PTP ports on the IEEE 1588 PTP clock:

1. Click on the IEEE PTP Port tab, and click Create. The IEEE 1588 PTP clock (Create) form opens.
2. Select a port and configure the required parameters.

The Cfg Sync Rate parameter is not supported if the Clock Type parameter was configured as Ordinary, Slave in Step 4 .

You can only configure the Local Priority and Master Only parameters if the PTP Profile parameter was configured as ITU-T G.8275.1 or ITU-T G.8275.2 in Step 4 .

3. Save your changes.

11

Save your changes and close the forms.

END OF STEPS

15.27 To configure an IEC/IEEE 61850-9-3 PTP clock on a 7705 SAR

15.27.1 Before you begin

You can only configure the IEC/IEEE 61850-9-3 PTP clock on a 7705 SAR with a CSM clock ID. Master clock is only supported for 7705 SAR devices with integrated GNSS. When you configure a 7705 SAR with the IEC/IEEE 61850-9-3 profile, the default settings are automatically applied.

15.27.2 Steps

1

On the equipment tree, expand Network→*NE*→ Shelf→Synchronization→Clock 32.

2

On the General tab, configure the PTP Profile parameter as IEC-61850.9.3.

3

Configure the Clock Type parameter, as required.

Note: Master clock is only supported for 7705 SAR devices with integrated GNSS.

4

Verify that the default values are correct for the following parameters: Domain, Network Type, Clock ID, Clock Announce Rx Timeout, and Clock Announce Interval.

5

To turn up PTP ports on the IEC-61850.9.3 PTP clock:

1. Click on the IEEE PTP Port tab. The IEEE 1588 PTP clock form opens.
2. Select a port and configure the required parameters.
3. Verify that the default values for the Cfg Sync Rate and Cfg Delay Req Rate parameters are configured as 1 pkt/sec.
4. Save your changes.

6

Save your changes and close the forms.

You can view the statistics for the PTP clock on the Statistics tab of the PTP clock form.

END OF STEPS

15.28 To configure alternate profiles under IEEE PTP Clock on a 7705 SAR

15.28.1 Before you begin

The 7705 SAR supports a number of PTP profiles. For networks deployed using specific profile, support to inter-work with other profiles is provided. Initially defined profile functions as a primary profile and alternate profiles can be configured to be assigned to desired PTP ports.

Note: Alternate profile configuration is supported only on a Master CSM clock and with primary clock profile type G.8275.1.

A maximum of 2 alternate profiles can be configured.

15.28.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Synchronization.

2

Right-click on the PTP Clock 32 and choose Properties. The ptp (Edit) form opens.

3

On the CSM Clock Alternate Profiles panel, click Create or choose a profile entry and click Properties. The PTP CSM clock Alternate Profile (Create|Edit) form opens.

4

Configure the required parameters on the General tab.

5

Save your changes and close the forms.

END OF STEPS

15.29 To associate an alternate profile to an IEEE PTP Port on a 7705 SAR

15.29.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Synchronization.

2

Right-click on the PTP Clock 32 and choose Properties. The ptp (Edit) form opens.

3

Click IEEE PTP Port tab.

4

Click Create or choose a port entry and click Properties. The IEEE 1588 PTP Port (Create|Edit) form opens.

5

Configure the required parameters.

6

On the Alternate Profile panel, configure an alternate profile to be associated with the port.

7

Save your changes and close the forms.

END OF STEPS

15.30 To configure an IEEE 1588 PTP port on a 7705 SAR

15.30.1 Before you begin

You can configure the IEEE 1588 PTP port on a 7705 SAR that is configured as an IEEE 1588 PTP client. See 15.20 “To configure timing synchronization” (p. 477) for information about how to configure a 7705 SAR as a PTP client.

15.30.2 Steps

1

Expand Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the Timing tab and click on PTP Port Properties. The PTP Port (Edit) form opens.

4

Configure the required parameters.

Note: The Master 1 and 2 clock recovery statistics are available on the Recovered Clock History Master tabs.

5

Save your changes and close the forms.

END OF STEPS

15.31 To configure system time on a 7705 SAR

15.31.1 Before you begin

You can configure system time updates, based on priority values. There are five system time reference types: PTP, GPS, GNSS, NTP, SNTP, and the default value, Holdover. PTP is supported on the 7705 SAR-8, 7705 SAR-18, 7705 SAR-A, 7705 SAR-A T1/E1, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, 7705 SAR-M ASAP, 7705 SAR-W, and 7705 SAR-Wx. GPS is supported on the 7705 SAR-H and 7705 SAR-Wx GPS variants. GNSS is supported on the 7705 SAR-8 CSMv2 with shelf v1 or v2, 7705 SAR-18, and 7705 SAR-Ax configured with a GNSS MDA.

Time references become qualified for use when they provide a valid input into system time. Of the qualified time references, the highest priority, which is the lowest numerical priority value, becomes the selected time reference. If all five time references are operationally up, the order of priority for the selected time reference is PTP or GPS, depending on which one has the lowest numerical priority value, then NTP, SNTP, and Holdover. When PTP or GPS are selected time references, the operational states of NTP and SNTP are automatically down.

15.31.2 Steps

1

In the navigation tree equipment view, right-click on the 7705 SAR on which you need to configure system time, and choose Properties.

2

Click on the Time tab and on the Input References tab.

3

Click Create. The Time References (Create) form opens.

4

Configure the required parameters.

If you configure the Ref Type parameter as PTP, a PTP clock and priority value must be assigned. If you configure GPS or GNSS, a GPS port or GNSS port and priority value must be assigned. See 15.26 “To configure an IEEE 1588 PTP clock on a 7705 SAR” (p. 483) for more information about configuring a PTP clock on a 7705 SAR.

5

Save your changes and close the form.

END OF STEPS

15.32 To configure SNTP on a 7705 SAR

15.32.1 Before you begin

SNTP is a compact, client-only version of NTP. The 7705 SAR supports SNTP in both unicast client modes and broadcast client modes. SNTP can only receive the time from SNTP or NTP servers and cannot be used to provide time services to other systems.

15.32.2 Steps

1

In the navigation tree equipment view, right-click on the 7705 SAR on which you need to configure SNTP, and choose Properties.

2

Click on the SNTP tab and configure the required parameters on the General tab.

3

Configure an SNTP server.

1. Click on the Server tab.
2. Click Create or select a server and click Properties. The SNTP Server (Create|Edit) form opens.
3. Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

15.33 To configure VPLS scalability on a 7705 SAR

15.33.1 Before you begin

The VPLS scalability can be configured only on 7705 SAR-8 and 7705 SAR-18 nodes. This configuration applies to SAR nodes from Release 21.4 R2 onwards.

15.33.2 Steps

1

On the equipment tree, choose Network→*NE* where *NE* is the object on which you need to configure VPLS scaling.

2

Right-click on the NE and choose Properties. The Network Element (Edit) form opens.

3

Click the System Forwarding Path Option tab and configure the Admin State parameter.

Note: By default, the Oper and Admin state parameters are set to Down. You must reboot the NE for the configured Admin State value to take effect. The Oper State parameter is updated when the NE is rebooted.

Once the Admin State is turned up, a total of 255 VPLS instances can be configured on the node.

4

Save your changes and close the form.

END OF STEPS

15.34 To configure NTP on supported devices

15.34.1 Before you begin

NEs can act as NTP servers, clients, or peers.

The 7705 SAR acts as an NTP client as well as an NTP server. The 7705 SAR as an NTP server, and the support for NTP on a VPRN service is supported from Release 21.4 R1 and later.

Note: For the Wavence SM, if the main and spare NTP servers are defined with the same IP address in NEtO, only the IP address for the main server will be displayed in the NFM-P. Main and primary NTP servers with the same IP address should be avoided.

For any failed operation related to NTP server addition or modification, you must close the current GUI session on the NTP tab and reopen it before proceeding. This action removes the failed entry from the GUI cache so that it is not applied again.

The tabs and parameters that are configurable vary depending on the NE.

15.34.2 Steps

1

On the equipment tree, choose Network→*NE* where *NE* is the object on which you need to configure NTP.

2

Right-click on the NE and choose Properties. The Network Element (Edit) form opens.

3

If you are configuring a 7705 SAR, Release 6.2 R1 or earlier, or a Wavence SM, click on the NTP tab and configure the NTP State parameter, and go to Step 6 . If you are configuring a 7705 SAR, Release 7.0 or later, go to Step 4 .

You must set the NTP State parameter to Enabled in order to configure the remaining NTP parameters.

4

Configure the NTP State parameter on the General tab.

You must set the NTP State parameter to Enabled in order to configure the remaining NTP parameters.

When you change the NTP State parameter from Disabled to Enabled, you must click Apply and confirm the action in order to activate the NTP tab.

5

Click on the NTP tab. The General tab is displayed.

6

Configure the required parameters.

7

If you are configuring NTP on an 1830 VWM, OS 6250 or OS 6450 device, go to Step 9. Otherwise, continue to Step 8.

8

Configure Authentication:

1. Click on the Authentication tab and click Create. The Network Time Protocol Authentication (Create) form opens.
2. Configure the required parameters.

   The values of the Authentication Key Id and Authentication Key parameters must be the same on the reference clock server and the NE.
3. Save your changes.
4. Go to Step 10.

9

Configure Authentication:

1. Click on the Authentication tab and click Load Authentication Keys. A dialog box appears. Click Yes.
2. Choose Resync→Resync All MIBs. A dialog box appears. Click Yes. The Authentication Keys are loaded.
3. Save your changes.

10

Configure an NTP server:

1. Click on the Server tab and click Create. The NTP Server (Create) form opens.
2. Configure the required parameters.

   Note:

   When the PTP Server parameter is enabled, the IP Address parameter becomes non-configurable. The address that is displayed in the field is not the PTP server address, it is an internal reference.

   The Authentication Key Id and NTP Version parameters are only displayed when the PTP Server parameter is not enabled.

   For 7705 SAR, NTP Version parameter supports version 2, version 3, and version 4 for IPv4 and only version 4 is supported for IPv6.

   If you are configuring NTP Server on an OS 6900 or OS 6860 or OS 6860E or OS 6860N or OS 6865 or OS 6465 or OS 6560 device, select Host mode as client.
3. Save your changes.

   If you are configuring Wavence SM, go to Step 21 .

11

If you are configuring NTP Peer on an OS 6900 or OS 6860 or OS 6860E or OS 6860N or OS 6865 or OS 6465 or OS 6560 device, continue to Step 12. Otherwise, go to Step 13.

12

Configure an NTP peer:

1. Click on the Server tab and click Create. The NTP Server (Create) form opens. Select Host mode as active.
2. Save your changes.
3. Go to Step 14.

13

Configure an NTP peer:

1. Click on the Peer tab and click Create. The NTP Peer (Create) form opens.
2. Configure the required parameters.

   Note:

For 7705 SAR, NTP Version parameter supports version 2, version 3, and version 4 for IPv4 and only version 4 is supported for IPv6.

3. Save your changes.

14

If you are configuring NTP Broadcast on an OS 6900 or OS 6860 or OS 6860E or OS 6860N or OS 6865 or OS 6465 or OS 6560 device, continue to Step 15. Otherwise, go to Step 16.

15

Configure NTP broadcast:

1. Click on the Server tab and click Create. The NTP Server (Create) form opens. Select Host mode as broadcast.
2. Save your changes.
3. Go to Step 17.

16

Configure NTP broadcast:

1. Click on the Broadcast tab and click Create. The NTP Broadcast (Create) form opens.
2. Configure the Router Instance parameter.
3. Select a source interface.
4. Configure the required parameters.

   A warning message is displayed if the Direction parameter is set to transmit when the NTP Server parameter is not enabled in Step 6 .

   When the value of the Direction parameter is set to transmit, the Authentication Key Id, NTP Version, and Time to Live parameters are displayed.
5. Save your changes.

17

Configure NTP multicast:

1. Click on the Multicast tab and click Create. The NTP Multicast (Create) form opens.
2. Configure the required parameters.

   When the value of the Direction parameter is set to transmit, the Authentication Key Id and NTP Version parameters are displayed.
3. Save your changes.

18

Click on the Associations tab and perform the following to view the properties of the NTP associations, if required:

1. Click Search. A list of NTP-associated NEs appears.
2. Choose an NE and click Properties. The Network Time Protocol Associations (View) form opens. You can view information about the status and mode of the NTP server.

3. Close the form.

19

Click on the Clients tab and perform the following to view the properties of the NTP clients, if required:

1. Click Search. A list of NTP clients appears.
2. Choose an NTP client and click Properties. The Network Time Protocol Clients (View) form opens. You can view information about the status and mode of the NTP client.
3. Close the form.

20

Configure an NTP Disabled interfaces:

1. Click on the Disabled Interfaces tab and click Create. The NTP Disabled interfaces (New instance) (create) form opens.
2. Add IP Address.
3. Save your changes.

21

Save and close the forms.

END OF STEPS

15.35 To configure NTP on 1830 VWM OSU devices

15.35.1 NTP on 1830 VWM devices

The 1830 VWM OSU has three NTP servers and the RMUs (Remote Managed Units include TLUs and PMUs) have one NTP server. You can configure the NTP parameters only on the 1830 VWM OSU device. The parameters are read-only on the other 1830 VWM devices.

15.35.2 Steps

1

On the equipment tree, expand Network→1830 VWM OSU→1830 VWM OSU Shelf .

2

Right-click on the 1830 VWM OSU Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the NTP tab and configure the required parameters in the General sub-tab.

4

Click on the Server tab, choose an NTP server, and click Properties. The NTP Server (Edit) form opens.

5

Configure the IP Address parameter.

6

Save and close the forms.

END OF STEPS

15.36 To configure SCADA on a 7705 SAR

15.36.1 Purpose

Perform this procedure to configure SCADA on a 7705 SAR-8 or 7705 SAR-18, Release 6.1 R1 or later. To configure SCADA, an ISC must be installed in the daughter card slot. See 6.8 “7705 SAR” (p. 222) for more information about the ISC.

15.36.2 Steps

1

In the navigation tree equipment view, navigate to the shelf object on which you want to configure SCADA. The path is Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the SCADA tab.

4

Select one or more ports that you want to configure as a SCADA bridge, and click Properties. The SCADA Bridge - MDDB (Edit) form opens.

All of the ports have the bridge type as MDDB by default. If you want to configure PCM, you must change the MDA mode on the ISC. See 15.78 “To configure an MDA” (p. 536).

5

Perform one of the following:

a. If you set the MDA mode of the ISC to MDDB, configure the general parameters and parameters in the MDDB panel.

b. If you set the MDA mode of the ISC to PCM, configure the general parameters and

parameters in the PCM panel.

The Squelch Reset Timeout (seconds) parameter only appears if you set the Squelch Reset Mode parameter to Auto.

The Force Active Master parameter only appears if the Redundant Mode parameter is set to Manual.

Note: The MDA mode of the ISC can be configured for VCB. See 15.37 “To configure voice conference bridging on a 7705 SAR” (p. 496) for more information.

6

Configure SCADA branches. You can create up to 32 branches per bridge.

1. Click on the Branches tab.
2. Click Create. The SCADA Branch (Create) form opens.
3. Configure the required parameters.
4. Click OK and confirm to close the form.
5. Configure more SCADA branches, as required.

   You can add SCADA MDDB and PCM branches as SAPs on an L2 access interface on a VLL Cpipe or Epipe service site. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for more information.

7

Click OK to save the configuration.

8

To activate or deactivate a SCADA bridge, select one or more SCADA bridges listed on the Shelf (Edit) form.

a. Click Turn Up and confirm to activate the SCADA bridge.

b. Click Shut Down and confirm to deactivate the SCADA bridge.

9

Save and close the form.

END OF STEPS

15.37 To configure voice conference bridging on a 7705 SAR

15.37.1 Purpose

Perform this procedure to configure VCB on a 7705 SAR-8 or 7705 SAR-18. To support VCB, the NE must be configured with an ISC.

The VCB application provides a simultaneous communication path between two or more voice circuits. VCBs are deployed in a central location with remote devices connected to the bridge via an NE over an IP/MPLS or TDM network.

VCBs can be used as a conference bridge with any-to-any connectivity (all branches participate) or as a bridge in broadcast mode where one branch broadcasts to the other branches that are in listen-only mode.

VCB can be configured in one of four applications. These applications are set at the card level. Each application uses a bridging algorithm that determines which branches control the management of the bridge and transmission of signals:

- VCB — One branch talks and all other branches on the bridge can hear.
- Broadcast — Only one branch on the bridge (fixed as branch 1) has control of the bridge to transmit, and all other branches are in listen-only mode.
- VCB Branch Initiate — Branches on the bridge are only enabled (unmuted) when the attached base station signals its presence by grounding the M-lead on the interface connected to the bridge. Upon receiving the grounded M-lead via T1/E1 ABCD bits or TDM PW signaling, the bridge unmutes the associated branch. When the ground is removed, the branch is muted again.
- Teleprotection — Each teleprotection relay transmits state information on discrete frequencies so that each relay can both hear what the other relays are transmitting as well as transmit its own information to the other relays.

15.37.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n.

3

Click the Daughter Card Slot with the Integrated Services card. The Daughter Card Slot form opens.

4

Click on the Daughter Card tab and configure the MDA Mode parameter as VCB.

5

Configure the VCB Application parameter. Choose one of the following options:

a. Broadcast

b. Tele-Protection

c. VCB

d. VCB Branch Initiate

6

In the Voice panel, configure the Companding Law and Signalling Type parameters. By default, the companding law is set to Mu-Law.

For information about which adapter cards support A-Law or Mu-Law, and which signalling types are configurable for each card, see the node documentation.

Note: Mu-Law encoding is used for North America, and provides up to 16 bridges per ISC, up to 24 branches per bridge, and up to 384 total branches per ISC. A-Law is used for the rest of the world. A-Law encoding provides up to 16 bridges per ISC, up to 32 branches per bridge, and up to 512 total branches per ISC.

7

Click OK and confirm the changes.

8

On the NE navigation tree, right-click on the Shelf and choose Properties. The Shelf (Edit) form opens.

9

Click on the SCADA tab.

10

Select one or more VCBs and set the Bridge Admin Status parameter to Up.

11

Click Turn Up.

12

Click on Properties. The SCADA Bridge - VCB form opens.

13

If you are configuring a VCB branch initiate, configure the VCB Idle Code and VCB Seized Code parameters as required.

14

Click on the Branches tab and click Create to create a SCADA branch. The SCADA Branch (Create) form opens

15

On the General tab of the SCADA branch, set the Administrative State parameter to Up, and configure the Description, Gain Input (dB) and Gain Output (dB) parameters.

16

Click OK and confirm the changes.

17

Configure more SCADA VCB branches as required. You can add SCADA VCB branches as SAPs on an L2 access interface on a VLL Epipe or Cpipe service site. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for more information.

END OF STEPS

15.38 To configure a 7705 SAR MW link

15.38.1 Purpose

Perform this procedure to configure a microwave link on a 7705 SAR-8 or 7705 SAR-18. To support an MW link, the NE must be configured with 4 × MWA Gig Ethernet (2 TX, 2 SFP) + 4 × Gig Ethernet SFP daughter cards. See 15.78 “To configure an MDA” (p. 536) for information about how to configure daughter cards.

Note: The NFM-P supports the creation of an MW link in 1+0 protection mode and 1+1 protection mode. You can only configure 1+1 protection on an NE in network mode.

15.38.2 Steps

1

In the navigation tree equipment view, choose Network→NE→*Shelf* where *Shelf* is the 7705 SAR shelf object on which you need to create an MW link.

2

Click on the shelf object to expand the tree view.

3

Right-click on the MW Links item and choose Create MW Link. The Create MW Link (Create) form opens, with the General tab displayed.

4

Configure the Description parameter, if required.

5

Click on the MW Link tab.

6

Configure the required parameters in the Hold Time panel.

7

Click on the States tab and configure the Administrative State parameter.

8

Click on the Protection tab and configure the Protection Scheme parameter.

If you choose the value One Plus One Hsb, the Protection form displays additional panels.

9

To configure 1+1 radio protection, configure the required parameters in the Equipment Protection Scheme, Radio Protection Scheme, and Transmission Protection Scheme panels.

10

Click OK. The MW Link (Create) form closes and the new MW link appears under the MW Links item in the navigation tree.

11

To modify the MW Link properties, right-click on the MW link and choose Properties. The MW Link (Edit) form appears with the General tab displayed.

12

Click on the MW Link tab and configure the required parameters in the Hold Time panel.

Note: The Peer IP Address parameter in the Peer Info panel is read-only and is displayed as follows:

- When a Wavence SA is in standalone mode, the Peer IP Address displays 0.0.0.0.
- When the Wavence SA is in single NE mode, the Peer IP Address displays the IP address of the other end Wavence SA if the other end of the radio link terminates on a Wavence SA in standalone mode.
- When the Wavence SA is in single NE mode, the Peer IP Address displays the IP address of the other end 7705 SAR if the other end of the radio link terminates on a Wavence SA in single NE mode.
- The Peer IP Address changes only when the 7705 SAR discovers a new peer after the radio link is operationally up between both Wavence SAs or when the 7705 SAR reboots.

13

Configure the Peer Discovery parameter in the Peer Discovery panel.

14

Click on the Statistics tab to view statistical information, if required.

15

Save and close the form.

END OF STEPS

15.39 To configure a 7705 SAR MW link member

15.39.1 Purpose

Perform this procedure to configure a 7705 SAR MW link member on the 7705 SAR-8 or 7705 SAR-18. To support an MW link and an MW link member, the NE must be configured with 4 × MWA Gig Ethernet (2 TX, 2 SFP) + 4 × Gig Ethernet SFP daughter cards. See 15.78 “To configure an MDA” (p. 536) for information about how to configure daughter cards.

If you configured 1+1 HSB protection in 15.38 “To configure a 7705 SAR MW link” (p. 499) , you can configure a main radio and a spare radio for each MW link. By default, the main radio assumes the active status and the spare radio assumes the standby status.

If you configured 1+1 HSB protection and you want to revert back to 1+0 protection, you must delete the spare radio. You cannot delete the main radio.

1+1 HSB protection is displayed on the NFM-P physical topology map. The protection state is shown on the links, and 1+1 links are grouped in a common link group. The expansion of links on the physical map displays eight links. See 4.1 “Topology map types” (p. 169) for more information about the NFM-P physical topology maps.

15.39.2 Steps

1

Choose Network→NE→Shelf→Card Slot n→Daughter Card Slot n→Port n/n/n where the port object is the port that you need to use for the MW link.

Note: Only the first four ports on the MDA are microwave-aware.
If a port is associated with an MW link, the port cannot be used as a timing reference. Also, if a port is already being used as a timing reference, the port cannot be associated with an MW link member.

2

Right-click on a port and choose Properties. The Physical Port (Edit) form opens.

3

Configure the Encapsulation Type parameter to be Dot1 Q.

4

Configure the Mode parameter.

Note: Setting the Mode parameter to Access allows you to later associate the MW link with an L2 or L3 access interface for services. Setting the Mode parameter to Access is the typical setting. Setting the parameter to Network allows you to associate the MW link with a network interface.
1+1 HSB radio protection can only be configured on a port in Network mode.

5

Click OK and confirm to close the form.

6

Create an MW link member:

a. If you configured 1+0 radio protection in 15.38 “To configure a 7705 SAR MW link” (p. 499) , right-click on the MW link you created and choose Create MW Link Member. The Create MW Link Member step form opens.

b. If you configured 1+1 HSB radio protection in 15.38 “To configure a 7705 SAR MW link” (p. 499) , right-click on the MW link you created and choose Create MW Link Main Radio. The Create MW Link Member step form opens.

7

Configure the required parameters on the Specify the Member Properties form.

The Database Filename parameter must be the same name as the database file in the config file of the cflash directory and must be in a *.tar format.

8

Click Next. The Select Ports form opens.

9

Choose the port from the list and click Finish. Only eligible ports are displayed in the list and only one port per link is supported.

To configure a main radio, you must choose a port in an odd numbered slot.

10

Select the View the newly created Port Termination check box if you need to view the MW link member properties, and click Close.

11

To configure a spare radio, right-click on the MW link where you created the main radio, and choose Create MW Link Spare Radio. Repeat Step 7 to Step 10 .

To configure a spare radio, you must choose a port in an even numbered slot.

12

Click on the Radio tab. The tab displays information about MPT hardware and the software state of the connected radio. You can enable PM statistics for G.826, ACM, and Power. Enable the checkboxes as required.

13

Associate the MW link with an L2/L3 access interface or a network interface, if required.

1. If you set the MW Link Member port Mode to Access in Step 4 , you can associate the MW link with a service, if required.

   The MW links can be used in the L2 Access Interfaces of the following service types:
   - VLL Epipe (see 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174))
   - VLL Ipipe (see 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174))
   - VPLS (see 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332))
   - MVPLS (see 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332))

   The MW links can also be used in the L3 Access Interfaces of the following service types:
   - VPRN (see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656))
   - IES (see 78.28 “To configure an L3 access interface on an IES site” (p. 2472))

   The created access interface appears in the list on the L2 Access Interface or L3 Access Interface tab of the MW Link (Edit) form, as appropriate.

   See the specific service chapter and forms when you create the required L2 or L3 access interfaces. Specify the MW Link Member as the Terminating Port.
2. If you set the MW Link Member port Mode setting to Network in Step 4 , the Network Interface tab is displayed. Click on the tab and then the Create button to create a network interface. The network interface that you create appears in the list. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for more information.

14

Click on the Statistics tab to view radio analog statistical information, if required.

15

Click on the Faults tab to view alarm information, if required.

16

To modify the MW Link Member properties, right-click on the MW link member and choose Properties. The MW Link Member (Edit) form appears with the General tab displayed.

17

Configure the required parameters.

18

Save and close the form.

END OF STEPS

15.40 To configure 7705 SAR auxiliary alarm definitions

15.40.1 Purpose

Perform this procedure to configure auxiliary alarm definitions for the auxiliary alarms daughter card on a 7705 SAR.

15.40.2 Steps

1

In the navigation tree equipment view, navigate to the shelf object on which you want to configure auxiliary alarm definitions. The path is Network→*NE*→Shelf.

2

Right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.

3

Click on the External Alarms tab.

4

As required, configure auxiliary alarm digital inputs.

1. Click on the Auxiliary Alarm Digital Inputs tab.
2. Select a digital input from the list and click Properties. The Auxiliary Alarm Digital (Edit) form opens.
3. Configure the required parameters.
4. Click OK. The Auxiliary Alarm Digital (Edit) form closes.
5. Repeat 1 to 4 to configure more auxiliary alarm digital inputs, as required.

5

As required, configure auxiliary alarm port inputs.

1. Click on the Auxiliary Alarm Port Inputs tab.
2. Click Create. The Auxiliary Alarm Ports (Create) form opens.
3. Click Select beside the CLI Name parameter and select a port reference. Only Ethernet ports can be selected.
4. Configure the required parameters.
5. Click OK and confirm to close the form.
6. Repeat 1 to 5 to configure more auxiliary alarm port inputs, as required.

6

Click on the Auxiliary Alarm Definitions tab on the Shelf (Edit) form.

7

Click Create. The Auxiliary Alarm Definition (Create) form opens.

8

Configure the required parameters.

9

Click on the Inputs tab and configure the required parameters.

10

Click on the Outputs tab and configure the required parameters.

11

Save and close the forms.

END OF STEPS

15.41 To manage an OmniSwitch running configuration

CAUTION

Service Disruption

When you reboot an OmniSwitch that is in service, it is service-affecting.

Ensure that the reboot activity occurs during a maintenance window.

15.41.1 Steps

1

In the navigation tree equipment view, expand the OmniSwitch object.

2

Right-click on a device and choose Properties. The Network Element (Edit) form opens.

3

Click on a shelf object. The Shelf (Edit) form opens.

4

Click on the Software Control Module tab.

5

Configure the Command to Apply parameter.

a. If you specify Copy certified to working, configure the Active Timeout parameter.

b. If you specify Reload from working, configure the Delayed Activation Timer parameter.

6

Save and close the form.

END OF STEPS

15.42 To configure OmniSwitch health monitoring

15.42.1 Before you begin

The health monitoring function monitors the consumable resources, such as bandwidth and CPU usage of the OmniSwitch at fixed intervals and collects the current values for each resource being monitored. You can specify the resource threshold limits. If a resource value is greater or less than the threshold, traps are sent to the NFM-P.

15.42.2 Steps

1

On the equipment tree, right-click on the OmniSwitch device for which you want to configure health monitoring and choose Properties. The Network Element (Edit) form opens.

2

On the Network Element (Edit) form navigation tree, click on a shelf object. The Shelf (Edit) form opens.

3

Click on the Health Monitoring tab and configure the required parameters.

4

Click on the Statistics tab and search for Card Health Stats (Physical Equipment) or Device Health Stats (Physical Equipment) to view card statistics or chassis statistics. Statistics can also be viewed from the card-level statistics tab.

5

Click on the Faults tab to view alarm information.

6

Save your changes and close the forms.

END OF STEPS

15.43 To configure a CCM on a 7950 XRS-20

15.43.1 Steps

1

On the equipment tree, right-click on the 7950 XRS-20 device for which you want to configure CCM properties and choose Properties. The Network Element (Edit) form opens.

2

On the Network Element (Edit) form navigation tree, expand the Shelf object.

3

Right-click on a CCM object and choose Properties. The CCM (Edit) form opens.

Alternatively, you can open a CCM properties form by right-clicking on the associated CPM object and choosing View CCM from the contextual menu.

4

To configure a flash memory module:

1. Click on the Flash Memory Modules tab.
2. Select a flash memory module in the list and click Properties. The Flash Memory (Edit) form opens.
3. Configure the Administrative State parameter.
4. Save your changes and close the form.

5

On the CCM Shelf (Edit) form, click View CPM to open the properties form of the associated CPM.

See 15.74 “To configure a CPM” (p. 531) for more information about the CPM properties form.

6

Close the forms.

END OF STEPS

15.44 To manage the internal fan on a 7210 SAS-D

15.44.1 Purpose

Perform this procedure to change the administration mode and view the fan operational state for the 7210 SAS-D 128 Mbyte variant.

15.44.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Fans→Fan Tray.

2

Right-click on Fan Tray and choose Properties. The Fan Tray (Edit) form opens.

3

Click on the Fans tab. A list form opens.

4

Choose a fan from the list. The Fan (Edit) form opens.

5

Configure the Admin Mode parameter. The operational state of the fan (On or Off) is displayed in the Oper Mode field.

6

Save your changes and close the forms.

END OF STEPS

15.45 To enable fan speed monitoring on a 7x50 device

15.45.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Fans→Fan Tray.

2

Right-click on Fan Tray and choose Properties. The Fan Tray (Edit) form opens.

3

Enable the Monitor Fan Speed parameter.

4

Save your changes and close the form.

TCA and MIB statistics are created.

Alarms will be generated if fan speed exceeds threshold limits (Minor - 75%, Major - 80%, Critical - 100%).

End of steps

15.46 To configure a power supply tray

15.46.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Power Supplies→Power Supply Tray.

2

Right-click on a power supply tray and choose Properties. The Power Supply Tray (Edit) form opens.

3

Configure the required parameters.

4

Configure the Power Supply Type parameter from the drop-down menu.

Note: If you are configuring for OS 6465-P6 and OS 6465-P12, depending on the power supply unit type (Third Party, ALE, or Phoenix Contact) the form displays different options. Select one of the following, as appropriate:

- Third Party - The Power (Watts) ranges from 50W to 200W.
- ALE - Select the Power Supply Input Type parameter from the drop-down menu. The default values for Power (Watts) are LO-AC = 75W, HI-AC = 240W, 24VDC = 240W, and 48VDC = 240W.
- Phoenix Contact - Select the Power Supply Input Type parameter from the drop-down menu. The default values for Power (Watts) are 24VDC = 240W, 48VDC = 240W.

5

Save your changes and close the form.

End of steps

15.47 To configure a power management zone

15.47.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Power Management Zones tab. A list of power management zones is displayed.

3

Select a power management zone and click Properties. The Power Management Zone (Edit) form opens.

4

Configure the power management parameters.

5

Click on the Utilization and Requirements tab to view information about power usage in the power management zone. Click on the ellipsis button (...) for detailed information about each entry.

6

Save your changes and close the forms.

END OF STEPS

15.48 To configure a PCM tray

15.48.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Power Connection Modules→PCM Tray.

2

Configure one or more PCM trays.

a. To configure one PCM tray, right-click on a PCM tray and choose Configure PCM. The PCM Tray (Edit) form opens.

b. To configure multiple PCM trays in bulk, right-click on empty PCM tray slots and choose Configure PCM. The PCM Tray (Multiple Instances) (Edit) form opens.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

15.49 To provision an APEQ

15.49.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Power Supplies→Power Supply Tray.

2

Right-click on a power supply tray and choose Configure PEQ. The Power Supply Tray (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

15.50 To configure a variable power supply APEQ

15.50.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Power Supplies→Power Supply Tray.

2

Right-click on a power supply tray and choose Properties. The Power Supply Tray form opens with the General tab displayed.

3

Click on the Power Supply tab and configure the Input Power Mode in the Power Details panel.

4

Save your changes and close the form.

END OF STEPS

15.51 To restart an 1830 VWM shelf

15.51.1 Before you begin

Ensure that the:

- 1830 VWM device is managed by the NFM-P.
- value of the Connection State parameter in the Shelf Specifics tab of the Shelf (Edit) form is Connected for all of the 1830 VWM devices except 1830 VWM OSU.

15.51.2 Steps

NOTICE

Service affecting

Cold restart is service-affecting.

Ensure that you perform cold restart only during a scheduled maintenance period.

1

On the equipment tree, expand Network→1830 VWM.

2

Right-click on the Shelf object and click Properties. The Shelf (Edit) form opens.

3

Click on the Shelf Specifics tab and configure the parameters in the Shelf Restart panel.

4

Save your changes and close the form.

END OF STEPS

15.52 To configure optical protection switching on an 1830 VWM OPS shelf

15.52.1 Steps

1

On the equipment tree, expand Network→1830 VWM OSU→1830 VWM OPS shelf.

2

Right-click on the OPS OSM card slot object and choose Properties. The Card Slot (Edit) form opens.

3

Click on the Card Specifics tab, then on the Protection Switch Management tab, and configure the Switch Command parameter.

Note: The OPS shelf supports the following external switch commands:

- Clear - clears all current external switch requests
- Forced Switch to Worker
- Forced Switch to Protection

4

Configure the remaining parameters.

5

Save your changes and close the form.

END OF STEPS

15.53 To configure bi-directional protection switching on an 1830 VWM OPS shelf

15.53.1 Steps

1

Configure an OPS OSM-DSV card on an 1830 VWM OPS shelf:

1. On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.
2. Assign the OPS OSM-DSV card, see 15.56 “To assign a card type” (p. 517).
3. On the Card Slot (Edit) form, click on the Card Specifics tab and configure the Assigned Status parameter as Assigned.
4. Click Apply and confirm your actions.

5. Configure the required parameters in the OPS OSM-DSV Specifics panel.
6. Save your changes and close the form.

2

Configure bi-directional protection switching for the PSE-R:

1. On the equipment tree, right-click on the 1830 VWM-OSU NE and choose Properties. The NE (Edit) form opens.
2. Click on the PSE-R PMU-D group tab and click Create. The PSE-R PMU-D group (Create) form opens.
3. Configure the PMUD Group parameter.
4. Select one to four PMUD in the PMUD panels.
5. Click on the OPS PSE-R tab and click Create. The OPS PSE-R (Create) form opens.
6. Configure the required parameters.

   Note: You must select a PMUD Group.
7. Click Apply. The OPS OSM PSE-R (Edit) form appears with the OPS OSM PSE-R protection switch management tab.
8. Configure the Activate parameter as True and configure the required parameters, if applicable.
9. Click Apply and confirm you actions.

3

Configure bi-directional protection switching for the PSE-L:

1. Click on the OPS PSE-L tab and click Create. The OPS PSE-L (Create) form opens.
2. Configure the Description parameter.
3. Select an OPS OSM-DSV card for the Slot parameter.
4. In the Worker Loop Monitor Interface panel, configure the port as OSC1.
5. In the Protection Loop Monitor Interface, configure the port as OSC2.
6. Click Apply. The OPS OSM PSE-L (Edit) form appears with the OPS OSM PSE-L protection switch management tab.
7. Configure the Activate parameter as True and configure the required parameters, if applicable.
8. Click Apply and confirm your actions.

4

Save your changes and close the forms.

END OF STEPS

15.54 To configure an OPS protection audit entity on an 1830 VWM OSU shelf

15.54.1 Steps

1

On the equipment tree, expand Network.

2

Right-click on the 1830 VWM OSU object and choose Properties. The Network Element (Edit) form opens.

3

Click on the OPS Protection Audit Entity tab and click Create. The OPS Protection Audit Entity (Create) form opens.

4

Choose the near-end and far-end 1830 VWM OPS OSM cards in Slot A and Slot Z. The shelf and card details are populated automatically.

5

Configure the remaining parameters.

6

Save your changes and close the forms.

END OF STEPS

15.55 To activate or deactivate a PAE

15.55.1 Steps

1

Configure a PAE. See 15.54 “To configure an OPS protection audit entity on an 1830 VWM OSU shelf” (p. 515).

2

On the equipment tree, expand Network.

3

Right-click on the 1830 VWM OSU object and choose Properties. The Network Element (Edit) form opens.

4

Click on the OPS Protection Audit Entity tab, choose an entry, and click Properties. The OPS Protection Audit Entity (Edit) form opens.

5

Set the Activate Parameter to True or False depending on whether you need to activate or deactivate the PAE.

6

Save your changes and close the forms.

END OF STEPS

Procedures for card and card slot object configuration

15.56 To assign a card type

15.56.1 Purpose

Perform this procedure to assign a card type to an empty card slot or SFM slot.

The available configurations vary depending on the device; not all NEs support all configurations.

15.56.2 Steps

1

On the equipment tree, right-click on the device where you want to assign a card type and choose Properties. The Network Element (Edit) form opens.

2

On the Network Element (Edit) form navigation tree, expand the Shelf object.

3

Right-click on an empty Card Slot object or SFM Slot object and choose Configure Card. The Card Slot (Create) or SFM Slot (Create) form opens.

4

Configure the required parameters.

Note: The correct chassis mode must be configured for the assigned card types. The chassis mode determines the behavior of the card and establishes the scaling limits and available features. See 15.13 “To configure the device chassis mode” (p. 471) and 15.19 “To enable mixed mode” (p. 476) for more information.

5

Configure the required parameters in the Virtual Scheduler panel.

6

Save your changes and close the form.

The card and slot appear in the navigation tree and in the equipment inventory list.

END OF STEPS

15.57 To configure an Xiom-s card slot

15.57.1 Purpose

Perform this procedure to configure an Xiom-s card slot on a 7750 SR-s series NE. Each Xiom-s supports up to two MDA-s.

Xiom-s card slots are indicated in the equipment tree using the index of ‘x1’ or ‘x2’. The MDA-s in an Xiom-s uses an index ‘x1/1’, ‘x1/2’, ‘x2/1’, or ‘x2/2’ reflecting the fact that these are contained in Xiom-s ‘x1’ or ‘x2’.

15.57.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot.

2

Perform one of the following:

a. To configure a new Xiom-s, right-click on the Xiom Card Slot object and choose Configure Card. The Xiom Card Slot (Create) form opens.

b. To modify an existing Xiom-s, right-click on the Xiom Card Slot object and choose Properties. The Xiom Card Slot (Edit) form opens.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

15.58 To assign a network queue policy to a forwarding plane

15.58.1 Before you begin

Perform this procedure to assign network queue policies to forwarding planes on 7450 ESS, 7750 SR, and 7950 XRS NEs.

15.58.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

On the General tab, select a network queue policy in the Network Queue Policy panel.

5

Save your changes and close the forms.

END OF STEPS

15.59 To assign an FP Resource policy to a forwarding plane

15.59.1 Before you begin

Forwarding Plane Resource policies allow custom allocation of card slot resources to queues. See 50.72 “To configure an FP Resource policy” (p. 1615).

15.59.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

On the General tab, select an FP Resource policy in the FP Resource Policy panel.

The FP Resource Policy panel is available only when FP Resource policies are supported on the NE.

5

Save your changes and close the forms.

END OF STEPS

15.60 To configure a network ingress pool on a forwarding plane

15.60.1 Before you begin

Perform this procedure to configure network ingress pools on forwarding planes on 7450 ESS, 7750 SR, and 7950 XRS NEs.

15.60.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

Click on the QoS Pool tab, choose a network ingress pool in the list, and click Properties. The QoS Pool form opens.

5

Configure the required general parameters.

6

Save your changes and close the forms.

END OF STEPS

15.61 To view the operational MC path management properties of a forwarding plane

15.61.1 Before you begin

MC path management properties are not configurable. However, you can perform this procedure to view the default operational MC path management properties of forwarding planes on 7450 ESS, 7750 SR, and 7950 XRS NEs.

15.61.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

Click on the QoS Pool tab, choose an MC path management pool in the list, and click Properties. The QoS Pool form opens.

5

View the default settings. MC path management properties are not configurable.

6

Save your changes and close the forms.

END OF STEPS

15.62 To configure egress WRED queue control on an XCM, IOM 3 or IMM forwarding plane

Note: You cannot configure a forwarding plane on an XCM card until you have provisioned the corresponding XMA card.

15.62.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot (Edit) form opens.

3

Click on the Forwarding Plane tab.

4

Choose an entry in the list and click Properties. The Forwarding Plane (Edit) form opens.

5

Configure the required parameters in the Egress WRED Queue Control panel.

6

Select a slope policy for the forwarding plane.

7

Save your changes and close the forms.

END OF STEPS

15.63 To configure ingress policy accounting policer limits on a forwarding plane

15.63.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

On the General tab, configure the Policer Limit parameter (under Ingress Policy Accounting).

The value for the Policer Limit parameter must not exceed the value configured for the Stats Pool limit parameter (under Policy Accounting).

5

Save your changes and close the forms.

END OF STEPS

15.64 To configure IMPM on an XCM, 2 x XP MDA IOM 3, or IMM forwarding plane

Note: See 15.79 “To configure IMPM on an MDA” (p. 539) for information about configuring IMPM on an MDA that is not installed in an XCM, 2 × XP MDA IOM 3, or in an IMM.

You cannot configure a forwarding plane on an XCM card until you have provisioned the corresponding XMA card.

15.64.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot (Edit) form opens.

3

Click on the Forwarding Plane tab.

4

Choose an entry in the list and click Properties. The Forwarding Plane (Edit) form opens.

5

Click Clear in the Ingress Multicast Path Management panel to remove the selected BW policy. The Select button is enabled.

6

Select a BW policy.

7

Configure the Administrative State parameter.

8

Configure the required bandwidth parameters.

9

Save your changes and close the forms.

END OF STEPS

15.65 To configure an ingress queue group on a forwarding plane

Note: You can configure ingress queue groups only on the forwarding planes of XCM, IOM3, and IMM cards.

You cannot configure a forwarding plane on an XCM card until you have provisioned the corresponding XMA card.

15.65.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot (Edit) form opens.

3

Click on the Forwarding Plane tab.

4

Choose an entry in the list and click Properties. The Forwarding Plane (Edit) form opens with the General tab displayed.

5

Perform one of the following:

a. Click on the Access Ingress Queue Group tab.

b. Click on the Network Ingress Queue Group tab.

6

Click Create. The Forwarding Plane Access|Ingress Queue Group (Create) form opens.

7

Select a Queue Group Template Policy in the Queue Group Instance panel.

Note: If you select a policy that contains queues, a configuration error message appears when you try to apply the changes.

8

Configure the Instance ID and Description parameters.

9

Select an Accounting Policy in the Accounting Policy panel.

Note: Only accounting policies with Queue Group Packets as record type are supported. You can also choose to create an accounting policy from this form by clicking the Create button instead of the Search button. See 87.37 “To configure an AA accounting policy” (p. 2849) for information on creating an accounting policy and the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting accounting statistics.

10

Configure the Collect Accounting Statistics parameter.

11

Select a Policer Control Policy in the Policer Control Policy panel.

12

Click Create in the Ingress Policer Control Override panel. The Ingress Policer Control Override (Create) form opens, with the General tab displayed.

13

Configure the required parameters and save your changes.

14

Click on the Level Override Policy Items tab.

15

Select an item from the list and click OK. The FP Ingress Policer Level Override form opens, with the General tab displayed.

16

Click on the Override tab and configure the Maximum Cumulative Buffer Space (bytes) parameter.

17

Save your changes.

18

Click on the Override Policy Items tab and click Create. The Forwarding Plane Ingress Queue Group Policer Override (Create) form opens.

19

Select a policer.

20

Click on the Override tab and configure the required parameters in the Overridden Queue Group Template Policer panel.

The parameters are configurable when the associated override check box is enabled.

21

Save your changes and close the forms.

END OF STEPS

15.66 To configure NE DDoS protection on a forwarding plane

15.66.1 Purpose

Perform this procedure to configure the maximum number of enforcement policers for a DDoS protection configuration on a forwarding plane. See the section on DDoS protection in the *NSP*

System Administrator Guide for more information.

Note: You cannot configure a forwarding plane on an XCM card until you have provisioned the corresponding XMA card.

15.66.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, choose an entry in the list, and click Properties. The Forwarding Plane form opens.

4

Configure the NE DDoS Protection - Dynamic Enforcement Policer Pool Size parameter. Enable the Default check box to specify a default value.

5

Save your changes and close the forms.

END OF STEPS

15.67 To enable named pool mode

15.67.1 Purpose

Perform this procedure to allow named pools to be created for an MDA. Set the Pool Mode parameter when you are configuring the card. Enabling the Named Pool mode is mutually exclusive with enabling Stable Pool sizing.

CAUTION

Service Disruption

The Pool Mode parameter can be enabled and disabled at anytime, however changing the pool mode resets the IOM when MDAs are provisioned on the slot.

If MDAs are not provisioned the IOM is not reset.

15.67.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the IO Card tab and configure the Pool Mode parameter.

4

Click OK and confirm your actions. The form closes.

END OF STEPS

15.68 To enable Stable Pool Sizing

15.68.1 Purpose

Perform this procedure to enable an even split of buffering amongst all possible MDAs.

When Stable Pool sizing is enabled, then for each MDA, the buffering is split between all ports based on the port maximum bandwidth and adjusted by the port's modify buffer allocation rate percentages. You cannot enable Stable Pool Sizing on any forwarding plane on a card while the card is either administratively or operationally in Named Pool mode.

15.68.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, select a forwarding plane, and click Properties. The Forwarding Plane form opens.

4

Configure the Stable Pool Sizing parameter and click OK. The form closes.

5

Save your changes and close the forms.

END OF STEPS

15.69 To enable Ingress Buffer Allocation

15.69.1 Purpose

Perform this procedure to enable the allocation of an ingress buffer percentage on a given forwarding plane. The allocation you specify will be set aside for use only by the ingress pools.

15.69.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, select a forwarding plane, and click Properties. The Forwarding Plane form opens.

4

Configure the Ingress Buffer Allocation (%) parameter and click OK. The form closes.

5

Save your changes and close the forms.

END OF STEPS

15.70 To initialize drop priority mode on card forwarding plane

15.70.1 Purpose

Perform this procedure to initialize the drop priority mode on card forwarding plane.

15.70.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the Card Slot object and choose Properties. The Card Slot form opens.

3

Click on the Forwarding Plane tab, select a forwarding plane, and click Properties. The Forwarding Plane form opens.

4

Configure the Initialize drop priority mode parameter, save your changes and close the form.

5

Save your changes and close the forms.

END OF STEPS

15.71 To select system resource profile policies for the 7210 SAS-R or 7210 SAS-S/Sx VC

15.71.1 Purpose

For the 7210 SAS-R or 7210 SAS-S/Sx VC, system resource allocations are configured in a system resource profile policy. See 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382) .

Perform the following procedure to assign a policy to the card slot of a 7210 SAS-R or 7210 SAS-S/Sx VC. You must distribute the policy to the device first.

Note: When the policy is assigned to a card slot and the SAP Aggregate Meter parameter is modified, you must reboot the card or the node for the modified value to take effect.

15.71.2 Steps

1

On the equipment tree, expand Network→7210 SAS-R or 7210 SAS-S/Sx VC→Shelf→Card Slot *n*.

2

Right-click on the card slot object and choose Properties. The Card Slot form opens.

3

Click on the IO Card tab and select a policy in the System Resource Profile Policy panel.

4

To view the currently effective resource profile for the card slot, click on the Active Resource Profile tab.

5

Click OK and confirm your changes. The form closes.

END OF STEPS

15.72 To configure OmniSwitch stacks

15.72.1 Steps

1

On the equipment tree, expand Network→OmniSwitch object.

2

Right-click on the OmniSwitch object and choose Properties. The Network Element form opens.

3

On the Network Element tree, expand the Shelf object.

4

Click on a slot object. The Card Slot form opens.

5

Click on the Stack Configuration tab and configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

15.73 To configure an OmniSwitch CPU temperature threshold

15.73.1 Steps

1

On the equipment tree, expand Network→OmniSwitch object.

2

Right-click on the OmniSwitch object and choose Properties. The Network Element form opens.

3

On the Network Element tree, expand the Shelf object.

4

Click on a slot object. The Card Slot form opens.

5

Click on the Hardware Environment tab and configure the Temperature Threshold (Celsius) parameter.

6

Save your changes and close the forms.

END OF STEPS

15.74 To configure a CPM

15.74.1 Purpose

Perform this procedure to configure CPM, CPM-s, CPM2, or CPM2-s objects.

15.74.2 Steps

1

On the equipment tree, right-click on the NE and choose Properties. The Network Element form opens.

2

On the Network Element tree, expand the Shelf object.

3

Right-click on Card Slot A or Card Slot B and choose Properties. The Card Slot form opens.

Alternatively, for the 7950 XRS, you can open a CPM Card Slot properties form by right-clicking on the associated CCM object and choosing View CPM from the contextual menu.

Note: Card Slot A and Card Slot B are automatically configured with CPMs. The CPMs cannot be removed, and they cannot be configured in any other card slot.

4

Configure the required parameters.

5

If the Flash Memory Module tab is available, you can perform the following steps. Depending on the device, only certain flash memory modules may be supported.

1. Click on the Flash Memory Modules tab, select a flash memory module, and click Properties. The Flash Memory form opens.
2. Configure the Administrative State parameter and click OK. The form closes.

Note:

You can configure flash memory modules on the 7950 XRS-20 from the CCM properties form. See 15.43 “To configure a CCM on a 7950 XRS-20” (p. 507) for more information.

6

If you are configuring a CPM on a 7950 XRS-20, the View CCM button is available. If required, click View CCM to open the properties form of the associated CCM.

See 15.43 “To configure a CCM on a 7950 XRS-20” (p. 507) for more information about the CCM properties form.

7

Save your changes and close the forms.

END OF STEPS

15.75 To configure a CPRI rate and channel management for a CDR channel on an 1830 VWM TLU or ITP card slot

15.75.1 Steps

1

On the equipment tree, expand Network→1830–VWM-OSU NE→1830–VWM-TLU shelf→Card Slot 1 TLU or Card Slot 1 ITP.

2

Right-click on the Card Slot 1 TLU or Card Slot 1 ITP object and choose Properties. The Card Slot (Edit) form opens.

3

Choose a CDR channel interface and click Properties. The CDR Channel (Edit) form opens.

4

Configure the CDR Channel Rate parameter.

5

Configure the CDR Channel Used For Mgmt parameter. The parameter is set to Disabled by default.

When the CDR Channel Used For Mgmt parameter is set to Enabled, the following configurations changes appear:

- The CDR Channel Rate parameter is set to 1 GBE.
- The port Administrative State parameter is set to UP.

- The LOS Propagation Mode parameter is set to Laser off.

Editing of the following parameters is not supported:

- Administrative State
- AINS mode
- Loopback, PRBS and monitoring settings
- CDR Channel Rate
- DWDM frequency

6

Save your changes and close the forms.

END OF STEPS

15.76 To modify the SFP profile and SFP profile rate on the 1830 VWM

15.76.1 Purpose

Perform this procedure to modify SFP profile or SFP profile rate associated with the CDR channel on the 1830 VWM TLU–9, 9M, and ITP devices.

15.76.2 Steps

1

Perform one of the following:

a. From the device:
   1. On the equipment tree, right-click on the 1830 VWM device object and choose Properties. The Network Element (Edit) form opens.
   2. Click on the SFP Profiles tab. A list of 32 default SFP profiles are displayed.
   3. Double-click on the SFP profile. The VWM SFP Profiles window opens with the General tab displayed.

      If required, edit the SFP Profile Name.

b. From the shelf:
   1. On the equipment tree, expand Network→1830 VWM.
   2. Right-click on the shelf object and choose Properties. The Shelf (Edit) form opens.
   3. Click on the Shelf Specifics tab.
   4. Go to Shelf SFP Profile section.

      If required, edit the SFP Profile Name. By default, Profile1 is assigned.
   5. Click Select. A list of 32 default SFP profiles are displayed.

2

Click on the SFP Profiles Rates tab to display the list of supported SFP rates.

3

Double-click on the selected SFP profile rate. The VWM SFP Profile Rate window opens.

4

Configure the SFP Profile Rate parameter from the drop-down menu..

5

Save your changes and close the form.

END OF STEPS

15.77 To upgrade a license path on an IOM-1, IOM5-e, and XIOM-s

15.77.1 Purpose

Perform this procedure to upgrade the license-level for IOM-1 or IOM5-e on 7750 SR platforms, Release 16.0 R4 and later, and for XIOM-s on 7750 SR-s platforms, Release 20.2 R1 and later.

Before you perform a license level upgrade, the NE bof should be pointed to the latest license file, and then you execute admin system license validate and admin system license activate.

15.77.2 Steps

1

On the equipment tree:

For IOM-1 or IOM5-e, expand Network→*NE*→Shelf→Card Slot.

For XIOM-s, expand Network→*NE*→Shelf→Card Slot→Xiom Card Slot.

2

Right-click on the IOM-1, IOM5-e, or XIOM-s and choose Properties. The Card Slot form opens.

3

Click on the Licenses tab. The Upgrade Path and Available Path tabs appear.

4

Click on the Upgrade Path tab and click Create. The Upgrade License Path (Create) form opens.

5

Click on Select beside the Assigned Level Upgrade ID parameters, and choose an ID.

6

Save your changes and close the forms.

END OF STEPS

Procedures for daughter card slot object configuration

15.78 To configure an MDA

15.78.1 Steps

1

On the equipment tree, right-click on a device and choose Properties. The Network Element form opens.

2

On the Network Element tree, expand the Shelf object→Card Slot object.

3

Perform one of the following:

a. To add a daughter card to an empty daughter card slot, right-click on the daughter card slot and choose Configure Daughter Card. The Daughter Card Slot form opens.

b. To modify an existing daughter card, click on the daughter card slot. The Daughter Card Slot (Edit) form opens.

Note: If a daughter card in your configuration is managed by NFM-P but becomes no longer supported by NFM-P (for instance, an obsoleted MDA), it will be displayed under the equipment tree as type Unspecified. To delete such a card from NFM-P, you must remove the MDA from the device. Once a card is removed from the device, it is also automatically deleted from NFM-P.

4

Configure the required parameters.

The Clock Mode parameter value must be set to Differential in order to configure the Differential Timestamp Frequency (kHz) parameter.

The Assigned Daughter Card Type and In MDA Carrier Module Slot parameters are configurable only during daughter-card creation.

Note: The ATM Mode parameter must be set to a value of max16k-vc for PPPoA and PPPoEoA to be supported.

5

Click on the Daughter Card tab and configure the required parameters.

Note: Before you configure the MDA Mode parameter, you must shut down all of the ports and remove all of the subchannels and bundles from the MDA.

6

To configure named buffer pools:

Note: You must enable pool mode on the parent I/O card to enable named buffer pool support. See 15.67 “To enable named pool mode” (p. 526) for more information.

1. Select a policy in the Ingress Pool Policy panel.
2. If you are configuring an HSMDA daughter card, go to Step 10 .

   Note:

   HSMDA daughter cards support only ingress named buffer pool policies, and not egress named buffer pool policies.
3. Select a policy in the Egress Pool Policy panel.

7

If required, select a policy in the Network Queue Policy panel.

For 7450 ESS, 7750 SR, and 7950 XRS NEs, network queue policies are assigned to forwarding planes; see 15.58 “To assign a network queue policy to a forwarding plane” (p. 518).

8

If required, select a Network Ingress Fabric Profile and an Access Ingress Fabric Profile in the Fabric Profiles panel.

9

If required, select a Network Ingress Security Queue and an Access Ingress Security Queue in the Security Queue Policies panel.

10

If required, select a policy in the Egress HSMDA Pool Policy panel.

Note: You must enable pool mode on the parent I/O card to enable HSMDA pool support. See 15.67 “To enable named pool mode” (p. 526) for more information.

11

Configure the required parameters in the HSMDA Aggregate Queue Burst panel.

12

To assign a network queue policy and network policy of ring type to a 7705 SAR:

1. Select a network queue policy (also called an add-drop queue policy) in the Ring Policies panel.
2. Select a network policy in the Network Policy panel.

   Note:

To assign a network policy of ring type to a 7705 SAR-8 or 7705 SAR-18 MDA, you must set the Network Policy Type for 7705 parameter to Ring when you configure a network policy on the MDA. See 50.41 “To configure a QoS network policy” (p. 1568) for more information.

13

Select an Access Ingress Shaper policy in the Shaper Policies panel.

Note: The Shaper policy is supported on all Ethernet MDAs on the 7705 SAR-8 and 7705 SAR-18 except the following:

- 2 × 10-Gig Bridged Ethernet + 1 × 2.5G Virtual Ethernet MDA
- 6 × 10/100 Ethernet + 2 × 10/100/1000 Ethernet SFP v2
  The Shaper policy is supported on all integrated MDAs on all 7705 SAR variants except the following:
- 7705 SAR-M — 2 × 10-Gig Bridged Ethernet XFP + 1 × 2.5G Virtual Ethernet
- 7705 SAR-H — 4 × 10/100 MDA

14

To configure MAC management:

1. Click on the MAC Operation tab and configure the required parameters.
2. To create an FDB Mac entry, click on the FDB Entries tab. The MDA MAC FDB Management form opens.
3. Configure the required parameters.
4. Select a port in the Port panel.

   Note:

   You can only bind a ring port to an FDB Mac entry.
5. Click OK.

15

To configure the cellular MDA on 7705 SAR-Hm, click on the Cellular tab and configure the required parameters.

When the 7705 SAR-Hm is operating in dual SIM mode, you can configure the SIM switchover behavior to be automatic based on specified failover criteria. When you set the Active SIM Card parameter to Automatic, you can configure the Down Recovery Interval (Min) and Down Recovery Criteria parameters for each SIM card. These parameters specify failover criteria to determine when to automatically switch to the backup SIM.

If you set the Down Recovery Interval (Min) parameter to be less than the Fail Duration (Min) value set on the cellular port, a node reboot happens instead of a SIM switchover.

16

If required, click on the Wlan Radio tab and configure the parameters.

17

If required, click on the Network Interfaces tab to add a network interface card.

18

Click Add. The Create Network Interface - Routing Instance form opens. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information on creating a network interface.

19

If required, perform 15.79 “To configure IMPM on an MDA” (p. 539) to configure IMPM on the daughter card.

20

For 7450 ESS, 7750 SR, and 7950 XRS NEs, network ingress pools are configured on forwarding planes; see 15.60 “To configure a network ingress pool on a forwarding plane” (p. 520).

21

To configure ATM:

1. Click on the ATM tab, select an ATM configuration, and click Properties. The ATM Configuration form opens.
2. Click on the VP Shaping tab and click Create. The ATM VP Shaper form opens.
3. Configure the VP ID parameter.
4. Select an ATM QoS Policy in the ATM QoS Policy panel.
5. Click on the QoS tab and configure the VP Shaping Weight parameter.
6. Save your changes and close the forms.
7. If required, repeat 1 to 6 to create additional ATM VP shapers.

22

Save your changes and close the forms.

END OF STEPS

15.79 To configure IMPM on an MDA

15.79.1 Purpose

Perform this procedure to configure IMPM on an MDA that is not installed in an XCM, 2 x XP MDA IOM 3 or an IMM.

Note: See 15.64 “To configure IMPM on an XCM, 2 x XP MDA IOM 3, or IMM forwarding plane” (p. 522) for information about configuring IMPM on an XCM, 2 × XP MDA IOM 3 or IMM.

15.79.2 Steps

1

Perform Step 1 to Step 6 of 15.78 “To configure an MDA” (p. 536) .

2

Configure the required parameters.

3

Select a policy in the BW policy panel.

Note: The multicast bandwidth policy named “default” cannot be modified or deleted.

4

Click OK and confirm your actions. The form closes.

END OF STEPS

15.80 To upgrade a license path on an FP4 MDA

15.80.1 Before you begin

Before you perform a license level upgrade, the NE bof should be pointed to the latest license file, and then you execute admin system license validate and admin system license activate.

15.80.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot→Daughter Card Slot.

2

Right-click on the MDA and choose Properties. The Card Slot form opens.

3

Click on the Licenses tab. The Upgrade Path and Available Path tabs appear.

4

Click on the Upgrade Path tab and click Create. The Upgrade License Path (Create) form opens.

5

Click on Select beside the Assigned Level Upgrade ID parameters, and choose an ID.

6

Save your changes and close the forms.

Note: If a higher upgrade path is deleted, then the Licensed Level reverts to the previous level.

END OF STEPS

15.81 To specify an event action for a 7x50 MDA

15.81.1 Purpose

Perform this procedure to specify the action to happen on the MDA when a specified event occurs. Supported event actions send a trap that is received by the NFM-P and will be raised as an alarm in NFM-P.

15.81.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot→Daughter Card Slot.

2

Right-click on the MDA and choose Properties. The Card Slot form opens.

3

Click on the Events tab and click Create. The Events (Create) form opens.

4

Configure the event type and the associated event action.

5

Save your changes and close the form.

END OF STEPS

15.82 To view the operational multicast channel properties of an MDA

15.82.1 Purpose

Perform this procedure to view specific information about an operational multicast channel on an MDA that is configured to use an Ingress Multicast Bandwidth policy.

For 7450 ESS, 7750 SR, and 7950 XRS NEs, the operational MC path management properties are available for viewing on forwarding planes; see 15.61 “To view the operational MC path management properties of a forwarding plane” (p. 520).

Note: You can also view operational multicast channel information in the following locations:

- the Mcast Path Mgmt tab on the properties form of a VPLS or VPRN service site
- the Mcast Path Mgmt Channels tab on the properties form of a routing instance.

Each properties form contains information about the ingress MDA of each operational channel.

15.82.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n.

3

Click the Daughter Card Slot object. The Daughter Card Slot form opens.

4

Click on the Mcast Path Mgmt Channels tab.

5

Search for channels, choose a channel from the list and click Properties.

6

View the channel information, which includes the service site or routing instance associated with the channel, and the following properties:

Note: Channel information is displayed only when the MDA receives multicast traffic for a previously configured multicast group.

- Group Address—the operational channel multicast group address
- Source Address—the operational channel multicast source address
- Bandwidth—the operational bandwidth of the channel
- Administrative Bandwidth—the administrative bandwidth of the channel
- Last Highest Bandwidth—the value of the multicast bandwidth that is currently allocated to the channel forwarding path on the MDA. This value is calculated based on periodic statistics polls and represents the highest recorded bandwidth value since the most recent restart of the bandwidth update timer for the channel.
- Second Highest Bandwidth—the second-highest recorded bandwidth value for the channel since the most recent restart of the bandwidth update timer. The value is calculated based on periodic statistics polls and is reset to zero every time the bandwidth update timer for the channel is restarted.

- BW Update Timer expiration—the time that remains before the bandwidth update timer for the channel expires
- Current Path—the path that the channel traffic uses to reach the switching fabric; the path value is primary, secondary, or ancillary
- Explicit Path—indicates whether the current path is explicitly specified by a Multicast Info Policy
- Preference Level—the preference level of the channel
- Black-Hole—whether the channel is in the Black-Hole state
- Black-Hole Rate—specifies the bandwidth rate, in kb/s, at which the channel enters the Black-Hole state

7

Close the form.

END OF STEPS

15.83 To configure a module card on a 7705 SAR-M/ME or 7705 SAR-H

15.83.1 Steps

1

On the equipment tree, right-click on a 7705 SAR-M/ME or 7705 SAR-H NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot *n*, right-click on the applicable Daughter Card Slot object and choose Configure Daughter Card. The Daughter Card Slot (Create) form opens.

Note: For the 7705 SAR-M and 7705 SAR-ME, only the third daughter card slot can be configured with module cards. The first two daughter card slots are configured with integrated cards.
For the 7705 SAR-H, only the second and third daughter card slots can be configured with module cards. The first daughter card slot is configured with an integrated card.

3

Configure the Assigned Daughter Card Type parameter.

4

Click Apply. The Daughter Card Slot form refreshes.

5

Click on the Daughter Card tab and configure the required parameters.

6

To add a fabric profile:

1. Select a fabric profile in the Network Ingress Fabric Profile panel.
2. Select a fabric profile in the Access Ingress Fabric Profile panel.

   If you are configuring a 1 Colour Optical Add/Drop Mux card, go to Step 21 .

7

Select a network queue policy in the Network Policy Queue panel, if required.

8

Click on the Ports tab.

9

Click on a port and choose Properties. The Physical Port (Edit) form opens.

Note: The first port on the DCM includes four SHDSL lines. The second port includes two XDSL lines.

10

Configure the required parameters.

11

Click on the States tab and configure the Administrative State parameter. If you are configuring a DSL or DCM module, go to Step 13 .

12

Click on the GPON tab and configure the required parameters. Go to Step 21 .

13

Click on the DSL tab and configure the required parameters.

14

To configure XDSL or SHDSL lines:

1. Click on the XDSL or SHDSL tab, as required.
2. Choose an XDSL or SHDSL line and click Properties. The XDSL Line (Edit) or SHDSL Line (Edit) form opens.
3. Configure the Administrative State parameter.
4. Save the configuration and close the form.
5. To configure more lines, repeat 1 to 4 .

15

Click on the Ethernet tab and configure the required parameters on the General tab.

16

Click on the EFM-OAM tab. To configure the parameters on this tab, perform the steps provided in 90.49 “To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form” (p. 3069) .

17

Click on the LLDP tab and configure the required parameters on the Nearest Bridge, Nearest Customer, and Nearest Non TPMR tabs.

18

Click on the Remote Peers tab under the LLDP tab to search for and display LLDP remote peers associated with the port. These remote peers are used to determine the physical topology of the network.

19

Click on the 802.1x Port and 802.1x Port Authenticator tabs, and configure the required parameters.

20

If you are configuring a 6 port Ethernet module card, click on the PoE tab and configure the required parameters.

21

Save your changes and close the forms.

END OF STEPS

15.84 To configure GNSS receiver functions on a 7705 SAR-Hm

15.84.1 Purpose

Perform this procedure to enable NMEA data streaming, view GNSS location information, and view satellite information for 7705 SAR-Hm, Release 16.0 R4 and later and 7705 SAR-Hmc, Release 20.5 R1 and later.

15.84.2 Steps

1

On the equipment tree, right-click on the NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot 1.

3

Click the Daughter Card Slot object and click Properties. The Daughter Card Slot form opens.

4

Click on the GNSS tab and configure the required parameters.

5

To enable NMEA, set the Sentence Streaming parameter to In Service and configure the Sentences and Sentence Interval (seconds) parameters, as required.

6

The form also displays GNSS location information.

7

Click on the Visible Satellites tab to display satellite information.

8

Close the form.

END OF STEPS

15.85 To perform a manual SIM switchover on a 7705 SAR-Hm

15.85.1 Purpose

Perform this procedure to manually switch from one SIM card to the other on a 7705 SAR-Hm in dual SIM mode.

15.85.2 Steps

1

On the equipment tree, right-click on a 7705 SAR-Hm and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot 1.

3

Click the Daughter Card Slot object and click Properties. The Daughter Card Slot form opens.

4

Click on the Cellular tab.

5

On the Cellular tab, configure the Active SIM Card to the other SIM

6

Click OK in the confirmation dialog.

The 7705 SAR-Hm changes cellular connection to the other carrier.

7

Close the form.

END OF STEPS

Procedures for bundle configuration

15.86 To create an FR group bundle

15.86.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n, right-click on the Bundles object and choose Create Bundle. The Create Multilink Bundle window opens.

3

Configure the general property parameters and click Next. The Configure Bundle Parameters form opens.

4

Set the Bundle Type parameter to FR and configure the rest of the required parameters, then click Next. The Configure Bundle Members form opens.

Note: To create a PPP group bundle, perform 15.90 “To create an MLPPP bundle” (p. 551) .
To create an IMA group bundle, perform 15.88 “To create an IMA group bundle” (p. 550)

5

Perform one of the following:

a. To create the multilink bundle without adding any members, go to Step 6 .

b. To add DS0 channel groups to the bundle:

1. Click Create. The Add Bundle Member form opens.
2. Configure the Show Only Compatible Channels parameter, as required.
3. Click Next. The Select Channels configuration form opens.
4. Search for channels, and then choose up to 8 channel groups from the list of channels to construct the bundle.

   Note:

   The channel group with the lowest Port ID is chosen as the primary member for the bundle. The Encap Type of the primary member is used for all other members. When adding members to a bundle at a later time, choose members with the same Encap Type as the primary member.
5. Click Finish. The form closes and the Create Multilink Bundle form reappears.

6

Click Finish to close the form.

END OF STEPS

15.87 To modify an FR group bundle

15.87.1 Purpose

Perform the following procedure to modify an FR group bundle already configured on a channelized MDA in the NE.

15.87.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n→Bundles, right-click on an FR group bundle object and choose Properties. The Multilink Bundle (Edit) form opens.

3

Configure the required parameters on the following tabs:

- General
- Multilink Bundle
- States
- MLFR
- Bundle Members (where you can create and add new bundle members)

4

To configure the FR interface:

1. Click Edit FR. The FR Interface Bundle form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Save your changes and close the form.

END OF STEPS

15.88 To create an IMA group bundle

15.88.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n, right-click on the Bundles object and choose Create Bundle. The Create Multilink Bundle form opens.

3

Configure the general property parameters and click Next. The Configure Bundle Parameters form opens.

4

Set the Bundle Type parameter to IMA and configure the rest of the required parameters, then click Next. The Configure Bundle Members form opens.

Note: To create a PPP group bundle, perform 15.90 “To create an MLPPP bundle” (p. 551) .
To create an FR group bundle, perform 15.86 “To create an FR group bundle” (p. 548) .

5

Perform one of the following:

a. To create the multilink bundle without adding any members, go to Step 6 .

b. To add DS0 channel groups to the bundle:

1. Click Create. The Add Bundle Member form opens.
2. Configure the Show Only Compatible Channels parameter, as required.
3. Click Next. The Select Channels configuration form opens.
4. Search for channels, and then choose up to 8 channel groups from the list of channels to construct the bundle.

   Note:

   The channel group with the lowest Port ID is chosen as the primary member for the bundle. The Encap Type of the primary member is used for all other members. When adding members to a bundle at a later time, choose members with the same Encap Type as the primary member.
5. Click Finish. The form closes and the Create Multilink Bundle form reappears.

6

Click Finish to close the form.

END OF STEPS

15.89 To modify an IMA group bundle

15.89.1 Purpose

Perform the following procedure to modify an IMA group bundle already configured on a channelized MDA in the NE.

15.89.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n→Bundles, right-click on an IMA group bundle object and choose Properties. The Multilink Bundle (Edit) form opens.

3

Configure the required parameters on the following form tabs:

- General
- Multilink Bundle
- States
- Bundle Members (where you can create and add new bundle members)

4

Save your changes and close the form.

END OF STEPS

15.90 To create an MLPPP bundle

15.90.1 Steps

1

On the equipment tree, right-click on the NE on which you need to create an MLPP group bundle and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n, right-click on the Bundles object and choose Create Bundle. The Create Multilink Bundle form opens.

3

Configure the general property parameters and click Next. The Configure Bundle Parameters form opens.

4

Set the Bundle Type parameter to PPP and configure the rest of the required parameters, then click Next. The Configure Bundle Members form opens.

Note: To create an FR group bundle, perform 15.86 “To create an FR group bundle” (p. 548)
To create an IMA group bundle, perform 15.88 “To create an IMA group bundle” (p. 550) .
If you create an MLPPP bundle and the DS1 BER exceeds a threshold, the PPP link is automatically removed from the MLPPP bundle.

5

Perform one of the following:

a. To create the multilink bundle without adding any members, go to Step 6 .

Note: If you need to configure the MLPPP group bundle for multiclass service transmission, do not add any bundle members. See 15.92 “To configure an MLPPP bundle for multiclass service transmission” (p. 554) for information about how to configure the MLPPP bundle for multiclass service transmission.

b. To add DS0 channel groups to the bundle:

1. Click Create. The Add Bundle Member form opens.
2. Configure the Show Only Compatible Channels parameter, as required.
3. Click Next. The Select Channels configuration form opens.
4. Search for channels, and then choose up to 8 compatible channel groups from the list of channels to construct the bundle.

   Note:

   If there are no compatible channels to choose from, edit some of the existing channels that are not compatible to make them compatible:
   - Click Back and disable the Show Only Compatible Channels parameter. Click Next.
   - Choose channels to edit from the list and click Properties.

   The channel group with the lowest Port ID is chosen as the primary member for the bundle. The Encap Type of the primary member is used for all other members. When adding members to a bundle at a later time, choose members with the same Encap Type as the primary member.
5. Click Finish. The form closes and the Create Multilink Bundle form reappears.

6

Save your changes and close the form.

END OF STEPS

15.91 To modify an MLPPP bundle

15.91.1 Purpose

Perform the following procedure to modify an MLPPP group bundle already configured on a channelized MDA in the NE.

15.91.2 Steps

1

In the Equipment view of the navigation tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n→Bundles.

3

Right-click on an MLPPP group bundle object and choose Properties. The Multilink Bundle (Edit) form opens.

4

Configure the required parameters on the following form tabs:

- General
- Multilink Bundle
- States
- MLPPP
- Bundle Members (where you can create and add new bundle members)

  Note:

  If you need to configure the MLPPP group bundle for multiclass service transmission, do not add any bundle members. See 15.92 “To configure an MLPPP bundle for multiclass service transmission” (p. 554) for information about how to configure the MLPPP bundle for multiclass service transmission.

5

Save your changes and close the form.

END OF STEPS

15.92 To configure an MLPPP bundle for multiclass service transmission

15.92.1 Purpose

Perform the following procedure to configure an existing MLPPP group bundle for multiclass service transmission.

Consider the following before you start to configure an MLPPP bundle for multiclass service transmission.

- Only channelized ASAP MDAs on the 7705 SAR and 7750 SR support multiclass MLPPP.
- You need to configure multiclass MLPPP before you add bundle members to the multilink bundle.
- Ports configured in network mode do not support multiclass MLPPP.
- Multiclass MLPPP is not supported when LFI is enabled.

15.92.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n→Bundles.

3

Right-click on an MLPPP group bundle object and choose Properties. The Multilink Bundle (Edit) form opens.

4

Click on the MLPPP tab and configure the required parameters.

Note: The Magic Number parameter is configurable only on channelized ASAP MDAs on 7750 SR NEs.
You can apply QoS profiles to an MLPPP bundle only if the Class Count parameter is set to 4.

5

Save your changes and close the form.

END OF STEPS

15.93 To configure an MLPPP bundle as a network interface on a channelized ASAP MDA

15.93.1 Purpose

Perform the following procedure if you need to configure an existing MLPPP group bundle as a network interface; see 15.90 “To create an MLPPP bundle” (p. 551) for details.

Consider the following when configuring an MLPPP as a network interface.

- Only channelized ASAP MDAs on the 7705 SAR and 7750 SR support MLPPP group bundles.
- The MLPPP network interfaces do not support LFI.
- The MLPPP network interfaces do not support MC MLPPP.
- IPv6 network interfaces are not supported.

15.93.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

On the NE navigation tree, expand Shelf→Card Slot n→Daughter Card Slot n→Bundles.

3

Right-click on an MLPPP group bundle object and choose Properties. The Multilink Bundle (Edit) form opens.

4

Click on the Bundle Members tab and configure the port mode:

1. Search for bundle members, choose a member from the list and click Properties. The Bundle Member (Edit) form opens.
2. Click on the Port tab.
3. Click Properties beside the port name. The DS0 Channel Group (Edit) form opens.
4. Set the Mode parameter to Network.
5. Save your changes and close the form.

5

Save your changes and close the form. The form refreshes with new tabs.

6

Click on the Network Interfaces tab and configure a network interface:

1. Click Create. The Create Network Interface - Routing Instance form opens.

2. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information about creating a network interface on the DS0 channel.

   Note:

   The port in Step 5 of 27.17 “To create an L3 network interface on a routing instance” (p. 856) is set to the DS0 channel by default. Go to Step 8 of 27.17 “To create an L3 network interface on a routing instance” (p. 856) .

   The L3 interface appears on the Network Interfaces tab.

7

Save your changes and close the form.

END OF STEPS

Procedures for extension shelf configuration

15.94 To create a satellite shelf

15.94.1 Steps

1

On the equipment tree, right-click on the NE where you need to create a satellite shelf and choose Create Satellite.

2

Configure the required parameters.

Configure the Dynamic Uplink and Uplink Distribution parameters to dynamically enable satellite port mapping.

3

Save your changes and close the form.

END OF STEPS

15.95 To configure satellite shelf uplink port topology

15.95.1 Overview

Use this procedure to configure uplink port topology between the connection ports on an NE and the uplink ports on a satellite shelf.

15.95.2 Steps

1

Perform one of the following:

a. To configure uplink port topology from the NFM-P main menu:
   1. From the NFM-P main menu, select Create→Equipment→Satellite Port Link.

b. To configure uplink port topology from the NE level:
   1. On the equipment tree, right-click on the NE and choose Properties.
   2. Click on the Physical Links tab, and then on the Satellite Port Topology tab. Click Create.

c. To configure uplink port topology from the port level:
   1. On the equipment tree, click on an uplink port on the satellite shelf.
   2. Hold the Ctrl key and click on the port on the NE to which the satellite shelf is connected.
   3. Right-click on one of the ports and select Create Satellite Port Link.
   4. Go to Step 4.

2

In the End Pointer A Port panel, click on the Select button and choose an uplink port on the satellite shelf.

3

In the End Pointer B Port panel, click on the Select button and choose the port on the NE to which the satellite shelf is connected.

4

Save your changes and close the forms.

END OF STEPS

15.96 To configure flexible satellite port mapping for Ethernet satellites

15.96.1 Overview

Use this procedure to configure any number of primary uplinks to be used between the 7x50 host, Release 15.0 R3 or later, and the Ethernet satellites to achieve uplink resiliency. You can also configure any number of secondary uplinks to be used between the 7x50 host, Release 15.0 R8 or later, and the Ethernet satellites to achieve uplink resiliency.

Note: For node versions prior to Release 16.0 R4, primary and secondary uplink ports of satellite access ports cannot be mapped to host ports of same MDA/Card in 7x50 nodes.

Note: The administrator must ensure that the new uplink port has sufficient resources to handle the re-assigned satellite ports and associated services.

Note: The satellite port mapping requires a manual refresh to get the updated information on the host port.

15.96.2 Steps

1

On the equipment tree, right-click on the Ethernet satellite and choose Properties.

2

Click on the Flexible Satellite Port Mapping tab, and then on the satellite port. Click Properties. The Flexible Satellite Port Mapping form opens.

3

Click Select beside the Primary Up Link Port - Port Name parameter and select a port from the list and click OK.

Note: The uplink association can only be changed if there are no services configured against any of the satellite client ports.

4

Click Select beside the Secondary Up Link Port - Port Name parameter and select a port from the list and click OK.

5

Repeat to configure multiple ports and then click Apply on the Flexible Satellite Port Mapping form.

6

To revert to the primary uplink click Clear beside the Secondary Up Link Port - Port Name parameter and click OK.

7

Save your changes and close the forms.

END OF STEPS

15.97 To configure tunable optics for Ethernet satellites

15.97.1 Before you begin

You must create a 64-port 10-Gig SFPP Ethernet Satellite with 4 100-Gig QSFP28 up-links or 64-port 10-Gig SFPP Ethernet Satellite with 4 100-Gig CFP up-links on the 7750 SR, 7450 ESS, and 7950 XRS, Release 16.0 R4 and later.

The tabs and parameters that display vary depending on the type of Ethernet Satellite, and the NE and release.

15.97.2 Steps

1

On the equipment tree, right-click on the esat port of the Ethernet satellite and choose Properties.

2

Perform one of the following:

a. Click Select beside the DWDM Channel parameter and choose the available channel. The other DWDM parameters are populated.

b. Click the DWDM Optics tab and configure the required parameters.

   You cannot configure the center frequency if the DWDM channel parameter is configured.

3

Save your changes and close the form.

END OF STEPS

15.98 To associate a port template on an Ethernet satellite shelf

15.98.1 Before you begin

You must create a port template before you can associate the port template on a satellite shelf. See 12.43 “To create a port template” (p. 375).

15.98.2 Steps

1

On the equipment tree, navigate to the Ethernet satellite shelf and choose Properties. The Satellite Configure Shelf form opens.

2

Set the Assigned Type parameter to a satellite that contains a port template.

3

Click Select beside Template Name and then choose a port template.

4

Save your changes and close the forms.

END OF STEPS

15.99 To configure local forward on Ethernet satellite shelf

15.99.1 Steps

1

On the equipment tree, navigate to the Ethernet satellite shelf and choose Properties. The Satellite Configure Shelf form opens.

2

Enable the Local Forward parameter.

3

On the equipment tree, right-click on the NE with the Ethernet satellite shelf and choose Properties.

4

Click on the Local Forwards tab and click Create. The Satellite Local Forward Create form opens.

5

Configure the required parameters on the General tab.

6

Click on the SAP tab and click Create. The Satellite Local Forward Sap Create form opens.

7

Configure the required parameters. A maximum of two SAP entries can be created.

8

Save your changes and close the forms.

END OF STEPS

15.100 To enable transparent clock functionality on an Ethernet satellite

15.100.1 Before you begin

You must create one of the following satellites on a 7450 ESS, 7750 SR or 7950 XRS, Release 16.0 R5 and later:

- 48-port 1-Gig Ethernet SFP satellite
- 64-port 10-Gig SFPP Ethernet satellite with 4 100-Gig CFP up-links
- 48-port 1-Gig Ethernet SFP satellite, LCS variant
- 64-port 10-Gig SFPP Ethernet satellite with 4 100-Gig QSFP28 up-links

See 15.94 “To create a satellite shelf” (p. 557).

15.100.2 Steps

1

On the equipment tree, navigate to the Ethernet satellite shelf and choose Properties. The Satellite Configure Shelf form opens.

2

Enable the Synchronous Ethernet and Ptp Transparent Clock parameters under Shelf Details, and the Ptp Transparent Clock parameter under Feature.

3

On the equipment tree, right-click on the NE with the Ethernet satellite shelf and choose Properties.

4

Click on the Clock item under the IEEE PTP and Synchronization items in the navigation tree.

5

Set the PTP Profile parameter on the General tab to IEEE1588-2008.

6

Associate the IEEE PTP ports. See 15.23 “To configure IEEE 1588 PTP ports on a 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, or 7950 XRS” (p. 481).

7

Save your changes and close the forms.

END OF STEPS

15.101 To perform a software upgrade on an extension shelf

15.101.1 Steps

1

If required, create a software repository policy for the upgrade files.

1. Choose Administration→NE Maintenance→Software Repository from the main menu.
2. Click on the Create button.
3. Configure the Primary Location parameter, and any other required parameters.
4. Save and close the form.

2

In the equipment tree, navigate to the host node for the extension shelf you need to upgrade.

3

Right-click on the node and choose Properties.

4

In the Software Repository panel, click on the Select button and select the software repository policy for the upgrade files.

5

Save your changes and close the form.

6

In the equipment tree, navigate to the extension shelf you need to upgrade.

7

Right-click on the shelf, and select Reboot Upgrade to perform the upgrade.

END OF STEPS

Procedures for power shelf configuration

15.102 To configure power shelves

15.102.1 Steps

1

On the equipment tree, navigate to the power shelf and choose Configure Power Shelf. The Power Shelf form opens.

2

Configure the required parameters.

3

Save your changes and close the form.

END OF STEPS

15.103 To configure power modules

15.103.1 Steps

1

On the equipment tree, navigate to the power module and choose Configure Power Module. The Power Module form opens.

2

Configure the required parameters.

3

Save your changes and close the form.

END OF STEPS

16 Port and channel object configuration

Configuring port and channel objects

16.1 Overview

16.1.1 Port and channel objects

Port objects are children of daughter card slot objects. They appear below the daughter card slot after the daughter card is configured. Channel objects are children of port objects. They appear on the navigation tree below the port object after the channel is configured.

Properties forms for port objects and channel objects are accessed using the NFM-P navigation tree.

This chapter contains the procedures to configure devices using the navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the navigation tree.

16.1.2 Working with port and channel objects

The types of ports that are available depend on the daughter cards that are configured in the chassis. For some OmniSwitch devices, the ports that are available depend on the chassis type, for others the type of card that is configured for the card slot. Ethernet ports cannot be channelized. SONET/SDH and TDM ports can be channelized.

The port syntax for most devices that support daughter cards is card slot/daughtercard/port. For example, Port 1/1/1 represents port 1 of daughter card 1 in slot 1. Connector ports are indicated in the equipment tree by the letter c preceding the port number; for example, 1/1/c1/1. The port syntax for the OmniSwitch is card slot/port.

In most cases, ports are created automatically when the daughter card is created. You must select one port object at a time and configure the properties of the port for the service that you need the port to provide. The properties vary depending on the port type. Channel objects are created on SONET/SDH or TDM ports for any type of channelization on the port whether it is a clear channel application or a sub-channel application.

Use the properties forms from the contextual menus in the navigation tree to configure port and channel parameters. You can configure the port mode as network, access, or hybrid.

- Network ports pass network-level traffic.
- Access ports are customer-facing and pass service-level traffic.
- Hybrid ports can pass network-level and access-level traffic.

Note: By default, a port is in network mode.
See 16.29 “To change the port mode” (p. 611) to change the port mode.

Clear channel ports (OC-192/OC-48c/OC-12c/OC-3c) are network or access ports. Channelized ports (CHOC-12/CHOC-3c/ DS3/E3) are always in access mode, however, the ASAP CHOC-3 supports both access and network modes. The ASAP DS1/E1 ports support:

- access mode for ATM/IMA and TDM

- network mode for PPP/MLPPP

Network ports are used in the service provider transport or infrastructure network, such as an IP/MPLS-enabled backbone network or uplink ports for rings using L2 Ethernet switches. A port that is in network mode can be assigned an IP address and host an L3 interface that can perform IP routing functions.

Access ports are associated with a SAP, a subscriber, and a service to provide connectivity. Access ports and channels are configured with encapsulation values to differentiate the service on the port or channel. When a port is access mode, one or more services can be configured on the port. A channelized port that is to act as an endpoint must be in access mode. You can convert access ports to hybrid ports and migrate any existing SAPs. See 16.30 “To migrate SAPs from access mode to hybrid mode” (p. 612) for more information.

Hybrid ports are ports that can host SAPs and network interfaces simultaneously. This enables a customer to use one uplink port for enterprise and end-user traffic. A hybrid port preserves the existing network and access port functionality, and uses the same QoS, scheduler, and port scheduler resources as other ports, but requires the configuration of weight values that allocate buffer percentages to the access and network traffic on the port.

The NFM-P supports hybrid mode configuration on Ethernet ports. Support for hybrid ports varies depending on the NE chassis, card, and release; see the NE documentation for information.

Note: Single-slot chassis such as the 7450 ESS-1 and 7750 SR-1 do not support hybrid port mode.

You can configure a hybrid port on an Ethernet MDA in an IOM-3XP, or on a CMA in an IOM-XP. The 7750 SR supports hybrid port configuration on an IMM in addition to an IOM-3XP.

Note: Hybrid port configuration is not supported on an HSMDA or a VSM MDA.

A hybrid port supports dot1q and QinQ encapsulation, but does not support null encapsulation, in order to accommodate single-SAP operation. The available VLAN tags are shared among the VLAN SAPs and VLAN network IP interfaces. When you create a SAP or L3 interface on a hybrid port, the outer VLAN tag must not be in use by another SAP or L3 interface on the port.

By default, the MTU of a hybrid port is set to the larger of the network and access MTUs to accommodate the creation of L3 interfaces and SAPs.

Note: A hybrid port can participate in a single-chassis LAG.
A hybrid port cannot participate in an MC-LAG or MC ring.

See the appropriate device documentation for more information about hybrid ports.

When working with a TDM port, you must specify the Line Buildout as either short or long. That is, for a DS3 port the Line Buildout parameter must be configured. If the TDM port is in the context of a SONET STS-1 sub-channel, for example, the DS3 channel is built on the STS-1 channel of a SONET port, the line buildout parameter is not required.

At the connection termination points, you are required to configure the Encap Type as required, the MTU size as required, and the configured MAC address as required when configuring the port or channel.

Policies can be added or deleted as required using the manage policy forms.

You can associate policies to ingress and egress access and network ports. Buffer policies are used to create and edit QoS buffer pool resources on network ports, access ports, and access channels. Egress network ports, access ports, and access channels have a dedicated buffer pool for queuing. The traffic is handled by a single buffer pool, one at the ingress, and one at the egress.

You can configure the amount of egress buffer space to be allocated to the port or channel. By default, all egress buffers are allocated fairly among the egress ports and channels based on their relative egress bandwidth.

The egress buffers for egress network ports and channels are put into per-port or per-channel egress buffer pools and are used by the egress network forwarding class queues on that port or channel. The ingress buffers allocated to network ports and channels are summed into a single pool and are used by the ingress network forwarding class queues (defined by the network ingress buffer policy).

The egress and ingress buffers allocated to access ports and channels are put into an egress buffer pool and ingress buffer pool for the port or channel. The access buffer pools are used by egress and ingress service queues created by the SAP-egress and SAP-ingress policies in use by services on the port or channel.

Changing the size of an egress buffer pool should be carefully planned. By default, there are no free buffers to increase the size of a pool. In order to increase a pool on one port or channel, the same amount of buffers must be freed from other egress buffer pools on the same daughter card.

16.2 Digital diagnostics monitoring

16.2.1 Overview

The NFM-P displays the following digital diagnostics monitoring data and alarm and warning information for ports on CFPs, QSFPs, SFPs and XFPs optical modular transceivers:

- temperature
- supply voltage (SFP)
- TX bias
- TX output power
- RX received optical power
- external calibration

The transceiver is programmed with warning and alarm thresholds for low and high conditions that can generate system events. The thresholds for CFPs, QSFPs, SFPs, and XFPs are programmed by the transceiver manufacturer. The NFM-P raises an alarm when these thresholds are exceeded.

You can view digital diagnostics monitoring information from the DDM tab and the Lane DDM tab on the property form of supported ports. The External Calibration tab is available if the ports on CFPs, QSFPs, SFPs, and XFPs optical modular transceivers support the external calibration functionality.

16.2.2 1830 VWM − DDM data retrieval

The NFM-P supports DDM data retrieval from each optical SFP in the data and control plane. See 16.77 “To retrieve 1830 VWM DDM data” (p. 680) for more information about how to retrieve 1830 VWM DDM data.

The following DDM data is retrieved from the local 1830 VWM devices or from the 1830 VWM RMUs:

- DDM received power
- DDM transmitted power
- DDM laser bias current
- DDM temperature
- DDM voltage

16.3 Remote fiber link monitoring in 1830 VWM devices

16.3.1 Overview

Passive WDM devices in remote radio sites cannot communicate with 1830 VWM OSU devices. You can still monitor and detect signal failures using RFLM ports.

16.3.2 RFLM

Perform the following to monitor and detect a signal failure in a passive WDM device:

1. Transmit an RFLM wave, that is, a 1627 nm signal, from an optical OSC port of an 1830 VWM OSU device to a passive WDM device in the remote radio site.
2. Loop the signal at the remote end back to the 1830 VWM OSU device. An alarm is generated if a signal path failure is detected.

NFM-P supports the configuration of an 1830 VWM OSU optical OSC port to operate as an RFLM port. The Interface Role parameter of the 1830 VWM OSU optical OSC port participating in the remote fiber link monitoring is set as RFLM after setting the administrative state to down. You can associate a label with the signal by entering relevant information in the RFLM Label text field. When a signal failure occurs, the text appears in the generated alarm and helps in identifying the RFLM alarm.

See 16.78 “To configure an OSC port of an 1830 VWM OSU as a RFLM port” (p. 680) for more information about configuring the OSC port as an RFLM port.

16.4 Tagged and untagged VLAN ports

16.4.1 Overview

The NFM-P supports tagged and untagged ports as VLAN access ports. The following table describes the behavior when ingress tagged or untagged traffic enters and then exits a VLAN.

Table 16-1 Tagged and untagged traffic behavior

Ingress traffic configuration	Ingress VLAN port configuration	Action	Egress VLAN port configuration and action
Untagged	Untagged	The VLAN allows the traffic and passes it based on the default VLAN ID. The MAC address of the destination is learned.	If the egress VLAN port is untagged, the traffic remains untagged.
			If the egress VLAN port is tagged, a tag is added to the traffic.
	Tagged		If the egress VLAN port is untagged, the traffic remains untagged.
			If the egress VLAN port is tagged, a tag is added to the traffic.
Tagged	Untagged	The VLAN allows the traffic if the tag matches the default VLAN ID. The MAC address of the destination is learned. If the tag does not match the default VLAN ID, the traffic is dropped.	If the egress VLAN port is untagged, the tag of the traffic is removed.
			If the egress VLAN port is tagged, the traffic remains tagged.
	Tagged	The VLAN allows the traffic: • if the tag matches the default VLAN ID • if the tag matches another VLAN ID The MAC address of the destination is learned. If the tag does not meet either condition, the traffic is dropped.	If the egress VLAN port is untagged, the tag of the traffic is removed.
			If the egress VLAN port is tagged, the traffic remains tagged.

The following figures show how tagged and untagged traffic is handled on the devices, based on the VLAN ID.

Figure 16-1 Untagged traffic and VLANs

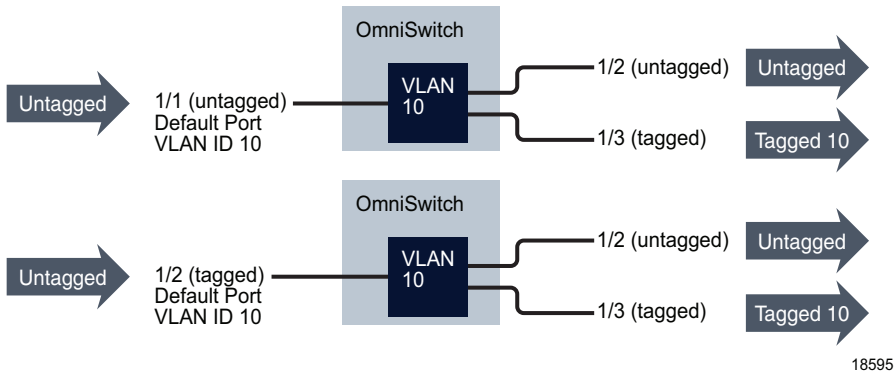

Figure 16-2 Tagged traffic and VLANs

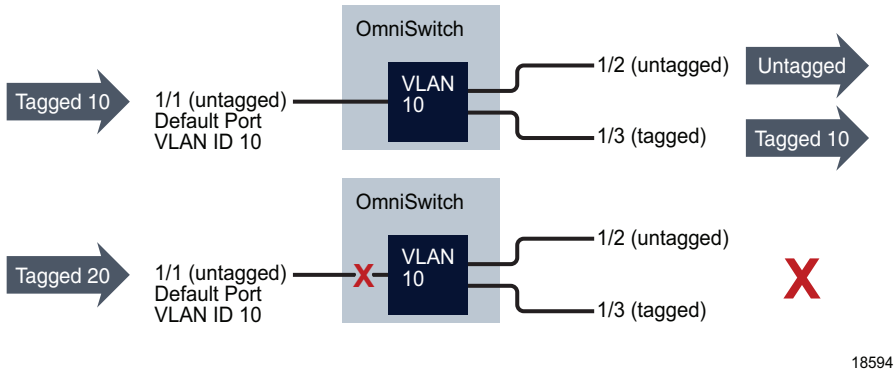

16.5 Connection termination points for services and interfaces

16.5.1 Overview

Connection termination points are objects that represent terminating endpoints for a service, for example the endpoint of a VLL service. Connection termination points can be Layer 2 or Layer 3 interfaces, depending on the type of service being created. At the connection termination points you must configure the mode as Access or Network, the Encap Type as required, the MTU size as required, and the configured MAC address as required when configuring the port or channel. The following objects can be used for connection termination points:

- STS-3 to STS-192 clear channels
- DS3 clear channel
- DS0 groups
- ports
- bundle

16.5.2 STS-3 to STS-192 clear channel

STS-3 to STS-192 clear channel SONET/SDH ports can be used to create SAPs or IP interfaces with one clear channel on each port that operates at the rate of the parent object. Clear channel SONET applications can be performed on any OC-*n* card. SONET channel termination for 1 × 10-Gig MDAs are not supported. See 16.9 “SONET clear channel applications” (p. 578) in this section for more information.

16.5.3 DS3 clear channel

A DS3 clear channel can be a connection termination point when it is explicitly configured as unchannelized, that is, when the Configuration Type is set to None, which is the default setting for a DS3. DS3 clear channel connections cannot be channelized to a lower level than the one full DS3 channel. See 16.10 “TDM channelization and clear channel applications” (p. 579) in this section for more information.

16.5.4 DS0 channel groups

A TDM channel group connects a group of DS0s by allocating a specific number of spans or interfaces and channels to a group of channels. The DS0 channel group defines the incoming and outgoing parameters for a group of channels such as IP profiles, routing tables, and translation tables to be assigned during the configuration of the specific channel group.

To use a DS1 or E1, you must create at least one DS0 group for the DS1 or E1. The NFM-P supports the automatic configuration of sub-channels and assignment of timeslots on the ports within DS0 groups. Depending on the TDM port selected, the NFM-P automatically creates the DS0 channel groups with the appropriate type of timeslots. For example, you can assign timeslots to a DS0 channel group as follows:

- For DS1 channel types: Use four timeslots per channel group, which results in six DS0 channel groups with IDs from 1 to 6. Each contains four timeslots.
- For E1 channel types: Use five timeslots per channel group, which results in six DS0 channel groups with IDs from 1 to 6. Each contains five timeslots, with one timeslot not assigned.

You can view the channels that are assigned in DS0 groups using the associate properties form. You can also disable, enable, or reassign the assignment of timeslots as required, in a DS0 group.

See 16.10 “TDM channelization and clear channel applications” (p. 579) in this section for more information about DS0 channel groups. See 16.61 “To perform a bulk channel creation on ports that support multiple sub-channels” (p. 654) for information about creating channels on ports for card types that support multiple sub-channels. See 16.66 “To configure TDM DS1 or E1 channels” (p. 663) for information about creating TDM DS1 or E1 channels and DS0 channel groups.

SONET STS-1 sub-channels

Only the DS0 group level can be used as a connection termination point for SONET STS-1 sub-channels. Channelization on the 1 × OC12 can be used to create up to 12 SONET STS-1 sub-channels. Each STS-1 channel can be used to create a DS3 frame on which you can build DS1 or E1 channels that can be configured to the DS0 channel group level. You can configure the DS0s of a DS0 group in any sequence and you do not need to use all DSOs. For example, you can use DS0 1, 3, 5, and 9. See 16.13.3 “SONET VT1.5 and VT2 payloads” (p. 587) in this section for more information.

Only the DS0 group level can be used as an endpoint on the channelized 12 × DS3 card. Channelization can be used on each DS3 port of the card to create independent TDM channels in the form of DS1 or E1 data channels that handle DS0 groups. The DS0s of a DS0 group can be configured in any sequence and you do not need to use all DS0s. For example, you can use DS0 1, 3, 5, and 9. See “SONET and SDH sub-channel applications and structure” (p. 586) in this section for more information.

16.5.5 Ethernet ports

Ethernet ports can be configured as connection termination points in SAPs and IP interfaces. They cannot be channelized.

You must configure the class of port, such as fast Ethernet, GigE, or 10G Ethernet. You must also configure the port encapsulation at the connection termination point. Ethernet access ports use:

- dot1q—supports multiple services on the port; the outer encapsulation value that distinguishes services is the VLAN ID in the IEEE 802.1Q header

- QinQ—supports multiple services on the port/channel; the inner and outer encapsulation values that distinguish services is the VLAN ID in the IEEE 802.1Q header
- null—supports a single service on the port

You must configure the duplex parameter from the Ethernet tab if the port is to be added to a LAG. Configure the Dot1 Q Ethertype and Q in Q Ethertype parameters from the Ethernet tab, if required. The range is 1536 to 65 535.

You must also configure the speed parameter from the General tab. The options are 10, 100, 1000, or 10 000, depending on the speed of the Ethernet interface.

Most OmniSwitch chassis offer four hybrid or combo ports. These ports consist of four paired 10/100/1000Base-T ports and four 1000 SFP ports. Preferences for these ports are configurable and, depending on the configuration, redundancy can be provided if a link fails.

16.5.6 PXC loopback ports

You can place a port in an internal loopback mode called a port cross-connect (PXC). Each PXC is associated with a single physical port, and contains two logical PXC sub-ports. One PXC sub-port is created per upstream or downstream path. PXC ports can be added to hybrid LAGs, and associated with the following interfaces:

- Base routing network interfaces
- VPRN network interfaces
- Epipe L2 access interfaces
- Ipipe L2 access interfaces
- VPLS L2 access interfaces
- VPRN L3 access interfaces
- VPRN tunnel interfaces
- IES L3 access interfaces
- IES IPsec interfaces

Supported breakout ports can be associated to a PXC. The PXC ports can be added to either a LAG or FPE. The LAG with the PXC can be associated to the FPE. See 12.41 “To create an FPE” (p. 374) for information about configuring FPE.

PXC ports are configured in the Equipment tree. See 16.59 “To configure PXC loopback ports” (p. 652) for information about configuring a PXC port.

16.5.7 Xconnect anchor ports

The PXC functionality uses the loopback mechanism of anchor ports to feed traffic from egress forwarding context into the ingress forwarding context on the same line card utilizing the E-chip functionality.

The Xconnect object appears automatically in the navigation tree under daughter cards supported NEs. After you create MAC on the Xconnect, you can create up to 2 loopbacks. Each loopback creates respective anchor ports in the navigation tree (Port 1/1/m1/1). The anchor ports can be associated to a PXC.

See 16.58 “To create and configure Xconnect anchor ports” (p. 650) for information about configuring anchor ports.

16.5.8 OmniSwitch learned port security

LPS provides a mechanism to control network device access on one or more OmniSwitch ports. Configurable LPS parameters allow you to restrict the source learning of host MAC addresses to:

- a specific amount of time in which the switch allows source learning to occur on all LPS ports
- a maximum number of learned MAC addresses allowed on the port
- a list of configured authorized source MAC addresses allowed on the port

The following options allow you to specify how the LPS port handles unauthorized traffic.

- Block only traffic that violates LPS port restrictions; authorized traffic is forwarded on the port.
- Disable the LPS port when unauthorized traffic is received; all traffic is stopped and a port reset is required to return the port to normal operation.

See 28.128 “To configure bridging on an OmniSwitch” (p. 1043) for information about enabling LPS on Ethernet ports and configuring LPS properties. See 16.56 “To configure OmniSwitch Ethernet ports” (p. 647) for information about configuring static MAC addresses on LPS enabled Ethernet ports.

Note: A deployment error is displayed in NFM-P while creating a LAG with ports. This deployment failure occurs if there is an LPS/VLAN configuration done on a specific port, using CLI.

16.5.9 MTU size and port configuration

You must specify the MTU size for an Ethernet port using the MTU (bytes) parameter on the General tab at the connection termination endpoint.

Consider the following when you configure MTU parameters.

- The managed devices must handle MTU limitations at many service points. The physical (access and network) ports, service, and service tunnel MTU values must be individually defined.
- The ports to be designated as network ports and the ports to be designated as access ports intended to carry service traffic must be identified.
- MTU values should not be frequently modified.
- Service MTU values must be less than or equal to the service tunnel MTU.
- Service MTU values must be less than or equal to the access port MTU.
- The MTU value for an in-band management port should be less than or equal to the MTU value for the peer port connected to the in-band management port.

See the device specific documentation for end-to-end considerations for configuring maximum MTU size throughout the managed network.

The Ethernet port MTU parameter indirectly defines the largest physical packet that the port can transmit or that the far-end Ethernet port can receive. Packets received that are larger than the MTU are discarded. Packets that cannot be fragmented at egress and that exceed the MTU are discarded.

The parameters for MTU configuration include the destination MAC address, source MAC address, Ethernet encapsulation type, length field, and complete Ethernet payload.

The MTU value for a port is associated with the port mode, such as access or network, and the port encapsulation type. If you change the mode or encapsulation type value for a port, the NFM-P adjusts the MTU value to a default value. If you do not want the MTU values for ports to revert to the defaults, you can configure the NFM-P to retain the currently configured MTU values for ports regardless of a mode or encapsulation type change. See 16.31 “To configure the NFM-P to retain non-default port MTU values” (p. 613) for more information.

16.5.10 HSMDA Egress Secondary Shapers

The egress port scheduler combines all subscriber queues of the same scheduling class and services the queues in a byte fair round robin fashion. This results in more packets being forwarded into the aggregation network towards a DSLAM than the DSLAM can accept. If the HSMDA egress port is congested, the egress bandwidth represented by the downstream discarded packets to the DSLAM may be allocated packets destined to other DSLAMs.

The HSMDA supports egress secondary shapers to provide a control mechanism to prevent downstream overruns without affecting the class-based scheduling behavior on the port. All subscribers destined to the same DSLAM have their queue groups mapped to the same egress secondary shaper. As the scheduler services the queues within the groups according to scheduler class, the destination shaper is updated.

After the shapers rate threshold is exceeded, scheduling for all queues associated with the shaper is stopped. When the dynamic rate drops below the threshold, the queues are allowed to be placed back on the scheduler service lists. By removing the queues from their scheduling context for a downstream congested DSLAM, the port scheduler is allowed to fill the egress port with packets destined to other DSLAMs without affecting class behavior on the port.

Egress secondary shapers are configured per port.

16.5.11 PoE

PoE provides power directly from the Ethernet ports. Powered devices such as IP phones, wireless LAN stations, Ethernet hubs, and other access points can be plugged directly into PoE-enabled Ethernet ports. The NFM-P supports PoE and PoE+ on supporting devices. See the following procedures:

- 16.49 “To configure PoE ports on a 7210 SAS” (p. 641)
- 16.52 “To configure PoE ports on a 7705 SAR” (p. 644)
- 16.57 “To configure OmniSwitch PoE Ports” (p. 650)

16.6 Copying and moving SAPs

16.6.1 Overview

You can copy and move SAPs between physical Ethernet ports, logical ports, or ports and LAGs.

You can also copy and move SAPs between endpoints configured with ATM encapsulation. For example, you can move a SAP on an IMA bundle to another bundle or to a TDM channel configured with ATM encapsulation. The function is typically used in redundancy scenarios, or to recover from a hardware failure.

Consider the following before you attempt to copy or move a SAP:

- Ports:
  - The source and destination ports can be on the same or on a different chassis. Inter-chassis copy or move is supported only when both NEs are of the same type, in the same chassis mode or from a lower chassis mode (source NE) to a higher chassis mode (destination NE), and at the same major software release.
  - The source and destination ports can be the same. This configuration allows you to change the encapsulation values for a group of SAPs on the same port using outer and inner encapsulation offset values. These values can be positive or negative, depending on whether the encapsulation values must increase or decrease.
  - The physical Ethernet ports can be of different types; for example, Fast Ethernet (10/100/1000 Base-T), Gigabit Ethernet (1000 Base-T), and 10 Gigabit Ethernet (10 GBase-T).
  - The source and destination ports must be configured as access ports.
  - The encapsulation type must be the same on the source and destination ports.
  - You cannot copy or move a SAP between an HSMDA port and a non-HSMDA port.
  - The copy and move operations fail if the encapsulation value for a SAP on the source port is used by a SAP on the destination port. You can modify the SAP encapsulation values used on the destination port, if required.
  - The selected L2 access interface SAPs on the source port can be moved only to another service, and not within the same service.
  - If a SAP on the source port belongs to a service not supported on the destination port, the SAP is not copied or moved.
  - The maximum number of SAPs varies for ports and MDAs. The SAP copy or move operation fails if the SAP capacity is exceeded for the destination port.
  - For L2 and L3 access interfaces, you can move all SAPs or a subset of the SAPs on a port. You can also copy all SAPs or a subset of the SAPs on a port for L2 access interfaces.
  - All SAPs associated with a specific port within a group interface contained in an L3 subscriber interface can be moved at the same time.
- LAGs:
  - LAG to SAP copy and move works on all services that support SAP to SAP copy and move. There are no restrictions.
  - LAG specific attributes such as lag-per-link-has are dropped during the copy or move.
  - The target can be a port or a LAG. The target LAG or port can be on the same or a different NE for L2 Services. For L3 Services the target must be the same NE.
  - Both the mode and the encap-type must match to be copied. The NFM-P only displays LAGs or ports that have matching mode and encap-type. You can copy from an Access LAG or port to a Hybrid LAG or port and vice versa.
  - Port speed is not checked. You must verify that the target of the copy or move has enough capacity.
- SAPs:
  - SAPs can be associated with either L2 access interfaces on services such as Apipes, Cpipes,

Epipes, Hpipes, Fpipes, Ipipes, VPLS, and MVPLS, or L3 access interfaces or subscriber interfaces that are associated with services such as IES and VPRN.

- L2 access interface SAPs can be copied or moved from one service to another service of the same service type. The source and target service must be on the same site, and the source and destination service site types must be the same, for example B-site, I-site.
- MEPs are copied or moved with the L2 SAP copy/move operation only if the Follow Service Topology Changes parameter is not enabled on the global MEG service. If the Follow Service Topology Changes parameter is enabled, the MEPs are automatically created by the OAM framework when the new SAP is created. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109).
- For an L2 SAP copy/move operation, only MEP CFM tests that are manually created are moved to the new SAP. MEP CFM tests created by Test Suite are not moved to the new SAP.
- For a SAP copy, the MEP ID is auto-assigned. For a SAP move, the MEP ID is copied only if the Node Wide MEP ID is not set on the destination NE. Otherwise the MEP ID will be auto-assigned based on the configured MEP ID. See 91.14 “To configure an automatic MEP ID assignment on an NE” (p. 3108).
- Endpoints associated with the specified SAPs are copied during an L2 SAP copy/move operation, but other objects associated with the endpoints such as other SAPs or spokes are not copied/moved. If the destination already has an endpoint with the same name, the endpoint associated with the source SAP will not be copied/moved with the source SAP.
- You cannot change SAP encapsulation from bridged to routed if ARP and DHCP options have been configured.

- SAP attributes:
  - VPLS and MVPLS dynamic FIB entries that are associated with copied or moved SAPs are discarded. Static FIB entries are transferred to the destination port after a successful SAP copy or move.
  - All SAP redundancy relationships at the source port, such as those provided by an MVPLS, are discarded during a copy or move operation.
  - All alarms, statistics, or OAM test results that are associated with a copied or moved SAP are discarded.
  - When a source SAP belongs to an SHG, the NFM-P attempts to preserve the SHG membership. The destination SHG must have the same name as the source SHG, and must have the same residential SHG status. If an SHG with the same name does not exist on the destination service site, the NFM-P creates it.
  - When a source SAP uses an aggregation scheduler, the scheduler is copied to the destination. The aggregation scheduler must be unique to the destination NE and must have the same scope, for example, MDA or port.
- MEPs
  - MEPs and associated CFM tests are included in SAP copy/move operations for L3 Access Interface and Subscriber Interface SAPs.

See 16.32 “To copy or move L2 SAPs between ports” (p. 614) , 16.33 “To copy or move L2 access interface SAPs between services” (p. 619) , 16.34 “To move L3 SAPs within or between ports or LAGs on the same NE” (p. 620) , and 16.35 “To move L3 subscriber interface SAPs between ports on the same NE” (p. 623) to move access interface SAPs between services or ports.

16.7 Configuring access SAP bandwidth CAC

16.7.1 Overview

You can configure the bandwidth CAC function on a port or LAG, based on the admin bandwidth configured on a SAP and on an associate port or LAG. In addition, a booking factor is provided to allow over/under booking of the sum of the SAP bandwidth compared to the port/LAG bandwidth. See 16.36 “To configure bandwidth CAC on an access SAP for services or a LAG” (p. 625).

The goal of the CAC function is to ensure that the sum of the admin SAP bandwidth on a port or LAG does not exceed the admin bandwidth configured on that port or LAG. This ensures there is sufficient bandwidth available on the associated port or LAG before a SAP is operationally up.

This feature is supported for SAPs on all 7x50 platforms with FP2-and higher hardware on the following service sites:

- VPLS
- VPRN
- MVPLS
- EPIPE
- IPIPE
- IES

The admin bandwidth is an abstract value which could represent either the ingress or egress bandwidth; the bandwidth is a statically configured value; there is no attempt to use any dynamic bandwidth information associated with a SAP, port or LAG.

An admin bandwidth can only be configured on a SAP connected to a port or LAG which itself has an admin bandwidth configured. When a LAG is configured, the admin bandwidth and booking factor on its constituent ports are ignored.

The NFM-P tracks the requested and available bandwidth per port or LAG; where the available bandwidth is equal to the admin bandwidth on the port or LAG, with the booking factor applied, minus the sum of admin bandwidth configured on its SAPs. The SAP bandwidth CAC state is maintained on a port or LAG basis as long as the NFM-P can associate a SAP to the port or LAG. The NFM-P will perform a server validation to make sure the booking bandwidth is not more than the available bandwidth.

16.8 SONET/SDH and TDM port encapsulation

16.8.1 Overview

SONET/SDH and TDM ports can be configured as connection termination points in SAPs and IP interfaces. They can be channelized.

An access port is used for customer-facing traffic on which services are configured. SAPs can only use an access port. When a port is configured for access mode, the appropriate encapsulation type must be specified to distinguish the services on the port.

You must configure the Encap Type parameter from the General tab at the connection termination point. SONET/SDH or TDM ports or channels support the following encapsulation types, depending on the MDA type:

- BCP Null
- IPCP
- BCP Dot1 Q
- FR — Support multiple services using the DLCI header to distinguish services
- ATM — Support multiple services using the VCI or VPI of the PVC
- PPP Auto
- CEM

MTU size considerations apply to SONET/SDH and TDM channels. See 16.5.9 “MTU size and port configuration” (p. 573) in 16.1.2 “Working with port and channel objects” (p. 565) for more information.

16.9 SONET clear channel applications

16.9.1 Overview

Ports on OC-*n* cards can be used for SONET clear channel applications. SONET or TDM clear channel applications allow you to create a full channel on a port which can be configured as access or network mode for SONET and access only for TDM. For example, when you create a 16 × OC3 SFP card, 16 ports are created. You can create one full channel on each of these ports. For more information about the 4 × OC3/STM1 ASAP SFP and 4 × Channelized OC3/OC12 ASAP SFP adapter cards for clear channel support, see 16.9.2 “OC3/STM1 clear channel support on the 7705 SAR-8 and 7705 SAR-18” (p. 579) in this section.

STS-192/48/12/3 clear channel applications use the following syntax:

```
card slot/daughtercard/port.STSType
```

For example, the clear channel STS-12 on slot 4, MDA 1, port 1 is named 4/1/1.sts12

The following table lists available SONET channel applications and parameters for clear channel, sub-channel, and TDM applications.

Table 16-2 SONET channel parameters

Applications	Channel ranges								
	STS-192	STS-48	STS-3	STS-1	DS3	DS1	DS0	E1	DS0
SONET Clear channel	1	1	1		1				
SONET Sub channel			1 to 4	1 to 3	1	1 to 28	1 to 24	1 to 21	2 to 32
TDM					1	1 to 28	1 to 24	1 to 21	2 to 32

16.9.2 OC3/STM1 clear channel support on the 7705 SAR-8 and 7705 SAR-18

The 7705 SAR-8 and 7705 SAR-18 support the 4 × OC3/STM1 ASAP SFP card. The 4 × OC3/STM1 ASAP SFP card has four SFP-based ports that can be independently configured for ATM in access mode or for Packet over SONET/SDH (POS) in network mode. Each port can be independently configured to be SONET (OC3) or SDH (STM1) framing. See 16.11 “ATM encapsulation” (p. 581) in this section for more information. The interface must be configured as UNI. The transmit clock rate for a port can be device- or loop-timed.

Note: Channels on the 4 × OC3/STM1 ASAP SFP adapter card cannot be configured as IES SAPs.

The section trace (J0) byte can be configured by the operator to check the physical cabling. The port can activate and deactivate local line and internal loopbacks. The MTU size for a 4-port OC3/STM1 ASAP SFP access port is fixed at 1572 for ATM encapsulation.

CSM activity switches for a 4-port OC3/STM1 ASAP SFP card on the 7705 SAR-8 and 7705 SAR-18 are hitless on the data path.

The 7705 SAR-8 and 7705 SAR-18 support the 4 × Channelized OC3/OC12 ASAP SFP adapter card. In OC3 mode, the card supports path sts3 under SONET and path sts3 with payload sts3 under SDH for clear channel POS and network APS. In OC12 mode, the card supports sts3 and sts12 under SONET and SDH for clear channel POS and network APS.

The 7705 SAR-8 and 7705 SAR-18 support the 2-port OC3/STM1 ASAP SFP card. The 2-port OC3/STM1 ASAP SFTP card has two SFP-based ports (optical or electrical) that can be configured for ATM/IMA or TDM in access mode or for MLPPP in access or network mode. The port type must be configured to be either SONET (OC3) or SDH (STM1).

16.10 TDM channelization and clear channel applications

16.10.1 Overview

When you create a 4 or 12 × DS3/E3 card, 4 or 12 DS3/E3 ports are created. You can then create DS3/E3 channels using the NFM-P, one per port. Each DS3/E3 channel can be channelized into 28 independent DS1 or 21 independent E1 data channels or, in clear channel applications, the DS3 can be the connection termination point. For channelized DS3 connections, each DS1 channel can be channelized to 24 DS0 groups and each E1 channel can be channelized to 31 DS0 groups. To use a DS1 or E1, you must create at least one DS0 group for the DS1 or E1.

By default, DS3 ports are automatically created for clear channel connections. To create a channelized DS3, you must configure the DS3 channel type as E3 or DS3 on the port. To channelize the DS3 to the DS0 level, you must set the Channelization Type to Channelized DS1 or E1. The following figure shows the channelized DS3 port structure for DS1 channels.

Figure 16-3 Channelized 12 × DS3 port structure for DS1 Channels

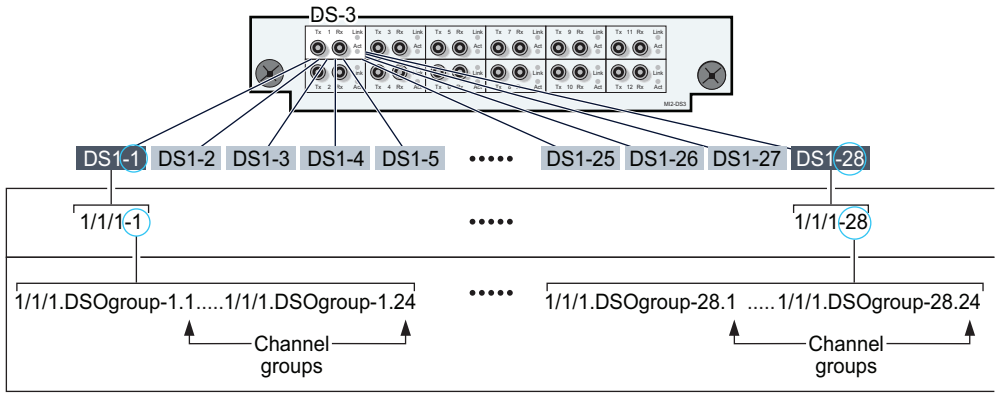

The channelized DS3 port structure for E1 channels has the following parameter values:

- each DS3 port supports 21 E1 channels
- each E1 channel supports 31 DS0 channel groups

TDM-based DS3 channelization uses the following syntax:

DS3 channel configured for TDM:

```
card slot/daughtercard/port.DS3-
```

DS1 channel from a TDM-based DS3 channel:

```
slot/daughtercard/port.DS1-[DS3#].[DS1#]
```

E1 channel from a TDM-based DS3 channel:

```
slot/daughtercard/port.DS1-[DS3#].[E1#]
```

DS0 group channel from the DS1 channel:

```
slot/daughtercard/port.DS0Grp-[STS3#].[STS1#].[DS1#/E1#].[Group#]
```

The following table provides an example of the naming conventions for a 12 × DS3 port.

Table 16-3 Example of TDM channel naming convention

Syntax	Description	Additional information
Channel 1/1/1.ds3	1/1/1 is the slot number/daughtercard number/port number .ds3 identifies the channel as DS3	Because DS3s are unchannelized by default, you must configure the Channelization Type as Channelized.

Table 16-3 Example of TDM channel naming convention (continued)

Syntax	Description	Additional information
Channel 1/1/1.ds1-1.2	.ds1 identifies the channel as DS1 1 is the DS3 number .2 is the DS1 number (1 to 28)	Identifies the DS1 channel and shows how the DS3 level acts as a place holder for the DS1s.
Channel 1/1/1.e1-1.2	.e1 identifies the channel as E1 1 is the DS3 number .2 is the E1 number (1 to 21)	Identifies the E1 channel and shows how the DS3 level acts as a place holder for the E1s.
Channel 1/1/1.ds0Grp-1.2.23	.ds0Grp- identifies the channel as a DS0 group 1 is the DS3 number .2 is the DS1 number (1 to 28) or E1 number (1 to 21) .23 identifies the DS0 channel group (1 to 24) on the DS1 or (2 to 32) on the E1	The DS0 group is configured on the DS1 or E1 channel. Only a DS0 can be used as a CTP.

16.11 ATM encapsulation

16.11.1 Overview

SONET/SDH clear channels can be provisioned so that ATM cells are encapsulated in SONET/SDH frames. The entire SONET/SDH path of the port is then used to carry ATM cells. The channel of the port becomes the ATM interface. ATM encapsulation is supported on the following daughter cards:

- 16 × ATM OC3 SFP
- 16 × ATM OC3 SFP B
- 4 × ATM OC12/OC3 SFP
- 4 × ATM OC12/OC3 SFP B
- 4 × OC3/STM1 ASAP SFP
- 4 × Any Service Channelized OC3
- 16 × Any Service Channelized DS1/E1

SONET/SDH clear channel applications with ATM encapsulation use the following syntax:

```
Daughter Card Slot - 3 (16 × ATMOC3), OK
Port 3/1/1 - Speed: OC3, State: OK
Channel 3/1/1.sts3, Mode: Access, Encap: ATM, State: Ok
```

Customer devices with ATM interfaces are connected directly or using the ATM access network to a 7750 SR that offers IES or VPRN services.

A SONET/SDH clear channel with ATM encapsulation can carry multiple ATM PVCs. ATM PVCs cannot be created from the NFM-P. They are created automatically by the router when a new Layer 3 interface over ATM is created. The following figure shows an example of ATM PVCs on a SONET/SDH clear channel.

Figure 16-4 ATM PVCs on a SONET/SDH clear channel

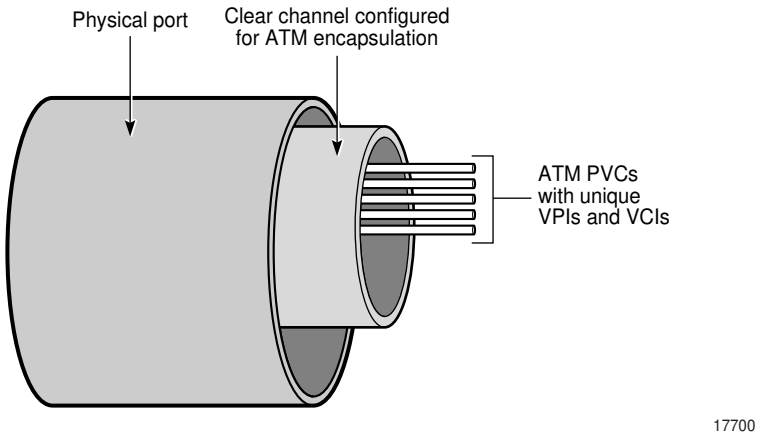

16.11.2 ILMI links between ATM interfaces

ILMI links can be configured between ATM interfaces. ILMI is a protocol that sets and sends physical layer, ATM layer, virtual path, and virtual channel parameters that are exchanged by ATM interfaces.

When ILMI is run between two ATM interfaces, the interfaces exchange ILMI packets that consist of SNMP messages across the physical connection. An ILMI link allows ILMI to run on an available virtual channel. When an ILMI link is created to carry ILMI messages, a PVC is automatically created with default PVC attributes. The PVC is deleted when the ILMI link is removed.

16.11.3 IMA

IMA is supported on channelized ASAP MDAs on the 7750 SR, 7705 SAR. For Wavence device IMA support, see the *NSP Wavence Device Support Guide*.

IMA group bundles aggregate E1 or DS1 ATM paths into a single logical ATM interface. Each IMA group bundle can have from 1 to 8 members, or links. Up to 56 IMA group bundles can be configured on a single 7750 SR MDA. Up to 8 IMA group bundles can be configured on a single 7705 SAR daughter card.

Each ATM Interface supports up to 48 VP/VC connection points. Each 16 × E1 ASAP Access (ETSI) board supports up to 128 VP/VC connection points.

An IMA group bundle has the following common interface properties:

- ATM encapsulation type is displayed on the ATM Interface properties form by clicking on the “Edit ATM button”.
- MTU value (of the primary link)

Member links inherit the common properties of the IMA group to which they belong. Only the properties of the primary link can be changed. The primary link is the member with the lowest interface index or initialization sequence. The primary link may change as member links are added and deleted from the IMA group.

Consider the following when you create a multilink IMA group bundle.

- When a path becomes a member of an IMA group bundle, it is no longer a physical ATM path interface.
- The members of the IMA group bundle inherit all of the properties of the primary link, such as the SONET configuration and MTU. If you modify the configuration of the primary link, the configuration of the bundle members is also modified.
- Member links that are added after the primary link has been added to the IMA group bundle must match the configuration of the primary link.
- You cannot configure services on a member link.

16.12 Workflow to manage port objects

16.12.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure port objects. This workflow assumes that the physical devices have been installed, commissioned, and discovered. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: After they are configured, port objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree. for more information.

16.12.2 Stages

1

As required, configure Ethernet ports; see 16.24 “To configure Ethernet ports” (p. 599) .

2

As required, change the port mode; see 16.29 “To change the port mode” (p. 611) .

3

As required, migrate SAPs from access mode to hybrid mode; see 16.30 “To migrate SAPs from access mode to hybrid mode” (p. 612) .

4

As required, configure the NFM-P to retain non-default port MTU values; see 16.31 “To configure the NFM-P to retain non-default port MTU values” (p. 613) .

5

As required, copy or move L2 access interface SAPs between ports; see 16.32 “To copy or move L2 SAPs between ports” (p. 614) .

6

As required, move L3 access interface SAPs within or between ports on the same NE; see 16.34 “To move L3 SAPs within or between ports or LAGs on the same NE” (p. 620) .

7

As required, move L3 subscriber interface SAPs between ports on the same NE; see 16.35 “To move L3 subscriber interface SAPs between ports on the same NE” (p. 623) .

8

As required, configure the bandwidth CAC function on a port or LAG; see 16.36 “To configure bandwidth CAC on an access SAP for services or a LAG” (p. 625) .

9

As required, add a queue group to an Ethernet port; see 16.37 “To add a queue group to an Ethernet port” (p. 627) .

10

As required, configure SONET ports; see 16.39 “To configure SONET ports” (p. 631) .

11

As required, configure an HSMDA port scheduler override; see 16.40 “To configure an HSMDA override” (p. 632) .

12

As required, configure TDM DS3 ports; see 16.41 “To configure TDM DS3 ports” (p. 633) .

13

As required, configure serial ports; see 16.42 “To configure serial ports” (p. 634) .

14

As required, configure PW ports; see 16.43 “To configure PW ports” (p. 634) .

15

As required, configure ports on 7210 SAS devices, as follows:

a. Configure a 7210 SAS-M channelized TDM DS1 or E1 port; see 16.44 “To configure a 7210 SAS-M channelized TDM DS1 or E1 port” (p. 635) .

b. Assign policies to a 7210 SAS Ethernet port; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) .

c. Configure an SHG for 7210 SAS NEs; see 16.46 “To create a 7210 SAS SHG” (p. 638) .

d. Configure PoE ports on supporting 7210 SAS NEs; see 16.49 “To configure PoE ports on a 7210 SAS” (p. 641) .

16

As required, configure ports on 7705 SAR devices, as follows:

a. virtual Ethernet ports on a 7705 SAR 2-port ring MDA; see 16.47 “To configure a virtual Ethernet port on a 7705 SAR 2-port ring MDA” (p. 639) .

b. PoE ports on a 7705 SAR; see 16.52 “To configure PoE ports on a 7705 SAR” (p. 644) .

c. GPS on a 7705 SAR; see 16.53 “To configure GPS on a 7705 SAR” (p. 645) .

d. 7705 SAR ASAP channelized TDM port; see 16.54 “To configure a 7705 SAR ASAP channelized TDM port” (p. 646) .

17

As required, configure the following on OmniSwitch devices:

a. OmniSwitch Ethernet ports; see 16.56 “To configure OmniSwitch Ethernet ports” (p. 647) .

b. OmniSwitch PoE ports; see 16.57 “To configure OmniSwitch PoE Ports” (p. 650) .

SONET and SDH sub-channel applications and structure

16.13 Overview

16.13.1 SONET sub-channel applications

SONET sub-channel applications allow you to create multiple STS-1 channels on deep channelized OC-12 and OC-3 ports, and to configure multiple DS0 connection termination points.

An STS-1 sub-channel is configured, or channelized, to carry one of the following payload types:

- DS3
- VT1.5
- VT2

Because of the differences between SONET and SDH in multiplexing and mapping of tributary payloads into higher digital levels, some building blocks of SDH have no SONET equivalent and therefore the containment hierarchy and terminology for SDH is different from SONET. See 16.15 “Comparison of SONET and SDH hierarchies” (p. 589) and 16.16 “SDH AU-4 and AU-3 sub-channel applications” (p. 590) in this section.

16.13.2 SONET DS3 payload

The following example, illustrated in Figure 16-5, “Channelized 1 × OC-12 port structure using STS-1/DS3 sub-channels” (p. 587) , shows the channelization sequence for a DS3 payload on a 1 × OC12 Deep Channel card.

1. One OC-12 port is created when you create a channelized 1 × OC12 Deep Channel card.
2. This port can be channelized into 12 SONET STS-1 sub-channels from the four STS-3s available on the port.
3. You can then configure each of these STS-1s to carry a DS3 frame
4. Each DS3 frame can be channelized into 28 independent DS1 data channels or 21 independent E1 data channels. Each channel must be created one at a time.

Figure 16-5 Channelized 1 × OC-12 port structure using STS-1/DS3 sub-channels

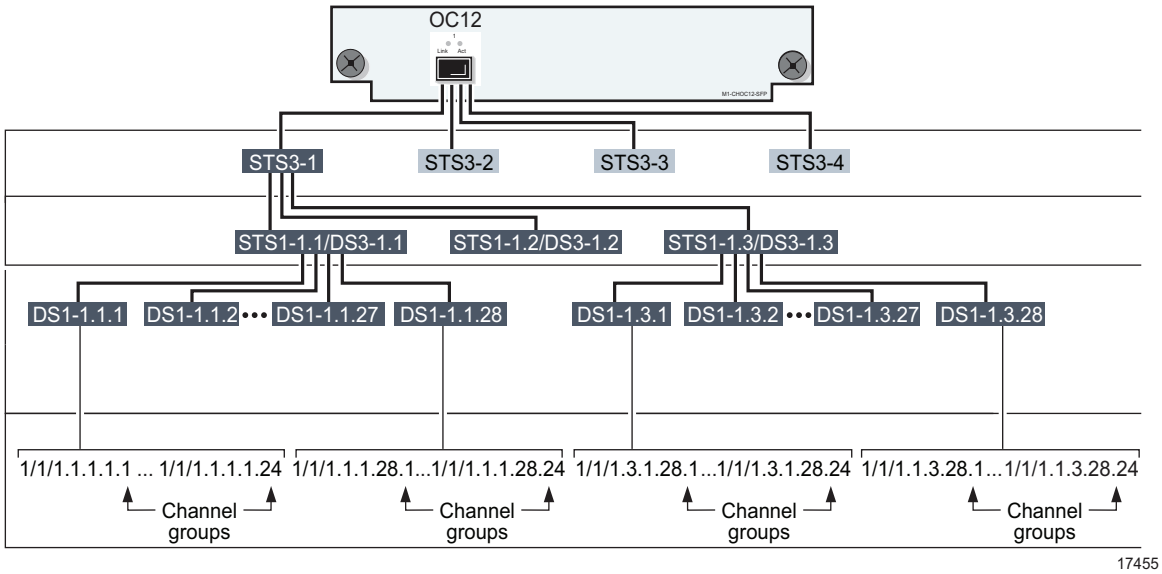

16.13.3 SONET VT1.5 and VT2 payloads

The following example, illustrated in Figure 16-6, “Channelized 4 × OC-3 port structure using VT sub-channels” (p. 588) , shows the channelization sequence for a VT1.5 or VT2 payload on a 1 × OC12 Deep Channel card.

1. One OC-12 port is created when you create a channelized 1 × OC12 Deep Channel card.
2. This port can be channelized into 12 SONET STS-1 sub-channels from the four STS-3s available on the port.
3. You can then configure each of these STS-1s to be channelized to carry up to 28 DS1. For SONET sub-channel configuration you must select payload type VT1.5 or VT2. When you choose VT1.5 payload type, seven VTG are implicitly created.
4. Each VTG is channelized into four independent VT1.5 channels or three VT2 channels, each of which can carry an independent DS1 data channel. Each channel must be created one at a time.
5. Each VT1.5 DS1 channel can be configured to handle up to 24 DS0 groups. Each VT2 DS1 channel can be configured to handle up to 32 DS0 groups. To use a DS1, you must create at least one DS0 group for the DS1 or E1.

Figure 16-6 Channelized 4 × OC-3 port structure using VT sub-channels

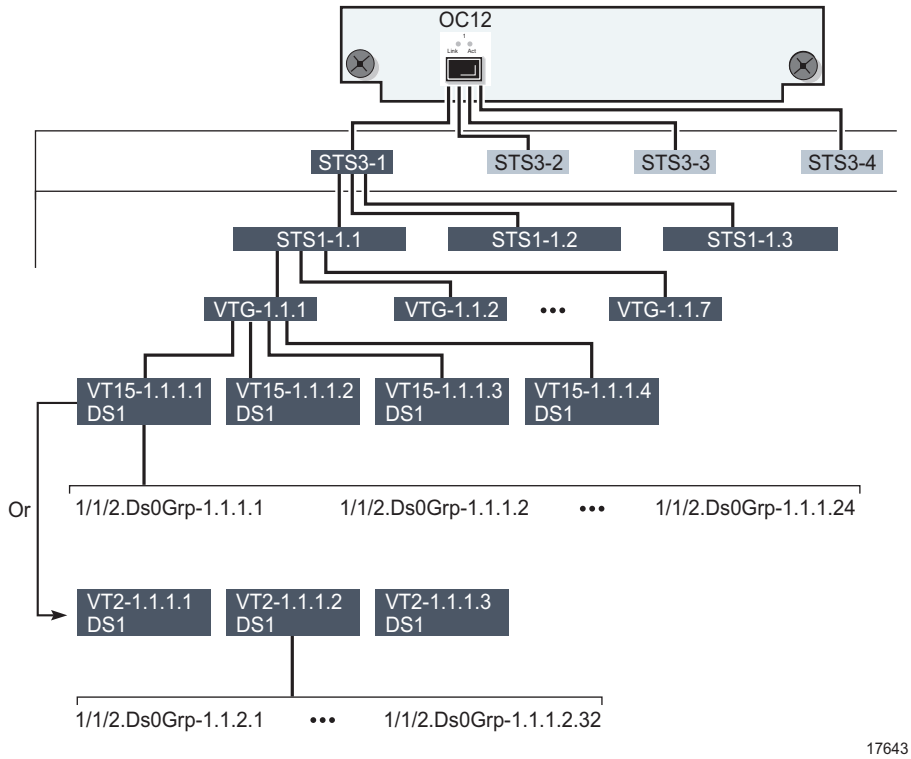

16.14 SONET sub-channel syntax

16.14.1 Syntax examples

The NFM-P uses the following SONET syntax for the STS-1 sub-channel:

```
card slot/daughtercard/port.STS1-[STS3#].[STS1#]
```

DS3 channels from an STS-1 sub-channel use the following syntax:

```
slot/daughtercard/port.DS3-[STS3#].[STS1#]
```

DS1 channels from a DS3 channel use the following syntax:

```
slot/daughtercard/port.DS1-[STS3#].[STS1#].[DS1#]
```

E1 channels from a DS3 channel use the following syntax:

```
slot/daughtercard/port.DS1-[STS3#].[STS1#].[E1#]
```

VT15 or VT2 from a VT group use the following syntax:

```
slot/daughtercard/port.VTG#-[STS3#].[STS1#]
```

DS1 channels from a VT15 or VT2 use the following syntax:

```
slot/daughtercard/port.DS1-[STS3#].[STS1#].[VT15#] or [VT2#}.[VTG#].[DS1#]
```

Table 16-4 Example of SONET sub-channel syntax for an OC-12 port

Syntax	Description	Additional information
Channel 1/1/1.sts1-2.2	1/1/1 is the slot number/daughtercard number/port number 2 is the STS-3 number (1 to 4) .2 is the STS-1 number (1 to 3)	The sts1 parameter is 2.2 which means that it is the fifth STS-1. There are four STS-3s for an OC-12 and each STS-3 has three STS-1s such that the fifth STS-1 is the second STS-1 of the second STS-3.
Channel 1/1/1.ds3-2.2	.ds3-2 is the STS3 number (1 to 4) .2 identifies the STS-1 number (1 to 3)	Identifies the DS3 channel and shows how the sts1 level acts as a place holder for the DS3
Channel 1/1/1.ds1 2.2.25	.ds1-2.2.25 identifies a DS1 channel (1 to 28) on the DS3	The DS1 is configured on the DS3.
Channel 1/1/1.e1 2.2.21	.e1-2.2.21 identifies a E1 channel (1 to 21) on the DS3	The E1 is configured on the DS3.
Channel 1/1/1.DS0Grp-2.2.20.5	.DS0Grp-2.2.20.5 identifies one of the DS0 channel groups: 1 to 28 for DS1; 2 to 32 for E1	The DS0 is configured on the DS1 or E1 channel.
Channel 1/1/1.ds1 2.2.3.18	ds1 2.2.3.18 identifies a DS1 channel on the VT15 or VT2 payload, where 3 is the VT group.	—

16.15 Comparison of SONET and SDH hierarchies

16.15.1 Overview

SONET and SDH are compatible digital hierarchies that support identical transmission rates with similar framing structures. However, SONET and SDH standards differ in the way they multiplex and map tributary payloads into higher digital levels. Because of these differences, some SDH framing blocks do not have equivalent blocks in SONET.

SONET provides one STS sub-channel, the STS-1, for mapping lower-capacity, deep-channel payloads while SDH provides two alternative sub-channels, the AU3 and the AU4. After an AU3 sub-channel is created, an AU4 cannot be created on that port in the case where the port is an OC-3 (STM-1). Conversely, if an AU4 sub-channel is created, an AU3 cannot be created on that port in the case where the port is an OC-3 (STM-1).

SONET and SDH transmission rates converge at 155.520 Mb/s where the SONET STS-3c is equivalent to the SDH STM-1 bit rate. For lower-capacity payloads, such as DS1, E1, DS3, and E3, SONET provides one unique mapping path for each payload while SDH permits two alternative paths for each payload as shown in the following figure.

Figure 16-7 Supported SONET/SDH multiplexing structures

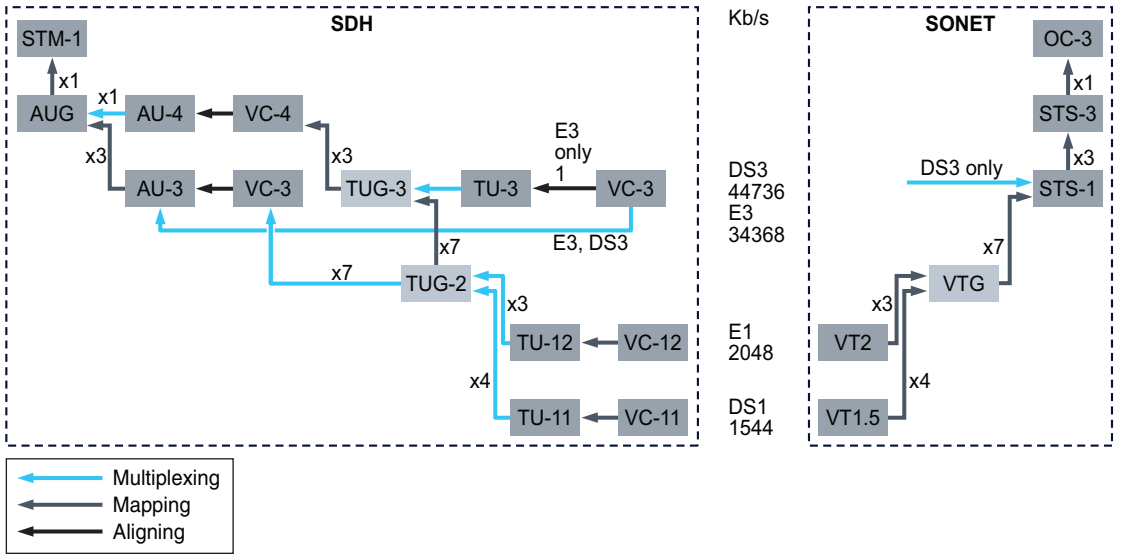

17610

16.16 SDH AU-4 and AU-3 sub-channel applications

16.16.1 Overview

SDH sub-channel applications allow you to create AU4 or AU3 channels on deep channelized OC-12 (STM-4) and OC-3 (STM-1) ports, and to configure multiple DS0 connection termination points.

After an AU4 sub-channel is created, it implicitly contains three TUG-3 groups. A TUG-3 is channelized to carry one of the following payload types:

- TU3
- TU11
- TU12

After an AU-3 sub-channel is created, it is channelized to carry one of the following payload types:

- DS3, E3
- TU11
- TU12

16.17 SDH TU3 payload

16.17.1 Overview

The following example, illustrated in Figure 16-8, “Channelized 4 × OC-3 port structure using AU4/TU3 sub-channels” (p. 590) , shows the channelization sequence for a TU3 payload on a 4 × OC3 Deep Channel card.

1. Four OC-3 ports are created when you create a channelized 4 × OC3 Deep Channel card.
2. The user creates each required AU4 sub-channel.
3. Each AU4 sub-channel implicitly contains three TUG3 groups.
4. You can then configure each TUG3 to carry a TU3 frame
5. Each TU3 frame is channelized into an independent E3 data channel and cannot be changed.

Figure 16-8 Channelized 4 × OC-3 port structure using AU4/TU3 sub-channels

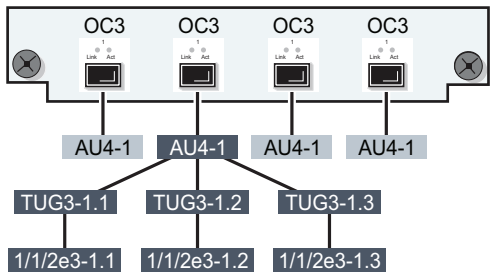

16.18 SDH E3 or DS3 payload

16.18.1 Overview

The following example, illustrated in Figure 16-9, “Channelized 4 × OC-3 port structure using AU3/E3 sub-channels” (p. 592) , shows the channelization sequence for an E3 or DS3 payload on a 4 × OC3 Deep Channel card using AU3 sub-channels.

1. Four OC-3 ports are created when you create a channelized 4 × OC3 Deep Channel card.
2. Each port can be channelized into up to three AU3 sub-channel.
3. Each AU3 frame can be channelized into an independent E3 or DS3 data channel.

Figure 16-9 Channelized 4 × OC-3 port structure using AU3/E3 sub-channels

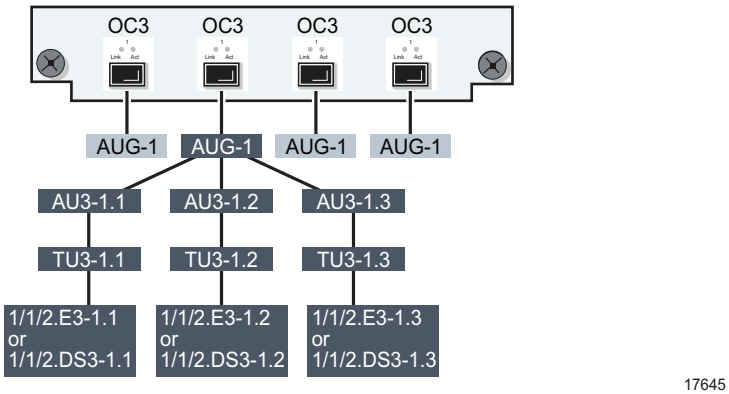

16.19 SDH TU11 and TU12 payloads

16.19.1 Overview

TU11 and TU12 payloads can be channelized using an AU4 sub-channel or an AU-3 sub-channel. An AU4 sub-channel implicitly contains three TUG3 groups. A TUG3 and an AU3 channelized to carry TU11 or TU12 payloads implicitly contain seven TUG2 tributary groups.

The following example, illustrated in Figure 16-10, “Channelized 4 × OC-3 port structure using AU4/TU sub-channels” (p. 593) , shows the channelization sequence for a TU11or a TU12 payload on a 4 × OC3 Deep Channel card.

1. Four OC-3 ports are created when you create a channelized 4 × OC3 Deep Channel card.
2. Each port can be channelized into one AU4 sub-channel or three AU3 sub-channels.
3. Each AU4 sub-channel implicitly contains three TUG3 groups.
4. You can then configure each TUG3 group or each AU3 sub-channel to carry a TU11 or TU12 payload.
5. Each TUG3 group or AU3 sub-channel implicitly contains seven TUG2 tributary groups.
6. Each TUG2 group can be channelized into three independent TU12 channels or four independent TU11 channels.
7. Each TU12 can contain an E1 or DS1 data signal. Each TU11 can contain a DS1 data signal. Each E1 can be configured to handle up to 31 DS0 groups. Each DS1 can be configured to handle up to 24 DS0 groups. To use an E1 or a DS1, you must create at least one DS0 group for the DS1 or E1.

Note: On the same TUG3 or AU3, you cannot mix TU-1*X* with DS1 and E1 payload types. Every TU-1*X* on the same TUG3 or AU3 must use the same payload type, either E1 or DS1, not a mix of both. The mix is shown in for illustrating the structure.

Figure 16-10 Channelized 4 × OC-3 port structure using AU4/TU sub-channels

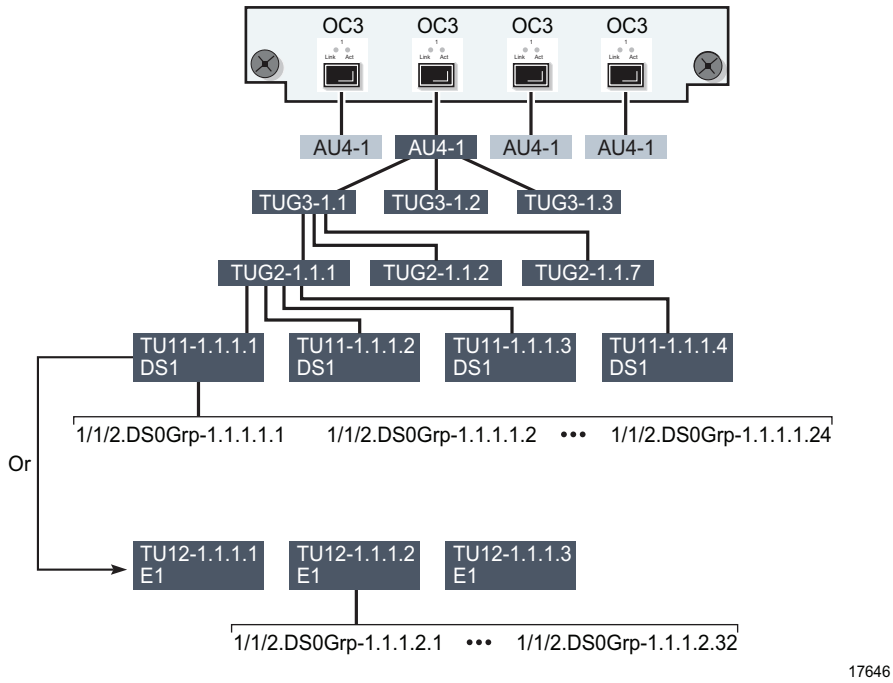

16.20 Workflow to manage channel objects

16.20.1 Purpose

The following workflow describes the sequence of high-level tasks required to manage and configure channel objects. This workflow assumes that the physical devices have been installed, commissioned, and discovered, and that the port objects have been configured. See Chapter 8, “Device commissioning and management” for more information about device commissioning. See Chapter 9, “Device discovery” for more information about device discovery.

Note: After they are configured, channel objects can be accessed using the equipment navigation tree. See Chapter 3, “NFM-P navigation tree” for more information about using the equipment navigation tree.

16.20.2 Stages

1

As required, configure SONET clear channels; see 16.60 “To configure SONET clear channels” (p. 653) .

2

As required, create channels on ports for card types that support multiple sub-channels; see 16.61 “To perform a bulk channel creation on ports that support multiple sub-channels” (p. 654) .

3

As required, configure SONET sub-channels; see 16.62 “To configure SONET sub-channels” (p. 655) .

4

As required, configure SDH sub-channels; see 16.63 “To configure SDH sub-channels” (p. 657) .

5

As required, create VT15 (TU11) or VT2 (TU12) sub-channels; see 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

6

As required, create TDM DS1 channels; see 16.67 “To create serial channels” (p. 665) .

7

As required, configure TDM DS1 or E1 channels; see 16.66 “To configure TDM DS1 or E1 channels” (p. 663) .

8

As required, configure serial channels; see 16.67 “To create serial channels” (p. 665) .

9

As required, create TDM DS3 channels; see 16.68 “To create TDM DS3 channels” (p. 666) .

10

As required, configure TDM DS3 channels; see 16.69 “To configure TDM DS3 channels” (p. 668) .

11

As required, configure a DS3 or E3 channel as a network interface on a channelized ASAP MDA; see 16.70 “To configure a DS3/E3 channel as a network interface on a channelized ASAP MDA” (p. 671) .

12

As required, configure an L3 interface on a DS3 or E3 channel on a channelized ASAP MDA; see 16.72 “To configure an L3 interface on a DS3/E3 channel on a channelized ASAP MDA” (p. 674) .

13

As required, configure a PVC; see 16.73 “To configure a PVC” (p. 676) .

14

As required, create an ILMI link; see 16.74 “To create an ILMI link” (p. 677) .

15

As required, modify an ILMI link; see 16.75 “To configure an ILMI link” (p. 678) .

Procedures for port configuration

16.21 To configure 1830 VWM ports

16.21.1 Steps

1

On the equipment tree, expand Network→1830–VWM-OSU NE→Shelf→Card Slot→Port object and choose Properties. The Physical Port (Edit) form opens.

2

Configure the Description parameters on the General tab.

3

Click on the States tab and configure the Administrative State parameter.

4

Click on the Media Adaptor tab and view the parameters listed. You can verify the parameters using the device CLI.

5

Click on the Port Specifics tab, and configure the required parameters.

- Configure the Topology String 1 and Topology String 2 parameters.
- Configure the AINS timer parameter.

  The AINS timer can be configured at the system level or at the port level. By default, all ports are AINS enabled and align to the AINS timer set by the system timer.

If you are configuring for the 1830 VWM TLU–200, configure the TLU–200 Specifics parameters.

6

Save your changes and close the form.

END OF STEPS

16.22 To configure Ethernet LAN ports on the 1830 VWM OSU and 1830 VWM SMM

16.22.1 Steps

1

On the equipment tree, expand Network→1830–VWM NE→Shelf→Card Slot→Port object and choose Properties. The Physical Port (Edit) form opens.

2

Configure the Description parameters on the General tab.

3

Click on the States tab and configure the Administrative State parameter.

4

Click on the Media Adaptor tab and view the parameters. You can verify the parameters using the device CLI.

5

Click on the Port Specifics tab and configure the required parameters.

The AINS timer can be configured at the system level or at the port level. By default, all ports are AINS enabled and align to the AINS timer set by the system timer.

6

Perform one of the following:

a. For an 1830 VWM OSU, configure the Interface role from the drop-down menu.

   If Userdata is selected as the Interface Role, the TPID parameter is enabled. The TPID parameter is not available if you set the Interface Role parameter to Normal.

b. For an 1830 VWM SMM, configure the MAU and Userdata parameters.

7

Save your changes and close the form.

END OF STEPS

16.23 To configure connector ports and breakout ports

16.23.1 Before you begin

Connector ports are indicated in the equipment tree by the letter c preceding the port number; for example, 10/1/c1/. Each connector port has the capability to support multiple Ethernet breakout ports.

DWDM channel and coherent optics are supported under connector ports and OTU is supported under breakout ports.

Depending on the NE and release, the tabs and parameters displayed may vary.

16.23.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

For XIOM-s, expand Network→*NE*→Shelf→Card Slot *n*→Xiom Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n/n*.

2

Right-click on the Port icon and choose Properties. The Connector Port (Edit) form opens.

3

Configure the required parameters.

Use the Connector Breakout Type parameter to configure the number and class of Ethernet ports to be housed by the connector port. For example, c4-10g indicates four 10 GB Ethernet ports.

For FP4-based QSFP28 MDAs on 7x50 NEs, the configuration of breakout connector types c1-10g and c1-25g (or just c1-10g for E4 platforms) implies you will be installing an adapter with a SFP+ or SFP28 optical module.

When you enable the CFP2-DCO parameter, the Coherent Optics tab becomes enabled and the Optical Transport Channel Unit tab becomes enabled under breakout port.

You can only set the DWDM channel after you enable the CFP2-DCO parameter. After you select the DWDM channel the wavelength and frequency associated with the channel are automatically populated. You cannot configure the DWDM channel under the General tab if frequency is configured under the DWDM Optics tab.

For 400G ZR optical module support, DWDM must be configured on 400G ZR QSFP-DD and DCO must be enabled on the transceiver before NFM-P DWDM and coherent optical configuration.

For the x2-1000g-wdm XMA, the CFP2-DCO parameter does not appear, as DWDM and coherent optics are supported by default.

4

To configure the center frequency and view the minimum/maximum frequencies and supported frequency grid for a particular optical module, click the DWDM Optics tab and configure the required parameters.

You cannot configure the frequency here if the DWDM channel is configured under the General tab. DWDM channel and frequency configuration are mutually exclusive.

5

As required, configure the breakout ports.

1. Right-click on the port icon for the breakout port on the equipment tree and choose Properties. The Physical Port (Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

6

To view information about the SFP, click on the Media Adaptor tab.

7

To enable coherent optics, click on the Coherent Optics tab and configure the required parameters.

8

Save your changes and close the form.

END OF STEPS

16.24 To configure Ethernet ports

16.24.1 Note

Before you can configure 802.1X on an Ethernet access port, 802.1X must be enabled on the device and an 802.1X policy must be created and distributed to the device.

See 12.17 “To enable or disable 802.1X” (p. 353) for information about enabling 802.1X on a device. See Chapter 59, “802.1x policies” for information about creating and distributing 802.1X policies.

Depending on the NE and release, the navigation to the Physical Port form or the tabs and parameters displayed may vary.

16.24.2 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required general parameters.

Only one MAC address can be assigned to a port. When a new MAC address is configured while the port is operational, IP issues an ARP, if appropriate, and BPDUs are sent with the new MAC address. A default MAC address is assigned by the system.

If the port is a LAG member, the Hold Time Up (seconds) parameter is configurable only if the port is the primary member. The Hold Time Down (seconds) parameter is not configurable when the port is a LAG member.

When you set a DWDM Channel, the other DWDM parameters are populated. You cannot configure the DWDM channel if frequency is configured under the DWDM Optics tab. DWDM channel and frequency configuration are mutually exclusive.

The DDM Event Suppression parameter is configurable only on ports on SFPs and XFPs optical modular transceivers.

You cannot change the Mode parameter when the Enable DEI parameter is selected.

When you set the Mode parameter to Hybrid, you cannot set the Encapsulation Type parameter to Null.

The MTU value for a port is associated with the port mode and the port encapsulation type. By default, the MTU of a hybrid port is set to the larger of the network and access MTUs.

You cannot deselect the Enable DEI parameter on a port when SAPs on that port are configured with color-aware meters. If you deselect the Enable DEI parameter, the ingress policies assigned to SAPs on the port must be configured with meters whose Color Mode parameter is set to Color Blind. See 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .

Consider the following when you configure the Enable DEI parameter on a hybrid port.

- When the Enable DEI parameter is selected, the network policy assigned to the port must be configured to use DEI as the profile. The Default FC Profile parameter and the Profile parameter on the Ingress Dot1p tab in the policy must be set to the DEI option. See 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .
- You cannot select Enable DEI and assign a DEI-configured network policy in the same configuration process. You must first assign the network policy with the Default FC Profile and Profile parameters set to the DEI option, and apply the change. Then you can subsequently select the Enable DEI parameter.

4

If you set the Mode parameter to Hybrid, configure the required parameters in the Hybrid Ingress Buffer Allocation and the Hybrid Egress Buffer Allocation panels, as required.

5

If you change the Mode parameter from Access to Hybrid and there are SAPs on the port, perform 16.30 “To migrate SAPs from access mode to hybrid mode” (p. 612) to migrate the SAPs.

6

If you are configuring a 7210 SAS Ethernet port, you can add the port to an SHG on supporting NEs. See 16.46 “To create a 7210 SAS SHG” (p. 638) for information about creating an SHG.

Note: A port cannot be added to or deleted from an SHG if it has a SAP configured on it or is a member of a LAG.
A port cannot belong to more than one SHG; you must remove any SHGs from the port before you can add the port to a different SHG.

7

Configure the Named Pool Buffer Policy for the port in the Named Buffer Pool panel.

1. Select an Ingress Pool Policy.
2. Select an Egress Pool Policy.
3. Configure the required parameters.

   Note:

Ensure that the named pool mode is enabled. See 15.67 “To enable named pool mode” (p. 526) in Chapter 15, “Shelf and card object configuration” for more information.

8

Click on the States tab and configure the required parameters.

9

Click on the Policies tab.

Note: The types of policies that you can choose depend on the type of NE, daughter card, and the mode type of the port you are configuring.

- The Network Queue and Accounting Policies panel are only displayed for network and hybrid ports.
- The Port Scheduler panel is only displayed for ports on non-HSMDA daughter cards.
- The HSMDA Egress Scheduler panel is only displayed for ports on HSMDA daughter cards. HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier.

10

Select an accounting policy for the port, if applicable.

11

Configure the Collect Accounting Statistics parameter, if required.

12

If you are configuring an access port on a 7210 SAS-E for the collection of packet forwarding statistics on egress queues, configure the required parameters in the Egress Packets Forwarding panel.

For the 7210 SAS-E, access ports always contain eight queues. You can enable up to eight queues on any port for the collection of packet forwarding statistics. However, only eight counters can be enabled for each 7210 SAS-E device. If no queues are enabled on a port, no accounting statistics are collected for that port.

13

If you are configuring a 7210 SAS Ethernet port, perform 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) , then go to Step 20 .

14

Select a network queue policy for the port, if applicable.

15

Select a 7250 SROS Port QoS policy for the port, if applicable.

16

Select a port scheduler policy for the port, if applicable.

17

Select a port scheduler policy override for the port, if applicable.

1. Click Create in the Port Scheduler Policy Override panel. The Port Scheduler Policy Override form opens.
2. Click on the Override tab.
3. Configure the Override, MAX, CIR, and PIR parameters for each required level.
4. Save your changes and close the form.

18

Configure the Collect Stats parameter in the Port Scheduler Monitor panel, if required.

19

Select a scheduler policy for the port, if applicable.

20

Click on the DDM tab to view digital diagnostics monitoring information, if required.

21

Click on the Media Adaptor tab to view information about the SFP, if required.

22

If you are configuring a 7705 SAR-Wx and the port supports DSL, the DSL tab is displayed. Configure DSL.

1. Click on the DSL tab.
2. Configure the required parameters.
3. Click on the XDSL Line tab.
4. Choose an XDSL line and click Properties. The XDSL Line (Edit) form opens.
5. Configure the Administrative State parameter.
6. Save your changes and close the form.
7. Repeat 2 to 6 to configure more lines, if required.

23

Click on the Ethernet tab.

24

Configure the required parameters.

For access or hybrid ports on the 7210 SAS-X, configure the SAP QoS Marking parameter to enable or disable SAP-based remarking.

For hybrid ports on the 7210 SAS-K12 and 7210 SAS-K30, configure the Egress Network Aggregate Shaper Rate parameter.

The RS-FEC Mode parameter is configurable in the NFM-P for supported physical ports and breakout ports. The RS-FEC Mode is configurable for physical ports under the Ethernet tab. For breakout ports, the RS-FEC Mode parameter is configurable either under the Ethernet tab of the breakout port or under the General tab of its connector port. After the RS-FEC Mode is configured, the RS-FEC Oper Mode can be viewed in the General tab.

On the 7210 SAS-X, configure the Egress Scheduler Mode parameter to choose FC-based or SAP-based scheduling for access ports.

For access ports on supporting 7210 SAS NEs, configure overrides for queues in the Port Access Egress policy assigned to the port; see50.100 “To configure QoS policy overrides on port access egress queues for a 7210 SAS” (p. 1660).

For 7210 SAS NEs, ports can forward multicast data packets to either L2 or L3 interfaces, but not both. Configure the Multicast Ingress parameter according to your forwarding requirements.

Management connectivity between the NFM-P and the 7705 SAR may be lost if a line loopback is applied to an Ethernet port that is carrying in-band management traffic. The line loopback remains in effect until the timer expires and the 7705 SAR removes the loopback.

Statistics are collected for IEEE 1588 PTP timestamps only when the Timestamp Capable parameter is enabled on the 7705 SAR-A, 7705 SAR-Ax, and the 7705 SAR-W.

If you are configuring a 7705 SAR-Wx and the port supports DSL, the EFM-OAM tab appears. Perform the required steps in 90.49 “To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form” (p. 3069) to configure the EFM-OAM tab.

On the 7250 IXR, the Timed Loopback Type parameter must be configured while enabling Swap Mac Address. Internal loopback is only supported in access or hybrid mode, and line loopback is supported in access or network mode. You cannot configure the Timed Loopback Type parameter when the Received Remote Loopback Requests parameter is set to Processed.

Select the Hardware Aggregate Shaper Scheduler policy.

25

Configure a port loopback on a 7210 SAS device.

1. If you are configuring a port loopback with MAC swap, you must first configure a no-service loopback port. See 12.46 “To configure a no-service loopback port on the 7210 SAS” (p. 378) .
2. Set the Type parameter to Internal. See Step 24 .
3. Select a SAP in the SAP panel.
4. In the MAC panel, configure the required parameters.

26

For supported 7210 SAS devices, associate the port with an operational group. See 12.9 “To configure a 7210 SAS operational group” (p. 346) for information about 7210 SAS operational groups.

Perform one of the following:

a. For an L2 uplink port, assign the port to an operational group. Select an operational group in the Oper Group panel.

b. For an access port, choose an operational group to monitor. Select an operational group in the Oper Group panel.

27

If you are configuring an Ethernet port on a 2-port ring MDA on a 7705 SAR-8 or 7705 SAR-18, select a VLAN filter.

28

If you are configuring a 7705 SAR Ethernet port, select a Shaper policy in the Shaper Policy panel as required.

Note: When the port mode is access, the Shaper policy can be associated with the access egress of the port. When the port mode is hybrid, the Shaper policy can be associated with the access egress and network egress.

29

If you are configuring a 7705 SAR Ethernet port, configure Ethernet Bandwidth Notification as required. See 16.48 “To configure Ethernet Bandwidth Notification on a 7705 SAR Ethernet port” (p. 640) for more information.

Note: ETH-BN can only be enabled if the egress rate is configured.

30

To transparently forward LACP packets, enable the LACP Tunnel parameter, as required.

31

Configure the Dot1x Tunneling Enabled parameter, as required.

32

In the Synchronous Ethernet panel, configure the required parameters.

33

In the E-LMI panel, configure the required parameters.

34

In the CRC Failure panel, configure the required parameters.

35

In the Oper Group panel, configure the required parameters.

36

Select the sFlow check box to enable sFlow on the Ethernet port. See 12.29 “To configure sFlow on an NE” (p. 364) for more information.

37

If the port is an access or hybrid port with QinQ or dot1q encapsulation, or a network port with dot1q encapsulation, the MEPs tab is displayed. Add a facility MEP, if required.

Note: Only one facility MEP can be configured on a port.

1. Click on the MEPs tab.
2. Click Create. The MEP (Create) form opens.
3. Select a global MEG for the MEP.

   Note:

   For port-based facility MEPs, the Maintenance Domain must be set for a Level of 0 or 1. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) for information on creating the MD.
4. Configure the required parameters.

   Only the parameters that are supported on the NE, daughter card, or port are displayed.

   You cannot enable the collection of LMM statistics on a LAG facility MEP or on a tunnel facility MEP if a service MEP is already configured to collect the LMM statistics information.

   To enable fault propagation on a LAG facility MEP or on a tunnel facility MEP, the Facility Fault Notify parameter must be enabled.

   To link the facility tunnel status to the SAP, the Facility VLAN ID parameter must match the outer encapsulation value of the SAP.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.
5. If the MD for the MEP has a Name Type of none the Y.1731 Tests tab is configurable. Click on the Y.1731 Tests tab and configure the required parameters.
6. Save your changes and close the form.

38

Click on the EFM-OAM tab to configure and initiate the 802.3ah EFM OAM diagnostic as described in 90.49 “To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form” (p. 3069) . Otherwise, continue to Step 40 .

39

To configure Ethernet link monitoring on the port, perform 90.53 “To configure link monitoring on an Ethernet port” (p. 3080) .

40

Click on the 802.1x Port and 802.1x Port Authenticator tabs, and configure the required parameters.

You can enable the Tunneling parameter only when:

- the port is in access mode
- the Controlled Port Control parameter is set to Force Authorized
- the port is not used as a SAP (except in an Epipe service)

Configure the RADIUS Server Policy Authentication and RADIUS Server Policy Accounting parameters to associate a RADIUS server policy to the Ethernet port.

For 7250 IXR-x1/xs/x3 NEs, Release 21.5.R1 and later, and for 7250 IXR-6d/6dl NEs, Release 22.7.R1 and later, you can only enable the Per Host Authentication State parameter when the forwarding path options for Ethernet 802.1x host authentication are enabled on the NE Properties form. See 12.57 “To configure forwarding path options or resource allocation on a 7250 IXR” (p. 388) .

You can click on the Per Host Authentication sub-tab to configure Allowed Source MACs and view authenticated hosts. The configured MAC addresses do not take part in authentication and are always allowed.

41

Go to 16.37 “To add a queue group to an Ethernet port” (p. 627) to add a queue group to the port.

42

If the port is an HSMDA access port, the Egress Secondary Shapers tab is present. Add an egress secondary shaper, if required.

Note: An HSMDA port is initially created with a default egress secondary shaper. The default shaper cannot be deleted, but it can be modified.
Non-default egress secondary shapers can only be created if the port’s Mode parameter is not set to the network option.
You cannot switch the Mode of an HSMDA port to network if there are non-default egress secondary shapers on that port. An error message will be issued. You must delete any non-default egress secondary shapers from the port before switching its Mode to network.

1. Click on the Egress Secondary Shapers tab.
2. Click Create. The HSMDA Egress Secondary Shaper (Create) form opens.
3. Configure the required parameters.
4. Click on the Classes tab.
5. Configure the Class Rate (kbps), Class Monitor Threshold (Kbytes) and Class Burst Threshold (bytes) parameters for Classes 1 through 8, as required.
6. Save your changes and close the form.

43

Click on the LLDP tab and configure the required parameters on the Nearest Bridge, Nearest Customer, and Nearest Non TPMR sub-tabs.

Care must be taken when configuring the Port ID SubType parameter, as the setting may affect the ability to build the Layer 2 topology map using LLDP. See the NE documentation for more information.

44

Click on the Remote Peers sub-tab under the LLDP tab to search for and display LLDP remote peers associated with the port. These remote peers are used to determine the physical topology of the network.

45

Configure a QoS pool, if required.

1. Click on the QoS Pool tab.
2. Choose a QoS pool from the list and click Properties. The QoS Pool (Edit) form opens.
3. Select a slope policy for the QoS pool.
4. Configure the required parameters.
5. Save your changes and close the form.

46

Click on the Override Policy Items tab to configure HSMDA overrides, if required. See 16.40 “To configure an HSMDA override” (p. 632) for more information.

47

To configure the center frequency and view the minimum/maximum frequencies and supported frequency grid for a particular optical module, click the DWDM Optics tab and configure the required parameters.

You cannot configure the frequency here if the DWDM channel is configured under the General tab. DWDM channel and frequency configuration are mutually exclusive.

48

To configure a tunable optic DWDM MDA or IMM on a 7750 SR or 7450 ESS, click on the Optical Transport Channel Unit tab and configure the required parameters.

49

Click on the OTU Alarms tab and select the OTU alarms for the Configured Alarms parameter.

50

Select a DWDM channel. The tabs and parameters that appear depend on the supported NE.

51

Click on the Optical tab and configure the required parameters.

52

Click on the Optical Amplifier tab and configure the Configured Alarms parameter.

53

Click on the Optical Tunable Dispersion Compensation Module tab and configure the required parameters.

54

Save your changes and close the form.

55

To configure coherent optics options for a 1-Port 100GE OTU4 DWDM tunable optic IMM on a 7750 SR or 7450 ESS, click on the Coherent Optics tab and configure the required parameters.

56

Click on the Operational Status tab to view information about the state of the optical port.

57

Save your changes and close the form.

END OF STEPS

16.25 To configure LLDP-MED

16.25.1 Before you begin

LLDP-MED (Media Endpoint Discovery) is an extension to LLDP to support interoperability between VoIP endpoint devices and other networking end-devices. LLDP-MED is focused mainly on discovery running between network devices and endpoints such as IP phones.

LLDP-MED configuration is supported only on 7210 SAS-S/Sx from Release 20.9 R1 onwards and on 7210 SAS-Dxp from Release 22.3 R1 onwards.

16.25.2 Steps

1

On the equipment tree, expand Network→NE→Shelf→Card Slot n→Daughter Card Slotn→Port n/n/n.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Click on the LLDP tab and then on Nearest Bridge tab.

4

Click on the LLDP MED sub-tab and configure the required parameters.

Associate LLDP network policies. To create LLDP MED Network Policy, See 16.26 “To configure LLDP MED Network Policy” (p. 608)

5

Save your changes and close the forms.

END OF STEPS

16.26 To configure LLDP MED Network Policy

16.26.1 Steps

1

Choose Administration→LLDP MED Network Policy from the NFM-P main menu. The LLDP MED Network Policy form opens.

2

Configure the required parameters.

3

Save your changes and close the form.

END OF STEPS

16.27 To configure a cellular port on a 7705 SAR-Hm

16.27.1 Before you begin

SIM cards must be installed in the 7705 SAR-Hm before the cellular ports can be configured.

16.27.2 Steps

1

On the equipment tree, expand Network→7705 SAR-Hm→Shelf→Card Slot 1→Daughter Card Slot 1→Port 1/1/1.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required general parameters.

The Mode and Encapsulation Type parameters are set to Network and Cellular, respectively, by default. These parameters cannot be changed.

4

Click on the States tab and configure the required parameters.

5

Click on the Cellular tab and configure the required parameters.

1. A PDN profile must be configured. See Chapter 69, “PDN profile policies” for more information.
2. To configure SIM switchover behaviour, click on the SIM Cards tab and configure the required parameters.

   The length of time configured for the Fail Duration (Min) parameter specifies when to perform an automatic failover from one SIM to the other. Set the Fail Duration (Min) parameter to a value lower than the Down Recovery Interval (Min) parameter on the cellular MDA, or a node reboot will happen instead of a SIM switchover.
3. View port bearer and port traffic flow template details by performing the following:
   - Click on the PortBearer tab, select a bearer in the list, and click Properties. The Cellular Port Bearer form opens.
   - Click on the PortTFT tab, select a traffic flow template from the available list, and click Properties. The Port Traffic flow Template (TFT) form opens.
4. To enable the 7705 SAR-Hmc to operate as a CBSD category A/B device in the CBRS spectrum band, click on the CBSD tab and configure the required parameters.

6

Click on the Statistics tab and perform one of the following:

a. Click Collect to collect performance statistics data on demand. The collected statistics are listed on the form.

b. Click Collect All to collect one on-demand statistics record for each statistic type that the object supports. The collected statistics are listed on the form.

7

Save your changes and close the form.

END OF STEPS

16.28 To configure a WLAN port on a 7705 SAR-Hm

16.28.1 Steps

1

On the equipment tree, expand Network→7705 SAR-Hm→Shelf→Card Slot 1→Daughter Card Slot 4→Port 1/4/1.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required general parameters.

4

Click on the WLAN tab and then click on the Network tab.

5

Click Create. The Wlan Network (Edit) form opens.

Configure the required parameters and click OK. The Access Point is created.

6

If required, configure the Access Point parameters:

1. Select the Access Point and click Properties.
2. Configure the parameters and click OK.
3. Click Apply.

7

If required, configure the Station parameters.

8

Save your changes and close the forms.

END OF STEPS

16.29 To change the port mode

Note: Save the device configuration before you perform this procedure. See the appropriate NE documentation for more information.

16.29.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Change the Mode parameter as required and click Apply.

4

If you are changing the mode from Access to Network and services are configured on the port, warnings are displayed in a dialog box, depending on the permissions that are associated with your user role. Perform one of the following:

a. The default warning dialog indicates that the object has dependencies that will be deleted if you apply the port mode change.
   1. Click on View Dependencies button to see the dependencies (services and interfaces) associated with the port.
   2. Click OK to confirm the configuration change.

      Any services and interfaces associated with the port, including SAPs, are deleted, and the mode for the port is changed to Network.

      Note:

      Changing the port mode from Access to Network without first removing the dependencies results in a TopologyMisconfigured alarm.

b. For users who have the port.RestrictNodeConfigModify permission associated with their role, a warning form appears, indicating that the object has dependencies that must be removed before proceeding.
   1. Click on View Dependencies button to see the dependencies (services and interfaces) associated with the port.
   2. Click OK to confirm the configuration change.
   3. Remove the services or interfaces associated with the port. You can manually delete them or move them to other ports as described in this chapter.
   4. Repeat Step 2 and Step 3 to change the port mode to Network.

5

Save your changes and close the form.

END OF STEPS

16.30 To migrate SAPs from access mode to hybrid mode

Note: Save the device configuration before you perform this procedure. See the appropriate NE documentation for more information.

16.30.1 Steps

1

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on an Ethernet port or SC-LAG port icon and choose Properties. The Physical Port (Edit) form opens.

3

Change the Mode parameter from Access to Hybrid. A dialog box opens.

4

Click OK to confirm the configuration change.

5

Perform an admin save.

6

Save your changes and close the form.

End of steps

16.31 To configure the NFM-P to retain non-default port MTU values

16.31.1 Before you begin

The MTU value for a port is associated with the port mode, such as access or network, and the port encapsulation type. If you change the mode or encapsulation type of a port, the MTU value changes to the associated default. If you do not want the port MTU values to change, you can use this procedure to configure the NFM-P to retain each port MTU value, regardless of a mode or encapsulation type change.

Note: You must perform this procedure on each main server in a redundant NFM-P deployment.

16.31.2 Steps

1

Log in to the NFM-P main server station as the nsp user.

2

Navigate to the /opt/nsp/nfmp/server/nms/config directory.

3

CAUTION

Service Disruption

Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file. Modifying the nms-server.xml file can have serious consequences that can include service disruption.

Open the nms-server.xml file using a plain-text editor.

4

Add the following line to the end of the file:

```
<customMTURefresh refresh="true"/>
```

5

Save and close the nms-server.xml file.

6

Open a console window.

7

Navigate to the /opt/nsp/nfmp/server/nms/bin directory.

8

If the main server is a standalone server, or the primary server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash read_config ↲
```

The main server reads the nms-server.xml file and puts the configuration change into effect.

9

If the main server is the standby server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash force_restart ↲
```

The main server restarts and puts the configuration change into effect.

10

Close the console window.

END OF STEPS

16.32 To copy or move L2 SAPs between ports

16.32.1 Before you begin

Note: If the source SAP has an MEP configured, and a CFM test is running on the service that has the new MEP, the test must be stopped and executed again to make the new MEP active.

Before you perform a SAP move operation, archive all existing test results associated with the source SAP MEP.

You can also use the NFM-P clipboard to copy or move SAPs. See Chapter 1, “NFM-P GUI” for information about using the NFM-P clipboard.

16.32.2 Steps

1

For the source NE, disable automatic configuration backup if it is enabled in the backup policy assigned to the NE. Configure the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

2

CAUTION

Service Disruption

An unsuccessful SAP copy or move may result in the deletion of one or more source SAPs from the source NE. The backed-up configuration file is required to restore the configuration on the source NE if a SAP copy or move operation fails.

On the source NE, use a CLI to back up the device configuration to a file. Ensure that you use a unique name for the backup file to ensure that it is not accidentally overwritten. Enter the following at the CLI prompt:

`admin save file_name` ↵

where *file_name* is a name that differs from the existing configuration file names on the NE and identifies this as the configuration in effect before the copy or move operation

See the appropriate device documentation for more information about saving the device configuration to a file.

3

For the destination NE, disable automatic configuration backup if it is enabled in the backup policy assigned to the NE. Configure the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

4

On the destination NE, use a CLI to back up the device configuration to a file. Ensure that you use a unique name for the backup file to ensure that it is not accidentally overwritten. Enter the following at the CLI prompt:

`admin save file_name` ↵

where *file_name* is a name that differs from the existing configuration file names on the NE and identifies this as the configuration in effect before the copy or move operation

See the appropriate device documentation for more information about saving the device configuration to a file.

5

CAUTION

Service Disruption

Nokia strongly recommends that you monitor the NE deployments and the alarm window during a SAP copy or move operation to ensure that faults that may arise during the operation are immediately detected.

The Incomplete Deployments tab of the Deployment form notifies the NFM-P operator of configuration changes that are not successfully deployed to the NE, and the NFM-P alarm window alerts the operator to object failures on the NE that are related to the copy or move operation.

Open the Deployment form to monitor deployments as they occur during the SAP copy or move operation. Choose Administration→NE Maintenance→Deployment from the NFM-P main menu. The Incomplete Deployments tab is displayed.

6

Choose Tools→Copy/Move→SAP. The SAP Copy/Move form opens.

7

Arrange the forms on the GUI so that the SAP Copy/Move form, the Deployment form, and the alarm window are all shown.

8

On the SAP Copy/Move form, click on the Result tab and then click Result Export Path. The Result Export Path form opens.

9

Specify the file name and location in which to save a text file that contains the results of the SAP copy or move operation.

10

Click Set. The Result Export Path form closes.

11

On the SAP Copy/Move form, click on the General tab.

12

Configure the Current Mode parameter as L2 Access Interface.

13

Configure the Service Type parameter.

14

Select a source node and a source port in the Source panel.

15

Configure the Outer Encap Value Start and Outer Encap Value End parameters, and the Inner Encap Value Start and Inner Encap Value End parameters, if applicable.

16

Select a destination node and a destination port in the Destination panel.

17

Configure the required parameters in the Execution Details panel.

18

Click Execute to begin the copy or move operation. The Execution Result panel displays the execution state of the copy or move operation and the number of successful and failed operations based on the input criteria.

Note: When you copy or move a SAP on an L2 interface and there is a non-zero value assigned to the PIM Snooping parameter Maximum Number of Groups, the value is not copied or moved to the new location. You must manually configure the Maximum Number of Groups parameter in the new location.
The Succeeded state only specifies that the SAP copy or move operation is successfully committed to the database. After the SAP copy or move operation is executed, you must monitor the Deployment form for any deployment errors.

19

Monitor the Deployments form and the alarm window for any deployment failures and faults related to the SAP copy or move operation. Click Refresh on the Deployments form to update the list, if required.

20

Click on the Result tab to view a list of the successful and failed operations as identified by the NFM-P.

You can double click the result to open it in a new window and view the details.

21

Click Refresh Result to view the most recent results.

22

To view the information for an individual SAP copy or move operation, move the mouse pointer over the Message field of the operation.

23

If a copy or move operation fails, you can restore the previous source NE configuration using the backup configuration file created in Step 2 .

1. Open a CLI on the source NE.
2. Enter the following at the command prompt to load the saved configuration file:

   exec *file_name* ↵

   where *file_name* is the name of the configuration file that was specified in Step 2

   The SAPs that were deleted during the unsuccessful copy or move operation are restored on the NE.
3. Enter the following at the command prompt to save the NE configuration:

   admin save ↵

24

If a copy or move operation partially fails, you can restore the previous destination NE configuration using one of the following methods.

Note: If the copy or move operation fails completely, there is no need to restore the previous configuration because no SAPs are created on the destination NE.

a. Restore the backed-up configuration file.
   1. Open a CLI on the destination NE.
   2. Enter the following at the command prompt to load the saved configuration file:

      exec *file_name* ↵

      where *file_name* is the name of the configuration file specified in Step 4

      The previous configuration is restored on the NE.
   3. Enter the following at the command prompt to save the NE configuration:

      admin save ↵

b. Manually delete the SAPs that the copy or move operation created on the destination NE.

25

Restore the previous automatic configuration backup setting in the backup policy of the source NE by configuring the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

26

Restore the previous automatic configuration backup setting in the backup policy of the destination NE by configuring the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

27

Close the form.

END OF STEPS

16.33 To copy or move L2 access interface SAPs between services

16.33.1 Before you begin

You can move L2 access interface SAPs between VPLS, MVPLS and VLL services.

Note the following:

- You can move the source SAP only to another service, and not within the same service. You cannot move the source SAP to a different site; it can be moved only to a different service on the same site.
- The source and destination service types must be the same. The source and destination site types (Site, B-Site or I-Site) must be the same.
- The source and destination services must have the same Customer ID.

VPLS L2 access interfaces can be used for MSAPs, access interfaces for SPB in control mode, and throughput tests. When you move an L2 interface, the NFM-P deletes the SAP from the source site. The NFM-P does not check that the source site has consistent configuration after the move, and for example, a throughput test might reference nonexistent SAPs.

16.33.2 Steps

1

Choose Manage→Services→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS, MVPLS, or VLL service and click Properties. The *Service_Type* Service (Edit) form opens.

3

On the service tree, expand Sites→Site→L2 Access Interfaces.

4

Right-click on the L2 access interface that you want to move and choose Move to Another Service. The Select Services form opens with a list of eligible target services.

5

Select a service. The NFM-P moves the L2 access interface SAP to the selected service and removes the service from the from the source service.

6

Close the forms.

END OF STEPS

16.34 To move L3 SAPs within or between ports or LAGs on the same NE

16.34.1 Before you begin

You can also use the NFM-P clipboard functionality to move SAPs. See Chapter 1, “NFM-P GUI” for information about using the NFM-P clipboard.

Note: L3 access interfaces are not bound to a port; for example, loopback interfaces cannot be moved.

16.34.2 Steps

1

For the NE, disable automatic configuration backup if it is enabled in the backup policy assigned to the NE. Configure the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

2

On the NE, use a CLI to back up the device configuration to a file. Ensure that you use a unique name for the backup file to ensure that it is not accidentally overwritten. Enter the following at the CLI prompt:

`admin save` ***`file_name`*** ↵

where *file_name* is a name that differs from the existing configuration file names on the NE and identifies this as the configuration in effect before the move operation

See the appropriate device documentation for more information about saving the device configuration to a file.

3

CAUTION

Service Disruption

Nokia strongly recommends that you monitor the NE deployments and the alarm window during a SAP move operation to ensure that faults that may arise during the operation are immediately detected.

The Incomplete Deployments tab of the Deployment form notifies the NFM-P operator of configuration changes that are not successfully deployed to the NE, and the NFM-P alarm window alerts the operator to object failures on the NE that are related to the move operation.

Open the Deployment form to monitor deployments as they occur during the SAP move operation. Choose Administration→NE Maintenance→Deployment from the NFM-P main menu. The Incomplete Deployments tab is displayed.

4

Choose Tools→Copy/Move→SAP. The SAP Copy/Move form opens.

5

Arrange the forms on the GUI so that the SAP Copy/Move form, the Deployment form, and the alarm window are all shown.

6

Click on the Result Export Path tab.

7

Specify the file name and location in which to save a text file that contains the results of the SAP move operation.

8

Click Set. The Result Export Path form closes.

9

Configure the Current Mode parameter as L3 Access Interface.

10

Configure the Service Type parameter.

11

Select a source node and a source port in the Source panel.

12

Configure the Outer Encap Value Start and Outer Encap Value End parameters, and the Inner Encap Value Start and Inner Encap Value End parameters, if applicable.

13

Select a destination port in the Destination panel.

14

Configure the required parameters in the Execution Details panel.

15

Click Execute to start the move operation. The Execution Result panel displays the execution state of the move operation and the number of successful and failed operations based on the input criteria.

16

Monitor the Deployments form and the alarm window for any deployment failures and faults related to the SAP move operation. Click Refresh on the Deployments form to update the list, if required.

17

Click on the Result tab to view a list of the successful and failed operations as identified by the NFM-P.

You can double click the result to open it in a new window and view the details.

18

Click Refresh Result to view the most recent results.

19

To view the information for an individual SAP move operation, move the mouse pointer over the Message field of the operation.

20

If a move operation fails, you can restore the previous NE configuration using the backup configuration file created in Step 2 .

1. Open a CLI on the NE.
2. Enter the following at the command prompt to load the saved configuration file:

   `exec file_name` ↵

   where *file_name* is the name of the configuration file specified in Step 2

   The SAPs that were deleted during the unsuccessful move operation are restored on the NE.
3. Enter the following at the command prompt to save the NE configuration:

   `admin save` ↵

21

Restore the previous automatic configuration backup setting in the backup policy of the NE by configuring the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

22

Close the forms.

END OF STEPS

16.35 To move L3 subscriber interface SAPs between ports on the same NE

16.35.1 Before you begin

You can also use the NFM-P clipboard functionality to move SAPs. However, you can only move all of the SAPs in a group interface; you cannot move a subset of SAPs. See Chapter 1, “NFM-P GUI” for information about using the NFM-P clipboard.

16.35.2 Steps

1

For the NE, disable automatic configuration backup if it is enabled in the backup policy assigned to the NE. Configure the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

2

On the NE, use a CLI to back up the device configuration to a file. Ensure that you use a unique name for the backup file to ensure that it is not accidentally overwritten. Enter the following at the CLI prompt:

`admin save file_name` ↵

where *file_name* is a name that differs from the existing configuration file names on the NE and identifies this as the configuration in effect before the move operation

See the appropriate device documentation for more information about saving the device configuration to a file.

3

CAUTION

Service Disruption

Nokia strongly recommends that you monitor the NE deployments and the alarm window during a SAP move operation to ensure that any faults that arise during the operation are immediately detected.

The Incomplete Deployments tab of the Deployment form notifies the NFM-P operator of configuration changes that are not successfully deployed to the NE, and the NFM-P alarm window alerts the operator to object failures on the NE that are related to the move operation.

Open the Deployment form to monitor deployments as they occur during the SAP move operation. Choose Administration→NE Maintenance→Deployment from the NFM-P main menu. The Incomplete Deployments tab is displayed.

4

Choose Tools→Copy/Move→SAP. The SAP Copy/Move form opens.

5

Arrange the forms on the GUI so that you can view the entire SAP Copy/Move form, the Deployment form, and the alarm window.

6

Click on the Result Export Path tab. The Result Export Path form opens.

7

Specify the file name and location in which to save a text file that contains the results of the SAP move operation.

8

Click Set. The Result Export Path form closes.

9

Configure the Current Mode parameter as L3 Subscriber Interface SAP.

10

Configure the Service Type parameter.

11

Select a source node and a source port in the Source panel.

12

Select a destination port in the Destination panel.

13

Configure the required parameters in the Execution Details panel.

14

Click Execute to start the move operation. The Execution Result panel displays the state of the move operation, and the number of successful and failed operations based on the input criteria.

15

Monitor the Deployments form and the alarm window for any deployment failures and faults related to the SAP move operation. Click Refresh on the Deployments form to update the list, if required.

16

Click on the Result tab to view a list of the successful and failed operations, as identified by the NFM-P.

You can double click the result to open it in a new window and view the details.

17

Click Refresh Result to view the most recent results.

Note: If one SAP fails, all SAPs under the same group interface fail because the SAPs can only move together. In this case, there is insufficient information to determine which SAP caused the problem using the NFM-P.

18

To view information about an individual SAP move operation, move the mouse pointer over the Message field of the operation.

19

If a move operation fails, you can restore the previous NE configuration using the backup configuration file created in Step 2 .

1. Open a CLI on the source NE.
2. Enter the following at the command prompt to load the saved configuration file:

   exec *file_name* ↵

   where *file_name* is the name of the configuration file specified in Step 2

   The SAPs that were deleted during the unsuccessful move operation are restored on the NE.
3. Enter the following at the command prompt to save the NE configuration:

   admin save ↵

20

Restore the previous automatic configuration backup setting in the backup policy of the NE by configuring the Auto Backup Scheme parameter appropriately. See 23.4 “To configure a backup policy” (p. 746) for configuration information.

21

Close the forms.

End of steps

16.36 To configure bandwidth CAC on an access SAP for services or a LAG

16.36.1 Before you begin

You can configure the bandwidth CAC function on a port or LAG, based on the admin bandwidth configured on a SAP and on an associate port or LAG. In addition, a booking factor is provided to allow over/under booking of the sum of the SAP bandwidth compared to the port/LAG bandwidth. See 16.7 “Configuring access SAP bandwidth CAC” (p. 577) for more information.

16.36.2 Steps

To configure bandwidth CAC for services

1

Choose Manage→Services→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search and choose a VPLS, VPRN, MVPLS, EPIPE, IPIPE, or IES service and click Properties. The *Service_Type* Service (Edit) form opens.

3

On the service tree, expand Sites→Site→L2 Access Interface or L3 Access Interface.

4

Right-click on the L2 access interface or L3 Access Interface that you want to configure the bandwidth CAC for and click Properties. The *Service_Type* L2 or L3 Access Interface form opens.

5

Click on the Port tab and configure the Mode parameter to Access or Hybrid mode to view the Access Bandwidth panel.

Note: By default, the No Bandwidth parameter is configured in the Access Bandwidth panel which excluded the port from the CAC function.

6

As required, configure the Bandwidth (kbps) and Book Factor (Percentage) parameters.

7

Save your changes and close the forms.

To configure bandwidth CAC for LAGs

8

Choose Equipment from the navigation tree view selector.

9

On the equipment navigation tree, expand Network→*NE*→Logical Groups→LAGs→LAG *n*.

10

Right-click on the LAG *n* object that you want to configure the bandwidth CAC for and choose Properties. The LAG (Edit) form opens.

11

On the General tab, configure the Mode parameter to Access or Hybrid mode to access the bandwidth CAC panel.

12

Click on the Access tab, and as required, configure the Bandwidth (kbps) and Book Factor (Percentage) parameters.

Note: By default, the No Bandwidth parameter is configured in the Access Bandwidth panel which excluded the LAG from the CAC function.

Note: When a LAG is configured, the admin bandwidth and booking factor on its constituent ports are ignored.

13

Save your changes and close the forms.

END OF STEPS

16.37 To add a queue group to an Ethernet port

16.37.1 Before you begin

You can create a port queue group on an Ethernet port after creating an Ingress/Egress Queue Group Template policy.

Note: You must use the same name for the port queue group and Ingress/Egress Queue Group Template policy.

16.37.2 Steps

1

On the equipment tree, navigate to the port object where you need to add a queue group. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required Ethernet port icon and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Ethernet tab.

4

Perform one of the following:

a. Click on the Network Egress Queue Group tab.

b. Click on the Access Ingress Queue Group tab.

c. Click on the Access Egress Queue Group tab.

Note: The Network Egress Queue Group tab is available if queue groups are supported on the port, depending on the NE release and IOM type. See 71.1.3 “Port queue groups” (p. 1963) in Chapter 71, “Queue groups” for more information about supported devices and IOM types.

5

Click Create. The Queue Group (Create) form opens.

6

Select a queue group template policy.

7

Configure the required parameters.

8

If you are adding a network egress queue group, go to Step 9 . Otherwise, continue to Step 10.

9

Select a policer control policy.

10

Select an accounting policy.

Note: Only the accounting policies with the Type parameter configured to one of the following options are listed:

- Queue Group Octets
- Queue Group Packets
- Combined Queue Group

11

Configure the Collect Accounting Statistics parameter, if required.

12

Select a scheduler policy.

13

If you are adding an access ingress queue group, go to Step 20 . Otherwise, continue to Step 14 .

14

Configure the required parameters.

The Aggregate Rate Limit (kbps) and Frame-Based Accounting parameters are configurable only if you did not select a scheduler policy. The Frame-Based Accounting parameter is configurable only if you configure the Aggregate Rate Limit (kbps) parameter.

15

If you are adding a network egress queue group (including queue groups for HSMDA), go to Step 20 . Otherwise, continue to Step 16 .

16

Click on the Host Matching tab.

Note: The Host Matching tab is available only if the Instance ID is set to 1.

17

Click Create. The Host String (Create) form opens.

18

Configure the Host String parameter.

19

Click OK. The Host String (Create) form closes and the information on the Host Matching tab is updated.

20

Click on the Overrides: *Type* Queue tab, where *Type* will be Network Egress, Access Egress, or Access Ingress, depending on your selection in Step 4 .

21

Click Create. The Queue Override (Create) form opens.

22

Select a Queue ID.

23

Click on the Override tab.

24

If you are configuring an access egress queue group queue override for HSMDA, then go to Step 28 .

25

Configure the required override parameters.

The Weight and CIR Weight parameters are only applicable to access egress queue group overrides.

The parameters are configurable when the Override check box is enabled.

26

In the Stats panel, configure queue depth monitoring. For more information, see 50.21 “Queue depth monitoring” (p. 1527).

27

Go to Step 30 .

28

Configure the required parameters.

The parameters are configurable when the associated Override check box is enabled.

29

If you need to configure an override for a Slope Policy, enable the associated Override check box and select a slope policy.

30

Save your changes and close the forms.

END OF STEPS

16.38 To configure queue group scheduler overrides

16.38.1 Before you begin

Perform this procedure to override the scheduler values for access ingress or access egress queue groups. The queue groups must already exist. To create queue groups, see 16.37 “To add a queue group to an Ethernet port” (p. 627).

16.38.2 Steps

1

On the equipment tree, navigate to the port object where you need to override the queue group scheduler values. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required Ethernet port icon and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Ethernet tab.

4

Perform one of the following:

a. Click on the Access Ingress Queue Group tab.

b. Click on the Access Egress Queue Group tab.

5

Choose a queue group and click Properties. The Queue Group (Edit) form opens.

6

Click on the Overrides: Ingress Scheduler tab or Overrides: Egress Scheduler tab, depending on your selection in Step 4 .

7

Choose a scheduler and click Create Override. The Access Ingress (or Access Egress) Queue Group Scheduler Override (Create) form opens.

8

Click on the Override tab and configure the required parameters.

9

Save your changes and close the forms.

END OF STEPS

16.39 To configure SONET ports

16.39.1 Steps

1

On the equipment tree, navigate to the port object where you want to configure a SONET port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters.

The DDM Event Suppression parameter is configurable only on ports on SFPs and XFPs optical modular transceivers.

4

Click on the States tab and configure the Administrative State parameter.

5

Click on the Channels tab to configure the SONET/SDH channel.

a. Click Create to create a new channel.

b. Choose a channel from the list and click Properties to view and edit the channel parameters.

c. Choose a channel from the list and click Delete to remove a channel from the port.

6

Click on the SONET tab and configure the required parameters.

7

Click on the SONET Monitoring tab and configure the required parameters.

8

Click on the SONET Overhead tab and configure the required parameters.

Enter a J0 byte that identifies the circuit. This byte is inserted continuously at source. This can be checked against the expected value by the receiver. If no byte is entered, then null is used. The parameter is configurable when the SONET Section Trace Mode is set to Byte.

The J0 String parameter is configurable when the SONET Section Trace Mode is set to String.

9

Save your changes and close the form.

END OF STEPS

16.40 To configure an HSMDA override

16.40.1 Steps

1

On the equipment tree, navigate to the port object where you want to configure a SONET port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

Select an HSMDA card. HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier.

2

Click on the Override Policy Items tab.

3

Click Create. The Port Egress HSMDA Scheduler Policy Override (Create) form opens.

4

Click on the Override tab and configure the required parameters.

The Rate (Mbps) parameter is configurable when a Group is not specified for the applied HSMDA scheduler policy. The Weight in Group parameter is only configurable when Group 1 or Group 2 is specified for the HSMDA scheduler policy.

5

Save your changes and close the forms.

END OF STEPS

16.41 To configure TDM DS3 ports

16.41.1 Steps

1

On the equipment tree, navigate to the port object where you want to configure a TDM DS3 port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click on the States tab and configure the Administrative State parameter.

5

Click on the DS3/E3 tab and configure the required parameters.

6

Save your changes and close the form.

END OF STEPS

16.42 To configure serial ports

16.42.1 Purpose

Perform this procedure to configure serial ports on a 12 × Serial Data card on the 7705 SAR.

16.42.2 Steps

1

On the equipment tree, navigate to the port object where you want to configure a serial port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a port icon and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters.

4

Click on the States tab and configure the Administrative State parameter.

5

Save your changes and close the form.

END OF STEPS

16.43 To configure PW ports

16.43.1 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the PW Ports tab. A list of existing PW ports is displayed.

3

Click Create to create a PW port or choose a port and click Properties. The PW Port (Create|Edit) form opens.

4

Configure the required general parameters.

5

Configure node redundancy for the PW port, if required.

1. Click Node Redundancy. The Manage Node Redundancy form appears, preconfigured for MC peer groups.
2. Click Search and configure an existing MC peer group, or click Create and configure a new MC peer group. See 40.4 “To configure an MC peer group” (p. 1330) .

   The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add PW port sync tags.

6

Save the changes and close the forms.

END OF STEPS

16.44 To configure a 7210 SAS-M channelized TDM DS1 or E1 port

16.44.1 Steps

1

On the equipment tree, navigate to the port object where you want to configure a 7210 SAS-M channelized TDM DS1 or E1 port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a 7210 SAS-M port object that contains a channelized DS1 or E1 and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters.

4

Select a split horizon group.

Note: To create or delete a 7210 SAS SHG, see 16.46 “To create a 7210 SAS SHG” (p. 638) .

5

Click on the DS1/E1 tab and configure the required parameters.

The Line Impedance parameter only appears when the Port Type parameter is set to E1.

6

Click on the QoS Pool tab.

7

Choose a QoS pool from the list and click Properties. The QoS Pool (Edit) form opens.

8

Select a slope policy.

9

Save your changes and close the form.

10

Click on the States tab and configure the Administrative State parameter.

11

Save your changes and close the forms.

END OF STEPS

16.45 To assign QoS policies to a 7210 SAS Ethernet port

16.45.1 Before you begin

Perform this procedure to assign QoS policies to a 7210 SAS Ethernet port, or to enable table-based ingress classification on the port for supporting 7210 SAS NEs.

The policies that can be assigned to a 7210 SAS Ethernet port vary depending on the port mode.

Table 16-5 7210 SAS QoS policies and port modes

7210 SAS QoS policy	Supported port mode
Port scheduler policy	All port modes
Port access egress policy	Access
Network policy (of port type)	L2 uplink, network, hybrid
Network queue policy	L2 uplink, network, hybrid
7210/7250 DSCP classification policy	Access

Support for port modes and policies also varies depending on the chassis type. For more information, see 50.23 “7210 SAS QoS policies” (p. 1528) and the NE documentation.

Slope policies for Ethernet ports are assigned to QoS pools (buffer pools) on the port; see 16.24 “To configure Ethernet ports” (p. 599) .

Network policies of network interface type are assigned to L3 network interfaces; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863) .

When you assign a 7210 and 1830 network policy that has an associated MPLS LSP-EXP classification policy, the port must be in network mode.

For ports on 7210 SAS-K12 NEs, when you assign a 7210 and 1830 network policy that has an associated remarking policy, the port must be in network mode if the remarking policy is one of the following types:

- DOT1P-LSP-EXP
- DOT1P-LSP-EXP-DSCP
- DSCP-LSP-EXP
- LSP-EXP

16.45.2 Steps

1

On the equipment tree, navigate to a configured Ethernet port object on a 7210 SAS device. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a configured Ethernet port object and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Policies tab, then click on the 7210 Specific tab.

Note: The policies that can be assigned vary depending on the chassis type and port mode.

4

Select a port scheduler policy.

For information about how to create a 7210 and 1830 Port Scheduler policy, see 50.60 “To configure a 7210, 7250 and 1830 Port Scheduler policy” (p. 1601) .

5

Select an access egress policy.

For information about how to create a 7210 and 1830 Port Access Egress policy, see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556) .

6

Configure table-based ingress classification for the port. Perform the following:

1. Configure the parameters in the Ingress Classification Policy panel.
2. Select a 7210/7250 DSCP classification policy in the DSCP Classification Policy panel.

   The selected policy is effective only when the Enable Table Classification parameter is enabled.

For information about how to create a 7210/7250 DSCP classification policy, see 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638).

7

Select a network policy.

For information about how to create a 7210 and 1830 Network policy, see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .

When the Enable DEI parameter is selected on a hybrid port, the network policy assigned to the port must be configured to use DEI as the profile. The Default FC Profile parameter and the Profile parameter on the Ingress Dot1p tab in the policy must be set to the DEI option. See 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .

8

Select a network queue policy.

For information about how to create a 7210 and 1830 Network Queue policy, see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583) .

9

Save your changes and close the forms.

END OF STEPS

16.46 To create a 7210 SAS SHG

16.46.1 Before you begin

Traffic that arrives on a access or access uplink port or LAG within an SHG is not copied to other ports and LAGs in the same SHG; the traffic is copied to ports and LAGs in other SHGs if they exist within the same VPLS. Ports and LAGs are added to an SHG when you configure a port or LAG on a 7210 SAS. SHGs are not supported for a VPLS instance.

Note: SHG creation is only supported on specific 7210 SAS variants; see the NE documentation for more information.

16.46.2 Steps

1

On the equipment tree, right-click on the device where you want to configure the SHG and select Properties. The Network Element (Edit) form opens.

2

Click on the Split Horizon Groups tab.

3

Click Create. The Split Horizon Group (Create) form opens.

4

Configure the required parameters and click OK.

Note: Only one SHG can be configured per service.
A SHG and a Mesh SDP Binding cannot exist simultaneously on a single service.

5

Repeat Step 3 to Step 4 to add another SHG, if required.

6

Save your changes and close the form.

END OF STEPS

16.47 To configure a virtual Ethernet port on a 7705 SAR 2-port ring MDA

16.47.1 Purpose

Perform this procedure to configure a virtual Ethernet port on a 7705 SAR 2 × 10-Gig Bridged Ethernet XFP + 1 × 2.5G Virtual Ethernet card on a 7705 SAR-8 and 7705 SAR-18. See 50.41 “To configure a QoS network policy” (p. 1568) for information about configuring a network policy of ring type for a 7705 SAR. See 15.78 “To configure an MDA” (p. 536) for information about configuring a ring port on a 7705 SAR 2 × 10-Gig Bridged Ethernet XFP + 1 × 2.5G Virtual Ethernet card.

16.47.2 Steps

1

On the equipment tree, choose Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Virtual Port *n/n/n*, where the virtual port is the port object you need to configure.

2

Right-click on the Virtual Port icon and choose Properties. The Virtual Port (Edit) form opens.

3

Configure the required parameters.

4

Click on the States tab and configure the Administrative State parameter.

5

Click on the Policies tab.

6

Select a network queue policy.

7

Click on the Ethernet tab and configure the required parameters.

You must configure the Down When Looped parameter value as Enabled to see the Down When Looped Configuration panel parameters.

8

Click on the Network Interfaces tab to create a network interface on the virtual Ethernet port.

9

Click Add. The Create Network Interface - Routing Instance form opens. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information on creating a network interface.

Note: You can only bind a virtual Ethernet port to a network interface.

10

Save your changes and close the form.

END OF STEPS

16.48 To configure Ethernet Bandwidth Notification on a 7705 SAR Ethernet port

16.48.1 Before you begin

The 7705 SAR can interop with ETH-BN communication from the Wavence SM over Ethernet links (without MWA). ETH-BN is supported on the following ports:

- 8p GigE v1/v2/v3
- XMDA v1/v2
- 8p PMC
- 6p 10GigE/GigE
- Ethernet modules in the 7705 SAR-H and 7705 SAR-M
- Virtual port of Ring MDA and 7705 SAR-M module
- all Ethernet ports for the 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, 7705 SAR-W and 7705 SAR-Wx

ETH-BN enables the detection and extraction of ETH-BN messages to the CSM. If the ETH-BN indicates a new throughput value, the CSM programs the new value into the egress-rate of the port.

The 7705 SAR only supports the client side of ETH-BN. When enabled on a port, the egress rate that previously was a fixed bandwidth, may be dynamically altered based on the available bandwidth indicated by the ETH-BN server.

16.48.2 Steps

1

On the equipment tree, navigate to the port object on the required 7705 SAR NE. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required port and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Ethernet tab and enable the Bandwidth Rate Change Notification parameter.

Note: ETH-BN notification can only be enabled if the Egress Rate (kbps) parameter is configured. The maximum egress rate is the native port rate, e.g. 1 Gbit/s, and the minimum egress rate for the port can be as low as 1 kbit/s. Any rate request outside of this range, including 0, cannot be followed. The egress rate setting has a +/- 1% accuracy, as rate conversion to the hardware is necessary. Any request to change less than 1% is ignored.

Note: The bandwidth indicated in the Bandwidth Notification Message from the Wavence SM includes the FCS. You must enable the Egress Rate Include FCS parameter in the Egress Rate panel or the bandwidth will not match the intended rate.

4

To reduce the number of bandwidth changes, configure the Bandwidth Rate Hold Time (seconds) parameter.

Note: Any Bandwidth Notification Message received before the hold-timer expires, after the last bandwidth change, is ignored (not stored).

5

Save your changes and close the forms.

Operational information on the number of BN changes, messages statistics, previous bandwidth value etc. is displayed in the Egress Rate panel and the Ethernet Bandwidth Notification Counters panel.

END OF STEPS

16.49 To configure PoE ports on a 7210 SAS

16.49.1 Before you begin

Some 7210 SAS chassis variants support PoE on specific ports. See the NE documentation for more information about 7210 SAS support of PoE.

You can view information about power consumption and availability on the PoE tab of the shelf properties form for supporting 7210 SAS NEs.

16.49.2 Steps

1

On the equipment tree, navigate to the port object on the required 7210 SAS NE. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required port and choose Properties. The Physical Port (Edit) form opens.

3

Click on the PoE tab and configure the PoE Mode parameter.

4

View the information about the PoE connection.

5

Save your changes and close the form.

END OF STEPS

16.50 To enable or disable hardware timestamps for ports on the 7210 SAS

16.50.1 When to use

For ports on the supported 7210 SAS, perform the following to enable or disable hardware timestamps for ingress and egress PTP packets.

16.50.2 Steps

1

On the equipment tree, navigate to the port object on the required 7210 SAS NE. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required port and choose Properties. The Physical Port (Edit) form opens.

3

On the General tab, in the PTP panel, configure the Enable PTP Hardware Timestamp parameter.

4

Save your changes and close the form.

END OF STEPS

16.51 To configure MAC or VLAN authentication

16.51.1 Before you begin

Consider the following before you configure MAC or VLAN authentication:

- You cannot configure MAC authentication and VLAN authentication simultaneously on the same port.
- VLAN authentication is supported only on dot1q ports.
- Before you can configure MAC authentication, the system resource profile for the NE must be appropriately configured. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) and 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).

16.51.2 Steps

1

On the equipment tree, navigate to the port object where you want to configure MAC or VLAN authentication. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on the required port and choose Properties. The Physical Port (Edit) form opens.

3

Click on the 802.1x Port Authenticator tab.

4

Set the Controlled Port Control parameter to Auto.

5

Perform one of the following:

a. To enable MAC authentication, set the MAC Auth Enabled parameter to True.

   Configure MAC Auth Wait parameter.

b. To enable VLAN authentication, set the VLAN Auth Enabled parameter to True.

6

Save your changes and close the forms.

END OF STEPS

16.52 To configure PoE ports on a 7705 SAR

16.52.1 Before you begin

You can configure PoE on a 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-W, and 7705 SAR-Wx.

Note: You can view information about power consumption and availability on the PoE tab of the shelf properties form.

16.52.2 Steps

1

On the equipment tree, navigate to the port object where you want to configure PoE. The path is Network→NE→Shelf→Card Slot→Daughter Card Slot→Port *n/n/n*.

2

Right-click on the port on which you want to configure PoE and choose Properties. The Physical Port (Edit) form opens.

3

Click on the PoE tab and configure the PoE Mode parameter.

4

Save your changes and close the form.

5

If you are configuring a 7705 SAR-H, perform the following steps to configure the PoE power supply source.

1. On the equipment tree, right-click on the 7705 SAR-H device on which you want to configure PoE and choose Properties. The Network Element (Edit) form opens.
2. On the Network Element (Edit) form navigation tree, expand the Shelf object and the Card Slot object.
3. If you are configuring a 7705 SAR-H, Release 6.1 or later, right-click on the Daughter Card Slot object and choose Properties. The Daughter Card Slot (Edit) form opens.

   If you are configuring an earlier release 7705 SAR-H, right-click on the Shelf object and choose Properties. The Shelf (Edit) form opens.
4. Configure the PoE Power Supply Source parameter in the PoE panel.

5. Save your changes and close the form.

End of steps

16.53 To configure GPS on a 7705 SAR

16.53.1 Before you begin

You can configure GPS on the 7705 SAR-8 CSMv2 with shelf v1 and v2, 7705 SAR-18, 7705 SAR-H, and on the following 7705 SAR-Wx variants:

- 4 GigE 4-port xDSL GPS Rx
- 5 GigE GPS Rx
- 5 GigE PoE+ GPS Rx

Only the 7705 SAR-8 CSMv2 with shelf v1 and v2, 7705 SAR-18, and 7705 SAR-Ax support a GPS MDA with a port type that is configurable for GPS, GLONASS, or GPS plus GLONASS.

GLONASS can be enabled on GPS ports of the p1 × GPS Rx module on the 7705 SAR-H and on the i1 × GPS Rx module on the 7705 SAR-Wx (5 GigE PoE+ GPS Rx).

From 7705 SAR Release 8.0 R1 or later, the GNSS MDA can be configured for GPS only mode or both GPS and GLONASS mode, simultaneously. The GLONASS only mode is not supported.

16.53.2 Steps

1

Ensure that the 7705 SAR NE is configured with the appropriate GPS MDA. See 15.78 “To configure an MDA” (p. 536) for more information about configuring an MDA.

2

On the equipment tree, navigate to the port object where you want to configure GPS. The path is Network→NE→Shelf→Card Slot→Daughter Card Slot→Port *n/n/n*.

3

Right-click on the port on which you want to configure GPS and choose Properties. The Physical Port (Edit) form opens.

4

Click on the GPS tab and configure the required parameters.

For the 7705 SAR-8 CSMv2 with shelf v1 and v2, 7705 SAR-18, and 7705 SAR-Ax, Release 8.0 R1 or later, you can configure the Type parameter as GPS only or both GPS and GLONASS, simultaneously. The GLONASS only mode is not supported.

5

Save your changes and close the form.

END OF STEPS

16.54 To configure a 7705 SAR ASAP channelized TDM port

16.54.1 Before you begin

This procedure applies to 7705 SAR channelized TDM DS1/E1 ports and channelized TDM DS3/E3 ports.

16.54.2 Steps

1

On the equipment tree, navigate to the port object where you want to configure a 7705 SAR ASAP channelized TDM port. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

2

Right-click on a port object that contains a channelized port (DS1/E1 or DS3/E3) and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters.

4

Click on the DS1/E1 or DS3/E3 tab and configure the required parameters.

All ports of the ASAP daughter card are automatically configured to the same port type when the first port is configured. A mix of DS1 and E1 ports or DS3 and E3 ports is not supported. The Port Type parameter can only be modified when no channels are configured on any port of the daughter card. If all of the channels are deleted on the daughter card, the port type is automatically set to DS1 or DS3 for all ports.

5

Click on the States tab and configure the Administrative State parameter.

6

Save your changes and close the form.

END OF STEPS

16.55 To configure a channelized TDM DS1 or E1 port

16.55.1 Steps

1

On the equipment tree, right-click on the device where you need to configure a channelized TDM DS1 or E1 port and select Properties. The Network Element (Edit) form opens.

2

On the Network Element (Edit) form navigation tree, expand to the port level and click on a port object that contains a channelized DS1 or E1. The Physical Port (Edit) form opens.

3

Configure the required parameters.

4

Click on the DS1/E1 tab and configure the required parameters.

All ports are automatically configured to the same port type when the first port is configured. A mix of DS1 and E1 ports is not supported.

5

Click on the States tab and configure the Administrative State parameter.

6

Save your changes and close the form.

END OF STEPS

16.56 To configure OmniSwitch Ethernet ports

Note: You can configure MVRP fixed ports, 802.1 Q ports, aggregate ports, and VLAN Stacking Network ports.

16.56.1 Steps

1

On the equipment tree, navigate to the port object where you want to configure an OmniSwitch Ethernet port. The path is Network→*NE*→Shelf→Card Slot *n*→Port *n/n/n*.

2

Right-click on a port object and choose Properties. The Physical Port (Edit) form opens.

3

Configure the required parameters.

An access port is used for customer-facing VLAN service traffic.

A network port or channel is used for network access in the service provider transport or infrastructure network. Typically, only one OmniSwitch network port is required for tree topology; ring topology requires two ports.

The Automatic VLAN Binding parameter is only supported on Stacked VLANs.

The Automatic VLAN Binding parameter is only supported on Stacked VLANs.

4

Click on the States tab and configure the Administrative State parameter.

5

Click on the Policies tab.

Note: A port must have a service access point associated with it before you can apply a UNI policy to the port.

6

Select a UNI policy.

7

Click on the Ethernet tab and configure the required parameters.

The Inter-Frame Gap (bytes) parameter is only configurable on Gigabit Ethernet ports.

The Default VLAN Restore, Default VLAN Enable, Ignore BPDU, Authenticate, and Ingress Filtering parameters are only configurable when the Enable Port Mobility parameter is set to Enable.

8

Click on the LLDP tab and configure the required parameters.

9

Click on the Remote Peers sub-tab under the LLDP tab to search for and display LLDP remote peers that are associated with the port. The remote peers are used to determine the physical topology of the network.

10

Click on the MVRP tab and configure the required parameters.

11

Configure VLAN restrictions, if required.

1. Click on the VLAN Restrictions tab.

2. Click Create. The MVRP VLAN Registration form appears.
3. Configure the required parameters.
4. Save your changes and close the form.

12

Click on the QoS tab, and configure the required parameters.

The Servicing Mode parameter does not apply to OS 6900 and OS 10K devices.

13

If you are configuring a port on an OS 6900 or OS 10K device, click on the Queue tab and configure the required parameters.

14

Click on the DHCP/DHCPv6 Snooping tab, and configure the required parameters. The tab is only available when the port is associated with a VLAN that has DHCP/DHCPv6 snooping enabled or DHCP/DHCPv6 snooping is enabled at the switch level.

15

Click on the Ethernet Hybrid tab, and configure the required parameters. The tab is only available for hybrid Ethernet ports.

16

Click on the PoE tab, and configure the required parameters. The tab is only available for PoE ports.

17

Click on the LPS Learned MAC Entries tab to view the MAC addresses learned on the port. LPS must be enabled on the port. See 28.128 “To configure bridging on an OmniSwitch” (p. 1043) for information about configuring LPS on an Ethernet port.

18

Click Create to add static MAC address entries, if required. The MAC Entries (Create) form opens.

Note: A port must be LPS-enabled and belong to a VLAN before you can add static MAC addresses to the port.

19

Select a VLAN Site.

20

Configure the MAC Address parameter.

21

Save your changes and close the form.

22

Repeat Step 18 to Step 21 to add another MAC address, if required.

23

Save your changes and close the forms.

END OF STEPS

16.57 To configure OmniSwitch PoE Ports

16.57.1 Steps

1

On the equipment tree, navigate to the card object where you want to configure an OmniSwitch PoE port. The path is Network→*NE*→Shelf→Card Slot *n*. The Card Slot (Edit) form opens.

2

Click on the PoE tab and configure the required parameters.

3

Save your changes and close the forms.

END OF STEPS

16.58 To create and configure Xconnect anchor ports

16.58.1 Purpose

Perform this procedure to create and configure anchor ports on FP4 daughter cards on 7750 SR and 7950 XRS NE variants, Release 21.5 and later. These anchor ports can be associated to a PXC; see 16.59 “To configure PXC loopback ports” (p. 652).

16.58.2 Steps

1

On the equipment tree, navigate to the Xconnect object under the FP4 daughter card where you need to configure an anchor port.

Right-click on Xconnect and choose Create Xconnect Mac. The Create MDA Xconnect MAC wizard opens.

2

Configure the MDA Xconnect MAC Properties.

Configure the required parameters and click Next. The Xconnect Mac object appears under the Xconnect object in the equipment tree.

3

Configure the MDA Xconnect Loopback.

1. Click Create.
2. Configure the required parameters and click Finish.

The Anchor Port object appears under the Xconnect Mac object in the equipment tree (Port 1/1/m1/1).

4

Close the wizard forms.

5

Create a second anchor port, if required.

1. Right-click on Xconnect Mac and choose Create Xconnect Loopback.
2. Configure the required parameters and click Finish.

The Anchor Port object appears under the first created anchor port in the equipment tree (Port 1/1/m1/2).

6

Close the wizard forms.

7

To view details for Xconnect Mac, right-click the Xconnect Mac and choose Properties.

8

Configure the anchor ports.

1. Right-click on Anchor Port and choose Properties.
2. Configure the required parameters.
3. Save and close the forms.

END OF STEPS

16.59 To configure PXC loopback ports

16.59.1 Steps

1

On the equipment tree, navigate to the node where you need to configure a PXC port. The path is Network→*NE*→Logical Groups. Right-click on PXCs and choose Create PXC. The Port Cross Connect configuration form opens.

2

Select a physical port to host the port cross-connect in the PXC Port panel.

Note: The port must be in shutdown state.

Note: A maximum of eight PXCs can be associated to each anchor port.

3

Save your changes and close the form.

END OF STEPS

Procedures for channel and framing link configuration

16.60 To configure SONET clear channels

16.60.1 Purpose

Perform this procedure to configure SONET clear channels on OC3 to OC192 ports that provide clear channel services. Each port supports one clear channel.

16.60.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device on which you want to create a clear channel port and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the port level and click on a clear channel port object. The Physical Port (Edit) form opens. SONET is the default framing scheme.

Note: A clear channel port is a port on a card that is not designated Deep Channelized.

5

Configure the required parameters and click Apply.

6

In the navigation tree, right-click on the port and choose Create Channel. The Channel Type Selection form opens.

7

Choose the channel type. The Sts*n* SONET Channel (Create) form opens with the General tab displayed.

If you are configuring a 4 x Channelized OC3/OC12 ASAP SFP adapter card on a 7705 SAR, Release 7.0 R5 or later:

- If the MDA mode is configured as OC3, the NFM-P supports channel type Sts3 under SONET, and channel type Sts3 with payload type Sts3 (non-channelized) under SDH for clear channel POS and Network APS.
- If the MDA mode is configured as OC12, the NFM-P supports channel types Sts3 and Sts12 under SONET or SDH for clear channel POS and network APS.

You can create an SC APS or MC APS group to enable automatic protection switching. See 38.14 “To create an SC APS group” (p. 1302) and 38.15 “To create an MC APS group” (p. 1305) .

8

Configure the required parameters.

The Encapsulation Type parameter is configurable when the Mode parameter is set to Access.

If you configure the Payload Type parameter as Sts3 (non-channelized), you must configure the Mode parameter as Network and the Encap Type as PPP Auto.

9

Click Apply.

a. If the Encapsulation Type parameter of the port is ATM, the form is refreshed.

b. If the Encapsulation Type parameter of the port is not ATM, go to Step 12 .

10

If required, configure an ATM interface. Click Edit ATM. The ATM interface form opens.

11

Configure the required parameters.

12

If required, create an ILMI link, see 16.74 “To create an ILMI link” (p. 677) .

13

Click OK and close the forms.

END OF STEPS

16.61 To perform a bulk channel creation on ports that support multiple sub-channels

16.61.1 Purpose

Perform this procedure to create the maximum number of channels on ports that support multiple sub-channels. The following table lists the port types that support bulk channel creation.

Table 16-6 Bulk channel creation support

Node	OC3	OC12	DS3 ASAP	E3 ASAP	APS common port	APS SONET VTG
7450 ESS	X	X	X	X	X	X

Table 16-6 Bulk channel creation support (continued)

Node	OC3	OC12	DS3 ASAP	E3 ASAP	APS common port	APS SONET VTG
7705 SAR-8 7705 SAR-8v2 7705 SAR-18	X	X	—	—	—	—

Note: This procedure assumes you have preconfigured an NE with an appropriate daughter card that supports multiple sub-channels.

16.61.2 Steps

1

Choose Equipment from the view selector in the navigation tree.

2

Locate and expand the device on which you need to perform a bulk channel creation.

3

Perform one of the following:

a. For OC3, OC12, DS3 ASAP, and E3 ASAP ports, right-click on a port icon and choose Create Maximum # of Channels. The Create Maximum # of Channels form opens.

b. For APS common ports and APS SONET VTGs, right click on an APS Common Config SONET channel icon and choose Create Maximum # of Channels. The Create Maximum # of Channels form opens.

Note: APS common ports and APS SONET VTGs are created when you create APS groups. See 38.14 “To create an SC APS group” (p. 1302) for more information.

4

Configure the required parameters.

5

Click Execute.

END OF STEPS

16.62 To configure SONET sub-channels

16.62.1 Before you begin

SONET sub-channels are available on deep channelized OC12 and OC3 ports. SONET STS1 channels support the following payload types:

- DS3

- VT15
- VT2

16.62.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device on which you want to configure SONET sub-channels and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the port level, right-click on a port object, and choose Create Channel. The Channel Type Selection form opens.

5

Choose the channel type. The Sts*n* SONET Channel (Create) form opens with the General tab displayed.

If you are configuring a 4 x Channelized OC3/OC12 ASAP SFP adapter card on a 7705 SAR, Release 7.0 R5 or later, when the MDA Mode is OC12 with SONET framing, you can create an Sts1 channel with a DS1 channel.

6

Configure the required parameters.

For an OC3 port, use a value from 1 to 3. For an OC12 port, use a value in the format <1 to 4>.<1 to 3>.

The DDM Event Suppression parameter is configurable only on ports on SFPs and XFPs optical modular transceivers.

7

Click Apply.

8

Create the STS1 (AU3) sub-channels according to the payload type configured:

a. To create VT15 (Sdh TU11) sub-channels, go to 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

b. To create VT2 (Sdh TU12) sub-channels, go to 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

c. To create a DS3 or E3 (SDH framing only) clear channel, go to Step 9 .

9

To create a DS3 or E3 clear channel, perform the following steps:

1. From the navigation tree, select an STS1 or AU3 channel and choose Create Channel. The DS3/E3 Channel (Create) form opens.
   - For AU3 channels, DS3 and E3 channelization is available.
   - For STS1 channels, only DS3 channelization is available.
2. Configure the Description parameter.
3. Configure the Channelized parameter.
   - Choose Ds1 to channelize the DS3 to carry up to 28 DS1 channelized to the DS0 level.
   - Choose E1 to channelize the DS3 to carry up to 21 E1 channelized to the DS0 level.
   - Choose None to create a DS3 or E3 clear channel. Configure the required parameters.
4. Click OK.

   If you set the Channelized parameter to None, a DS3 clear channel in Access mode is created under the STS1 or the AU3 in the navigation tree.

   If you set the Channelized parameter to DS1, a DS3 channelized for DS1 is created under the STS1 (AU3). Go to Step 10 in 16.68 “To create TDM DS3 channels” (p. 666) to create the DS1 channels.

   If you set the Channelized parameter to E1, a DS3 channelized for E1 is created under the STS1 (AU3) in the navigation tree. Go to Step 10 in 16.68 “To create TDM DS3 channels” (p. 666) to create the E1 channels.

END OF STEPS

16.63 To configure SDH sub-channels

16.63.1 Purpose

SDH sub-channels are available on deep channelized OC12 and OC3 ports. SDH AU4 channels support the following payload types:

- TU3
- TU11
- TU12

SDH AU3 channels support the following payload types:

- DS3
- E3
- TU11
- TU12

Perform this procedure to configure SDH sub-channels on ports that provide channelized services. You must first configure port framing to SDH since default framing on SONET/SDH ports is SONET.

16.63.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device on which you want to configure SDH sub-channels and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the port level and click on a channelized port. The Physical Port (Edit) form opens.

5

Click on the SONET tab and set the Framing parameter to SDH.

6

Click OK.

7

Right-click on the port again and choose Create Channel. The Channel Type Selection form opens.

8

Choose one of the following channel types and click OK. The Sts*n* SONET channel (Create) form opens.

a. SONET Sts1 channelization and SDH Au3 channelization are equivalent. Go to Step 6 of 16.62 “To configure SONET sub-channels” (p. 655) .

b. If you choose SONET Sts3 (Sdh Stm1), the Sts3 SONET Channel (Create) form opens. Go to Step 9 .

The Payload Type parameter for SDH STM1 (AU4) is set to TUG3. There is no SONET framing equivalent. Go to Step 9 .

If you are configuring a 4 x Channelized OC3/OC12 ASAP SFP adapter card on a 7705 SAR, Release 7.0 R5 or later, when the MDA Mode is OC12 with SDH framing, you can create an Sts1 and Sts3 channel type with a DS1 or an E1 channel.

9

Configure the required parameters.

10

Click on the States tab and configure the Administrative State.

11

Click OK. The STM1 (AU4) channel appears in the navigation tree under the port of the daughter card. It contains three TUG3 groups.

12

Right-click on a TUG3 group and choose Properties. The Sdh Tug3 (Edit) form opens.

13

Choose one of the following options to configure the Payload Type parameter:

a. PDH Tu3. Click Cancel and go to Step 14 to create a TU3 channel.

b. SONET VT15 (SDH Tu11). Click Apply and go to 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

c. SONET VT2 (SDH Tu12). Click Apply and go to 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

14

To create a TU channel, right-click on the TUG3 group and choose Create Channel. The TU3 Channel (Create) form opens.

15

Configure the required parameters.

16

Click on the States tab and configure the Administrative State.

17

Click OK. A TU3 channel is created under the TUG3 group in the navigation tree.

18

Right-click on the Tu3 and choose Create Channel. The DS3/E3 Channel (Create) form opens.

The E3 is the only payload type available.

19

Configure the required parameters.

20

Click on the States tab and configure the Administrative State.

21

Click OK. An E3 clear channel is created under the AU4/TUG3 in the navigation tree.

END OF STEPS

16.64 To create VT15 (TU11) or VT2 (TU12) sub-channels

16.64.1 Before you begin

A SONET VTG and an SDH TUG2 are equivalent groups. A VTG contains four VT15 channels or three VT2. A TUG2 contains four TU11 channels or three TU12. An SDH TUG 3 has no SONET equivalent, however, it contains seven TUG2. (An STM1/AU4 contains three TUG3). See 16.15 “Comparison of SONET and SDH hierarchies” (p. 589) in “SONET and SDH sub-channel applications and structure” (p. 586) .

An STS1 channelized to carry a VT15 or VT2 payload type contains seven VTG. An AU3 or a TUG3 channelized to carry a TU11 or TU12 payload type contain seven TUG2.

Each VTG (TUG2) supports up to four VT15 (TU11) channels or three VT2 (TU12) channels.

16.64.2 Steps

1

Complete 16.62 “To configure SONET sub-channels” (p. 655) or 16.63 “To configure SDH sub-channels” (p. 657) to carry a VT15 (TU11) or VT2 (TU12) payload type.

2

To create a VT15 (TU11) channel, perform the following steps:

1. Choose Equipment from the navigation tree view selector.
2. Right-click on the device where you want to create a VT15 (TU11) channel and select Properties. The Network Element (Edit) form opens.
3. On the Network Element (Edit) form navigation tree, expand the Shelf icon.
4. Expand to the channel level, right-click on a VTG or TUG2 channel group that is to carry a VT15 (TU11) payload and choose Create Channel. The SONET VT (SDH Tu) Channel (Create) form opens.
5. Configure the required parameters.

   Set the Local Channel ID to a value in the range 1 to 4 since there are four VT15 (TU11) in a VTG (TU2).

   The default payload type for a VT15 (TU11) channel is PDH DS1. There is no other available payload type.
6. Click OK. A VT15 channel is created under the VTG or a TU11 channel is created under the TUG2 in the navigation tree.
7. Right-click on a VT15 or TU11 channel and choose Create Channel. The DS1/E1 Channel (Create) form opens.

The Local Channel ID is equal to the VT15 channel Local Channel ID parameter and the channel type is set to DS1 since a VT15 or TU11 channel supports one DS1.

8. Click OK A DS1 channel is created under the VT15 or the TU11 in the navigation tree.

   You can create up to 28 VT15 (TU11) channels on a VTG (TUG2).

9. Repeat Step 1 and Step 2 for each VT15 (TU11) channel that you want to create.
10. Since a DS1 channel is not used as a SAP, create a DS0 group for each DS1 channel created. Go to Step 11 of 16.68 “To create TDM DS3 channels” (p. 666) .

3

To create a VT2/TU12 channel, complete these steps:

1. From the navigation tree, select a VTG or TUG2 channel group to carry a VT2 (TU12) payload and choose Create Channel. The SONET VT (SDH Tu) Channel (Create) form opens.
2. Configure the required parameters.

   Set the Local Channel ID to a value in the range 1 to 3 since there are three VT2 (TU12) in a VTG (TUG2).

   PDH E1 is not available for SONET framing.

3. Click OK. A VT2 channel is created under the VTG, or a TU12 channel is created under the TUG2 in the navigation tree.
4. Right-click on the VT2 or TU12 channel and choose Create Channel. The DS1/E1 Channel (Create) form opens.

   The Local Channel ID is equal to the VT2 channel Local Channel ID parameter, since a VT2 channel supports one E1 or DS1.

5. Click on the States tab and configure the Administrative State parameter.
6. Click OK. A DS1 channel is created under the VT2 in the navigation tree. An E1 or DS1 is created under the TU12. You can create up to 21 VT2 (TU12) channels on a VTG (TUG2).
7. Repeat Step 1 and Step 2 for each VT2 (TU12) that you want to create.
8. Since a DS1/E1 channel is not used as a SAP, create a DS0 group for each DS1 or E1 channel created. Go to Step 11 of 16.68 “To create TDM DS3 channels” (p. 666) .

END OF STEPS

16.65 To create TDM DS1 or E1 channels

16.65.1 Purpose

Perform this procedure to create DS1/E1 channels and child objects on a channelized DS1/E1 port.

Note: If you are using the XML API to create a DS1 or E1 channel, the create request creates a channel with a G704 CRC framing type for E1 channels and ESF for DS1 channels. To modify the channel framing type requires a second request.

16.65.2 Steps

1

On the equipment tree, right-click on the device where you want to configure TDM DS1 or E1 channels and select Properties. The Network Element (Edit) form opens.

2

On the Network Element navigation tree, expand to the port level.

3

Click on a port. The Port form opens.

4

Configure the required general parameters.

5

Click on the DS1/E1 tab and configure the required parameters.

6

Click OK.

7

Right-click on the port and choose Create Channel. The DS1/E1 Channel (Create) form opens. If there are no channels available on that port, the Create Channel menu option is dimmed.

8

Click OK.

9

To create DS0 channel groups:

1. On the navigation tree, right-click on a DS1 or E1 channel under the required port and choose Create Channel. The DS0 Channel Group (Create) form opens.
2. Configure the required parameters.

   For the Local Channel ID parameter, for a DS0 group from a DS1 channel, choose from 1 to 24. For a DS0 group from an E1 channel, choose from 2 to 32.

   Note:

   All channel groups on a 7705 SAR ASAP port must have the same mode configured. When the first channel group is configured, all other channel groups on the port must be set to the same mode.

   The default encapsulation type for the 7705 SAR is N/A. The encapsulation type must be changed from N/A to a valid type. For instance, to support CAS signaling, the encapsulation type must be set to CEM. After a valid encapsulation type is configured for a channel group,

the encapsulation type cannot be changed for that channel group. To change the encapsulation type, the channel group must be deleted and recreated.

On a 16/32x Channelized ASAP, you must configure the MDA mode in order to use certain encapsulation types. See 15.78 “To configure an MDA” (p. 536) for more information on configuring the MDA mode.

3. Click on the Channel Group tab and configure the required parameters.

10

Save your changes and close the forms.

END OF STEPS

16.66 To configure TDM DS1 or E1 channels

16.66.1 Purpose

Perform this procedure to configure DS1/E1 channel parameters and child objects on a channelized DS1/E1 port, add a DS0 channel group, collect and view performance statistics, or view alarm information.

16.66.2 Steps

1

In the navigation tree equipment view, right-click on the device you need to configure TDM DS1 or E1 channels and choose Properties. The Network Element (Edit) form opens.

2

Expand the Shelf icon to the channel level and choose a DS1 or E1 channel. The DS1/E1 Channel (Edit) form opens.

3

Click on the States tab and configure the Administrative State parameter.

4

Click on the Channel tab. The DS1/E1 Channel (Edit) form opens.

5

Configure the required parameters.

The Synchronous Status Messages, E1 Tx DUS/DNU, and E1 SSM-Bit parameters are configurable only on E1 channels on a 7705 SAR.

The Channel Framing parameter must be set to E1-unframed to create a DS0 channel under the E1 channel.

The Clock Source parameter value Differential is only supported on the following 7705 SAR cards:

- 16/32 × Channelized DS1/E1 ASAP v2 on a 7705 SAR-8, 7705 SAR-8 v2, or 7705 SAR-18
- 4 × Channelized DS3/E3 ASAP on a 7705 SAR-8, 7705 SAR-8 v2, or 7705 SAR-18, Release 6.1 R3 or later
- 4 × Channelized OC3/OC12 ASAP SFP on a DS1/E1 channel, on a 7705 SAR-8, 7705 SAR-8 v2, or 7705 SAR-18, Release 6.2 R1 or later
- ASAP channelized card on a 7705 SAR-A

Note: If you create an MLPPP bundle and the DS1 BER exceeds a threshold, the PPP link is automatically removed from the MLPPP bundle.

6

Add a DS0 channel group, if required.

Note: For the 7250 IXR-R6, you can only create a DS0 channel under the E1 channel (not the DS1 channel) of a 32-port Any Service Channelized DS1/E1 Multi-Sync MDA.

1. Click on the SubChannels tab.
2. Click on the Create Cha... tab. The DS0 Channel Group (Create) form opens.
3. Configure the required parameters.

   Note:

   On a 16/32x Channelized ASAP, you must configure the MDA mode in order to specify the encapsulation type. See 15.78 “To configure an MDA” (p. 536) for more information about configuring the MDA mode.

   You must set the value of the Encapsulation Type parameter to CEM if you are configuring an MDA card of type 16 × Channelized DS1/E1 ASAP v2 or 32 × Channelized DS1/E1 ASAP v2 on a 7705 SAR-8 or 7705 SAR-18.
4. Click OK.
5. Click OK.

7

Perform one of the following:

a. Click Collect to collect performance statistics data on demand. The collected statistics are listed on the form.

b. Click Collect All to collect one on-demand statistics record for each statistic type that the object supports. The collected statistics are listed on the form.

8

Click on the Faults tab to view alarm information, if required.

Note: If are configuring a 7705 SAR and selected a clock source of adaptive, you can click on the Adaptive Clock History tab to view channel adaptive clock performance data for the preceding 15 min.

9

Save your changes and close the forms.

END OF STEPS

16.67 To create serial channels

16.67.1 Purpose

Perform this procedure to create serial channels and child objects from a Serial Data Interface port.

16.67.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Navigate to the device where you want to configure serial channels and expand to the port level.

3

Right-click on the port and choose Create Channel. The Serial Channel (Create) form opens. If there are no channels available on that port, the Create Channel menu option is dimmed.

4

Click on the States tab and configure the Administrative State parameter. Click Apply.

5

Click on the Channel tab and configure the required parameters.

Depending on the serial type of the configured port, different parameters are configurable in the Control Leads panel.

6

Create DS0 channel groups:

1. Click on the SubChannels tab.
2. Click on the Create Cha... button. The DS0 Channel Group (Create) form opens.
3. Configure the required parameters.

   Note:

   The default encapsulation type for the 7705 SAR is N/A. The encapsulation type must be changed from N/A to a valid type. After a valid encapsulation type is configured for a channel group, the encapsulation type cannot be changed for the channel group. To change the encapsulation type, the channel group must be deleted and re-created.

On an SDI MDA, you must configure the serial type to X.21 or V.35 to use the IPCP, FR, HDLC, or Cisco HDLC encapsulation types. See 16.42 “To configure serial ports” (p. 634) for more information.

4. Click OK.

7

Configure the DS0 channel according to the encapsulation type, if required:

a. For an IPCP channel, click Edit PPP and configure the required parameters.

b. For an FR channel, click Edit FR and configure the required parameters.

c. For a Cisco HDLC channel, click on the CISCO HDLC tab and configure the required parameters.

8

Save your changes and close the form.

END OF STEPS

16.68 To create TDM DS3 channels

16.68.1 Purpose

Perform this procedure to create DS3 channel and children objects from a channelized DS3 port.

16.68.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device on which you want to create a TDM DS3 channel and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the port level and click on a port. The Physical Port (Edit) form opens.

5

Configure the required parameters.

6

Choose the DS3/E3 tab and configure the required parameters. Click OK.

7

Right-click on the port again and choose Create Channel. The DS3/E3 Channel (Create) form opens. If there are no channels available on that port, the Create Channel menu option is dimmed.

8

To create the DS3/E3 channel, complete these steps:

a. Configure the Channelized parameter to DS1 or E1, if you want to channelize to the DS0 level.

b. Configure the Channelized parameter to None if this is a clear channel application.

Configure the required parameters.

The Local Channel ID parameter is configured automatically and the Mode parameter is always Access for TDM.

Note: On a 4x Channelized ASAP, you must configure the MDA mode in order to use certain encapsulation types. See 15.78 “To configure an MDA” (p. 536) for more information on configuring the MDA mode.
If you want to configure a DS3/E3 channel as a network interface on a channelized ASAP MDA, see 16.70 “To configure a DS3/E3 channel as a network interface on a channelized ASAP MDA” (p. 671) .

9

Click OK.

If this is a Channelized DS1 or E1 application, continue to Step 10 .

10

To create DS1 or E1 channels, complete these steps:

1. From the navigation tree, select a DS3 channel and choose Create Channel. The Create DS1/E1 Channel form opens.
2. Configure the Local Channel ID parameter.
   - For DS1 choose from 1 to 28.
   - For E1 choose from 1 to 21.
3. Click OK. The DS1 or E1 channels are created under the DS3 in the navigation tree.

11

To create DS0 channel groups, complete these steps:

1. From the navigation tree, select a DS1 or E1 channel and choose Create Channel from the contextual menu to create DS0 groups. The DS0 Channel (Create) form opens.
2. Configure the required parameters.

   For a DS0 group from a DS1 choose from 1 to 24 for the Local Channel ID parameter. For a DS0 group from an E1, choose from 2 to 32.

   Note:

On a 4x Channelized ASAP, you must configure the MDA mode in order to use certain encapsulation types. See 15.78 “To configure an MDA” (p. 536) for more information on configuring the MDA mode.

3. Click on the Channel Group tab and configure the required parameters.
4. Click OK. The timeslot is assigned to the DS0 channel group in the navigation tree.

END OF STEPS

16.69 To configure TDM DS3 channels

16.69.1 Purpose

Perform this procedure to configure DS3 channels and children objects from a channelized DS3 port.

16.69.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device where you want to configure TDM DS3 channels and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the channel level and click on a DS3 channel. The DS3/E3 Channel (Edit) form opens. Configure the required parameters.

5

Configure the Channelized parameter and click Apply.

If you choose DS1 or E1, perform Step 10 to Step 21 .

6

Configure the Encapsulation Type parameter and click Apply.

a. If you choose BCP Null, BCP DOT1 Q, or IPCP, go to Step 7 .

b. If you choose FR, go to Step 8 .

c. If you choose Cisco HDLC, go to Step 9 .

d. If you choose WAN Mirror, perform Step 11 to Step 14 .

7

Configure the PPP interface for the channel, if required.

1. Click Edit PPP. The PPP Interface form opens. Configure the required parameters.
2. Click OK.
3. Perform Step 11 to Step 14 .

8

Configure the FR interface, if required.

1. Click Edit FR. The FR Interface form opens. Configure the required parameters.
2. Click on the Frf12 tab and configure the required parameters. Set the Mode parameter to Enabled.
3. Click Select beside the MCFR Egress QoS Profile parameter. The Select MCFR Egress Qos Profile - FR Configuration form opens.
4. Choose a profile from the list and click OK. The Select MCFR Egress Qos Profile - FR Configuration form closes and the FR Interface form is refreshed with the MCFR Egress QoS Profile information.
5. Click OK.
6. Perform Step 11 to Step 14 .

9

Configure the Cisco HDLC information.

1. Click on the Cisco HDLC tab and configure the required parameters.
2. Perform Step 11 to Step 14 .

10

Configure the DS1 and DS0 channels. Perform the following steps:

1. Click on the SubChannels tab.
2. Click Create. The DS1/E1 Channel (Create) form opens. Configure the required parameters.
3. Click on the States tab and configure the required parameters.
4. Click OK.

11

Click on the States tab and configure the required parameters.

12

Click on the Channel tab. The DS3/E3 Channel (Edit) form opens. Configure the required parameters.

The Clock Source parameter value Differential is supported on a 4 × channelized DS3/E3 ASAP MDA for a DS3/E3 channel, on a 7705 SAR-8, 7705 SAR-8 v2, or 7705 SAR-18, Release 7.0 R1 or later.

13

Click on the Message Data Link tab and configure the required parameters.

The MDL Message Type parameter specifies the Line Message Data Link message for a DS3 and specifies the transmission method of a message over a channelized interface. The parameter is only applicable if the DS3 is using C-bit framing. The default is disabled. Click on the check boxes to enable transmission methods and enter text strings to choose the message options for this parameter as required. The transmission options are:

- Test Signal
- DS3 Path
- Idle Signal

Table 16-7 MDL message options

Option	String Length	Description
Port Number String	0 to 38 characters	specifies the port ID code
Generator Number String	0 to 38 characters	specifies the generator number to send in the MDL test signal message
Equipment ID Code	0 to 10 characters	specifies the Equipment ID code
Location ID Code	0 to 11 characters	specifies the Location ID code
Frame ID Code	0 to 10 characters	specifies the Frame ID code
Unit ID Code	0 to 6 characters	specifies the unit ID code
Facility ID Code	0 to 38 characters	specifies the facility ID code

14

Click OK.

15

Right-click on a DS1 channel in the navigation tree and choose Properties. The property form for the channel opens.

16

Click on the States tab and configure the required parameters.

17

Add a DS0 channel group, if required. Perform the following steps:

1. Click on the SubChannels tab.
2. Click Create. The DS0 Channel Group (Create) form opens.
3. Configure the required parameters.

   Note:

On a 4x Channelized ASAP, you must configure the MDA mode in order to use certain encapsulation types. See 15.78 “To configure an MDA” (p. 536) for more information on configuring the MDA mode.

4. Click OK.

18

Click on the Channel tab and configure the required parameters.

19

Click OK.

20

Configure timeslots for the DS0 channel group. Perform the following steps.

1. Choose a DS0 channel in the navigation tree and choose Properties. The property form for the channel opens.
2. Configure the required parameters.
3. Click on the States tab and configure the required parameters.
4. Click on the Channel Group tab and configure the required parameters.

   Choose one or more timeslots. You can use the Select All and Deselect All buttons.
5. If you set the Encapsulation Type parameter to BCP Null, BCP DOT1 Q, or IPCP in 2 , you can configure the PPP for the timeslot on the Channel Groups tab.
   - Click Edit PPP. The PPP Interface form opens. Configure the required parameters.

     The Compression parameter is configurable when the DS1 or E1 channel group is on an ASAP MDA in a 7750 SR.
   - Click OK.

21

Click OK.

END OF STEPS

16.70 To configure a DS3/E3 channel as a network interface on a channelized ASAP MDA

16.70.1 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device where you want to configure a DS3/E3 channel and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon to the port level.

4

Perform one of the following.

a. If you are configuring an OC-*n* channelized ASAP MDA, right-click on an available port and choose Create Channel. The Sts SONET Channel (Create) form opens; go to Step 5 .

b. If you are configuring a DS3/E3 channelized ASAP MDA, right-click on an available port and choose Create Channel. The DS3/E3 Channel (Create) form opens; go to Step 9 .

5

Configure the required parameters.

6

Click on the States tab and configure the required parameters.

7

Click OK.The Sts SONET Channel (Create) form closes and the STS1 channel appears in the navigation tree under the port of the channelized ASAP daughter card.

8

Right-click on the STS1 channel and choose Create Channel. The DS3/E3 Channel (Create) form opens.

9

Configure the required parameters.

10

Click on the States tab and configure the required parameters.

11

Click Apply. The form displays additional tabs.

12

Click on the Channel tab and configure the required parameters.

13

Click on the Message Data Link tab and configure the required parameters.

The MDL Message Type parameter specifies the Line Message Data Link message for a DS3 and specifies the transmission method of a message over a channelized interface. The parameter is only applicable if the DS3 is using C-bit framing. The default is disabled. Click on

the check boxes to enable transmission methods and enter text strings to choose the message options for this parameter as required. The transmission options are:

- Test Signal
- DS3 Path
- Idle Signal

Table 16-8 MDL message options

Option	String Length	Description
Port Number String	0 to 38 characters	specifies the port ID code
Generator Number String	0 to 38 characters	specifies the generator number to send in the MDL test signal message
Equipment ID Code	0 to 10 characters	specifies the Equipment ID code
Location ID Code	0 to 11 characters	specifies the Location ID code
Frame ID Code	0 to 10 characters	specifies the Frame ID code
Unit ID Code	0 to 6 characters	specifies the unit ID code
Facility ID Code	0 to 38 characters	specifies the facility ID code

14

Click OK. The DS3/E3 channel appears in the navigation tree under the port on the channelized ASAP MDA.

END OF STEPS

16.71 To configure data framing on a 7705 SAR

16.71.1 Purpose

Perform this procedure to configure data framing on a 7705 SAR-8 or 7705 SAR-18. Data framing is supported on TPIF channels on an 8 × C37.94 Teleprotection Interface MDA or TPIF ports on an 8 × Voice Teleprotection Interface MDA. See 15.78 “To configure an MDA” (p. 536) for information about how to configure daughter cards.

Note: Data framing can be configured as Unframed only on odd numbered ports of the 8 × C37.94 Teleprotection Interface MDA and TPIF ports on an 8 × Voice Teleprotection Interface MDA. For example, ports 1,3,5 and 7 on the C37.94 MDA and port 1 on the VT MDA can be used in Unframed mode.

Data framing can be configured as Framed on all ports of the 8 × C37.94 Teleprotection Interface MDA and TPIF ports on an 8 × Voice Teleprotection Interface MDA.

16.71.2 Steps

1

Choose Network→NE→Shelf→Card Slot n→Daughter Card Slot n→Port n/n/n→Channel where the channel object is the channel that you need to configure for data framing.

2

Right-click on the channel and choose Create. The Data Channel (Edit) form opens.

3

Click on the Channel tab.

4

Configure the Data Framing Mode parameter.

Note: In Unframed mode, all 32 timeslots will be selected for DS0 channels. The channels can be added as SAPs on a Cpipe service that has the VC Type parameter configured as SAToP Tpif.
In Framed mode, TPIF ports can be added as SAPs on a Cpipe service that has the VC Type parameter configured as CESoPSN.

5

Click OK and confirm to close the form.

6

Add the channels or TPIF ports as SAPs on a CPipe service, if required.

See 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161).

END OF STEPS

16.72 To configure an L3 interface on a DS3/E3 channel on a channelized ASAP MDA

16.72.1 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device on which you want to configure an L3 interface and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the channel level and click on a DS3/E3 channel on a channelized ASAP MDA port. The DS3/E3 Channel (Edit) form opens.

5

Click on the Policies tab to configure the policies for the DS3/E3 channel. Perform the following steps.

1. Click Select beside the Network Queue Policy Name parameter. The Select Network Queue Policy - DS3/E3 Channel list form opens.
2. Select a network queue policy from the list and click OK. The Select Network Queue Policy - DS3/E3 Channel list form closes and the DS3/E3 Channel (Edit) form reappears.
3. Click Select beside the ID parameter in the Accounting Policy panel. The Select Accounting Policy - DS3/E3 Channel list form opens.
4. Select an accounting policy from the list and click Ok. The Select Accounting Policy - DS3/E3 Channel list form closes and the DS3/E3 Channel (Edit) form refreshes with the accounting policy information.
5. Configure the Collect Accounting Statistics parameter.
6. Click Select in the Port Scheduler Policy panel to assign a port scheduler policy to the port. The Select Port Scheduler Policy - Physical Port form opens.
7. Select a port scheduler policy from the list and click OK. The Select Port Scheduler Policy - Physical Port form closes.

6

Click on the Network Interfaces tab.

7

Click Create. The Create Network Interface - Routing Instance form opens.

8

See 27.17 “To create an L3 network interface on a routing instance” (p. 856) to create a network interface on the DS3/E3 channel.

The port is set to the DS3/E3 channel by default.

9

The L3 interface appears on the Network Interfaces tab. Close the DS3/E3 Channel (Edit) form.

END OF STEPS

16.73 To configure a PVC

16.73.1 Before you begin

PVCs are automatically created by the device when a new L3 interface over ATM is created on the NFM-P, and when a new ILMI link is created between ATM interfaces. See 16.74 “To create an ILMI link” (p. 677) for more information about creating ILMI links.

16.73.2 Steps

1

Choose Equipment from the navigation tree view selector.

2

Right-click on the device where you want to configure a PVC and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4

Expand to the channel level and click on a SONET/SDH ATM clear channel. The property form for the channel opens.

5

Click on the L3 Interfaces tab. Choose the appropriate interface from the list and click Properties. The L3 Interface form opens.

6

Click on the ATM tab and configure the required parameters.

7

Click Apply.

8

To view more information about the PVC:

1. Click View PVC Connection. The ATM PVC Connection form opens.
2. Click on the tabs to view information for the ATM PVC connection.

9

Click Cancel to close the ATM PVC Connection form, the L3 Interface form, and the channel form.

END OF STEPS

16.74 To create an ILMI link

16.74.1 Purpose

Perform this procedure to create an ILMI link between device ATM interfaces. This functionality is supported on the 7750 SR. Ensure that ATM QoS policies are configured before you create the ILMI link. See Chapter 49, “Policies overview” for more information about creating ATM QoS policies.

16.74.2 Steps

1

Perform one of the following to open the ATM interface properties form.

a. Use the navigation tree

1. Choose Equipment from the navigation tree view selector.
2. Right-click on the device on which you want to create an ILMI link and select Properties. The Network Element (Edit) form opens.
3. On the Network Element (Edit) form navigation tree, expand the Shelf icon.
4. Expand to the port level and click on a clear channel port. The Physical Port (Edit) form opens.

   Note:

   A clear channel port is a port on a card that is not designated Deep Channelized.
5. Click on the Channels tab.
6. Choose the appropriate channel from the list and click Properties. The Channel (Edit) form opens.
7. Click Edit ATM. The ATM Interface form opens.

b. Use the Manage Equipment list form to search for a clear channel port with an ATM interface.

1. Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.
2. Choose Port.
3. Configure the filter criteria to search for ATM encapsulation types.
4. Click Search and select an entry by double-clicking. Alternatively, select an entry and click Properties. The Select form closes and the Channel (Edit) form opens.
5. Click Edit ATM. The ATM Configuration (Edit) form opens.

2

Click Create ILMI Link. The Configure ILMI Link form opens.

3

Configure the required parameters.

The Restore Keep-Alive Defaults parameter restores the default values of the Keep-Alive Polling Frequency (seconds), Keep-Alive Polling Count, and Keep-Alive Test Frequency (seconds) parameters.

The Keep-Alive Test Frequency (seconds) parameter is configurable when the Protocol Version parameter is set to 4.0.

4

Choose an ingress ATM QoS policy by clicking Select. The Select ATM Policy - ATM ILMI Link Configuration list form opens. Click Search and choose an ingress ATM policy from the list to associate with the ILMI link. Click OK. The Select ATM Policy - ATM ILMI Link Configuration list form closes.

5

Choose an egress ATM QoS policy by clicking Select. The Select ATM Policy - ATM ILMI Link Configuration list form opens. Click Search and choose an egress ATM policy from the list to associate with the ILMI link. Click OK. The Select ATM Policy - ATM ILMI Link Configuration list form closes.

6

Click Finish and close the forms.

END OF STEPS

16.75 To configure an ILMI link

16.75.1 Steps

1

Perform Step 1 of 16.74 “To create an ILMI link” (p. 677) to open the ATM Configuration (Edit) form.

2

Click Edit ILMI Link. The ILMI Link form opens.

3

Configure the required parameters.

The Restore Keep-Alive Defaults parameter restores the default values of the Keep-Alive Polling Frequency (seconds), Keep-Alive Polling Count, and Keep-Alive Test Frequency (seconds) parameters.

The Keep-Alive Test Frequency (seconds) parameter is configurable when the Protocol Version parameter is set to 4.0.

4

Choose an ingress ATM QoS policy by clicking Select. The Select ATM Policy - ATM ILMI Link Configuration list form opens. Click Search and choose an ingress ATM policy from the list to associate with the ILMI link. Click OK. The Select ATM Policy - ATM ILMI Link Configuration list form closes.

5

Choose an egress ATM QoS policy by clicking Select. The Select ATM Policy - ATM ILMI Link Configuration list form opens. Click Search and choose an egress ATM policy from the list to associate with the ILMI link. Click OK. The Select ATM Policy - ATM ILMI Link Configuration list form closes.

6

Click on the Peer Interface Configuration tab to view ILMI properties of the peer interface of the link.

7

Click OK to close the Configure ILMI Link form. The ATM Configuration (Edit) form opens.

8

To remove an ILMI link, click Remove ILMI Link. Removing the ILMI link also removes the PVC.

Note: An ILMI link can be removed only if its Administrative Status parameter is set to Disabled.

9

Click Cancel to close the forms.

END OF STEPS

16.76 To view the channels associated with a 1830 VWM TLU port

16.76.1 Steps

1

On the equipment tree, expand Network→1830–VWM-OSU NE→1830–VWM-TLU shelf→Card Slot 1 TLU→Port.

2

Right-click on the Port object and choose Properties. The Physical Port (Edit) form opens.

3

Click on the CDR Channel tab, choose the CDR channel interface, and click Properties. The CDR Channel (Edit) form opens.

Result: The mapping of the port to the CDR channel appears.

End of steps

16.77 To retrieve 1830 VWM DDM data

16.77.1 Steps

1

On the equipment tree, expand Network→1830 VWM→Shelf→Card.

2

Right-click on the port object and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Statistics tab, choose DDM Statistics (On Demand) (VWM) from the object drop-down menu, and click Collect. The DDM data is listed.

4

Close the form.

End of steps

16.78 To configure an OSC port of an 1830 VWM OSU as a RFLM port

16.78.1 Steps

1

On the equipment tree, expand Network→1830 VWM OSU→Shelf→Card.

2

Right-click on the OSC port and choose Properties. The Physical Port (Edit) form opens.

3

Click on the States tab and set the Administrative State parameter to Down.

4

Click on the Port Specifics tab and set the Interface Role parameter to RFLM.

5

Enter suitable text in the RFLM Label field.

6

Save your changes and close the form.

END OF STEPS

16.79 To perform CPRI monitoring using 1830 VWM TLU 9M MON ports

16.79.1 Before you begin

Ensure that the States parameter of the target ports (client and its corresponding line port, for example, if the target port is C1, its corresponding line port is L1) are configured as Maintenance, if you need the monitoring mode to be configured as tap Insert or terminate.

16.79.2 Steps

1

On the equipment tree, expand Network→1830 VWM OSU→Shelf (TLU 9M)→Card (TLU 9M).

2

Right-click on the MON port and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Interface Monitor tab.

4

Configure the Monitor Mode parameter.

5

Click Select in the Interface Monitor Details panel and the Select Target Port form opens.

6

Select the target port and click OK. The port details are updated in the Interface Monitor Details panel.

Note:

If...	then...
you configure the Monitoring Mode parameter as Listen in Step 4,	the target ports can be in up, down, or maintenance state.
you configure the Monitoring Mode parameter as Tap Inset or Terminate in Step 4,	only the target ports that have the States Parameter configured as Maintenance are listed for selection in the Select Target Port form.

7

Click Apply.

8

Click on the States tab and configure the Administrative State parameter as Up.

9

Save your changes and close the form.

END OF STEPS

16.80 To configure OTDR on 1830 VWM

16.80.1 Purpose

Perform this procedure to configure Optical Time Domain Reflectometer (OTDR) on the supported 1830 VWM. OTDR functionality is used as a diagnosis to detect fiber breaks and open connectors in the network, eliminating the delay in fault detection.

16.80.2 Steps

1

On the equipment tree, expand Network→1830–VWM-OSU NE→Shelf→Card Slot→Port object and choose Properties. The Physical Port (Edit) form opens.

2

Click on the OTDR tab, and then on the OTDR Specifics tab.

3

Configure the OTDR Mode in the OTDR panel.

4

Configure the Execution Measurement parameter in the OTDR Measurement panel.

Execution details are populated under the Execution Details panel.

5

Click on the OTDR Results tab. The results are displayed in a tabular format. Click Resync to display the latest OTDR results.

END OF STEPS

17 Inventory management

Managing inventory

17.1 Overview

17.1.1 Functional description

You can use an NFM-P GUI or OSS client to inventory the managed network equipment. The inventory information is available to GUI clients on equipment list, properties, and management forms. For information about using the XML API interface to perform inventory management, see the “Inventory management” chapter of the *NSP NFM-P XML API Developer Guide*.

Inventory information may be required for the following purposes:

- spare equipment management
- SLA audits
- license capacity tracking

17.1.2 License capacity tracking

An NFM-P license has one overall capacity value for managed equipment; each licensed equipment object consumes a number of license points based on the equipment type. You can use inventory management to anticipate the need for additional license capacity as the managed network grows.

For licensed equipment types, an inventory list includes the following fields:

- Licensed Product—whether the equipment consumes NFM-P license points
- License Points Consumed—the number of NFM-P license points that the equipment consumes

Creating a license points inventory

You can use the License Points Inventory button on the NFM-P License form to generate an inventory of the license point consumption for managed objects and subscriber types. The inventory entries are saved in a file and listed by device site ID or subscriber type; the information for each includes the following:

- object FDN, for equipment
- associated site ID, for equipment
- licensed product name or subscriber type
- number of license points that the object consumes

See the section on software and license configuration procedures in the *NSP System Administrator Guide* for information about creating a license point inventory.

17.1.3 Viewing inventory information in the GUI

Using the NFM-P Equipment Manager, you can list the license information for the following object types:

- cards, using Card (Physical Equipment) as the search criterion
- card child objects:
  - Blade
  - IO Card
  - IO Card→Base Card
  - IO Card→Daughter Card
  - IO Card→MCM Card
  - Processor Card
- chassis, using Network Element (Network) as the search criterion
- VWM shelves, using VWM Shelf (Physical Equipment) as the search criterion

You can use a GUI client to retrieve the inventory information for equipment objects in the following contexts:

- one managed NE
- entire managed network

Note: Not all columns in an inventory list apply to all equipment types. For example, not all card types have an associated CLEI code. In such a case, an inventory list column displays N/A.

Figure 17-1 Generating inventory data for all managed NEs

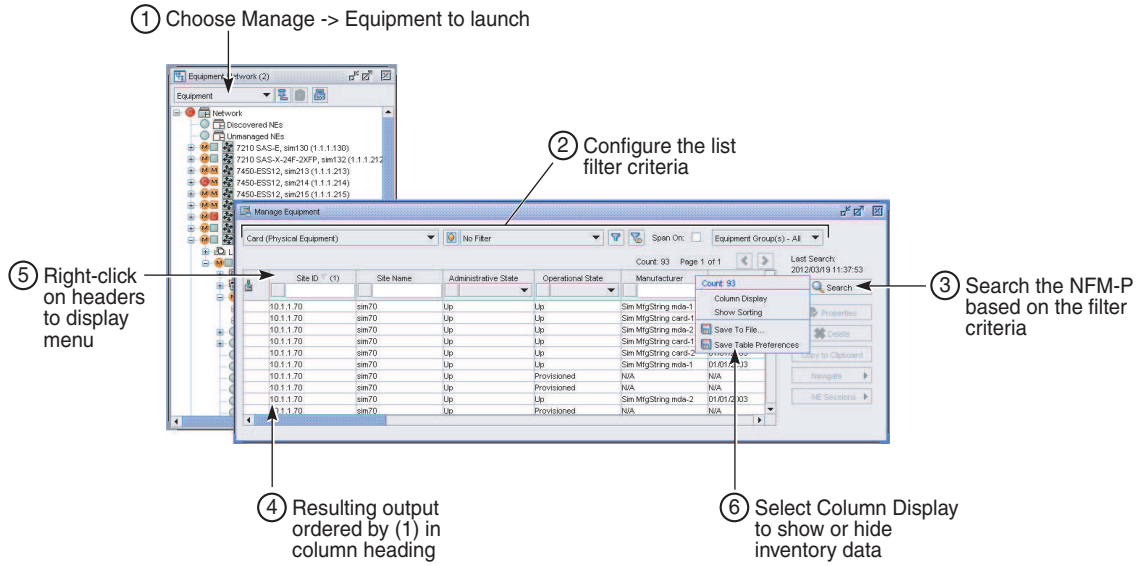

Figure 17-2 Generating inventory data for a managed NE

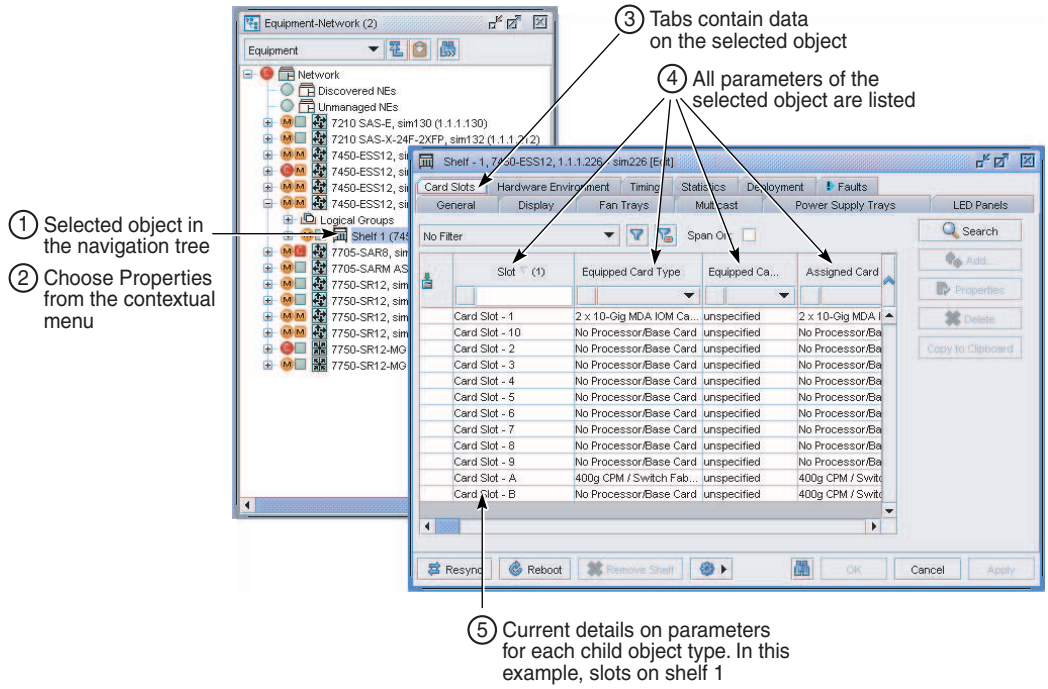

22864

You can save the inventory data in various formats, for example, CSV and HTML.

Figure 17-3 CSV inventory output

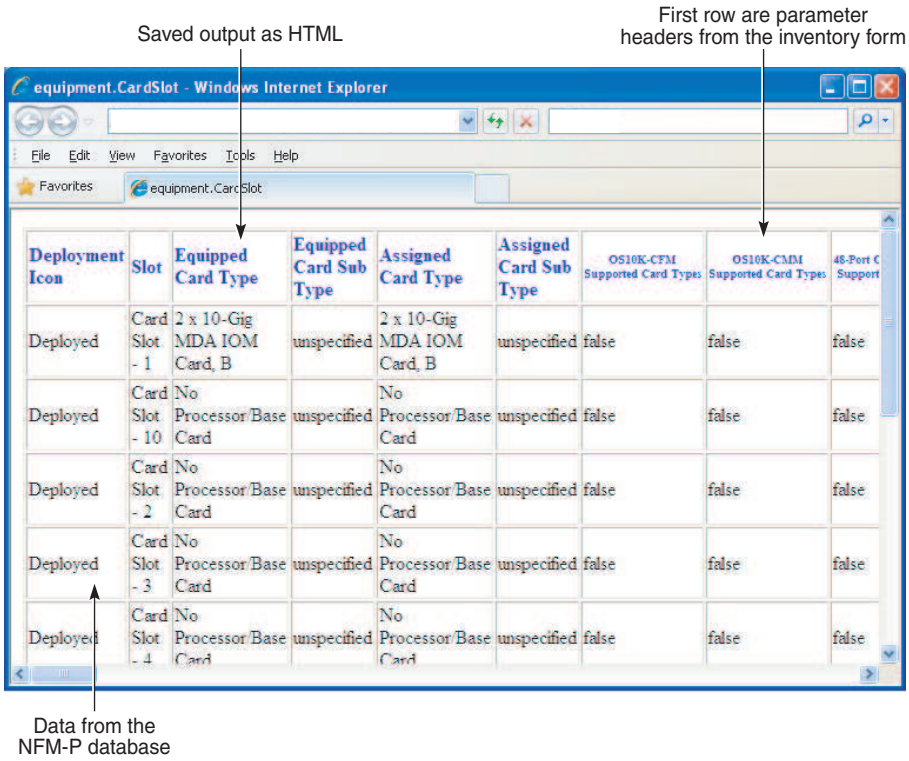

Figure 17-4 HTML inventory output

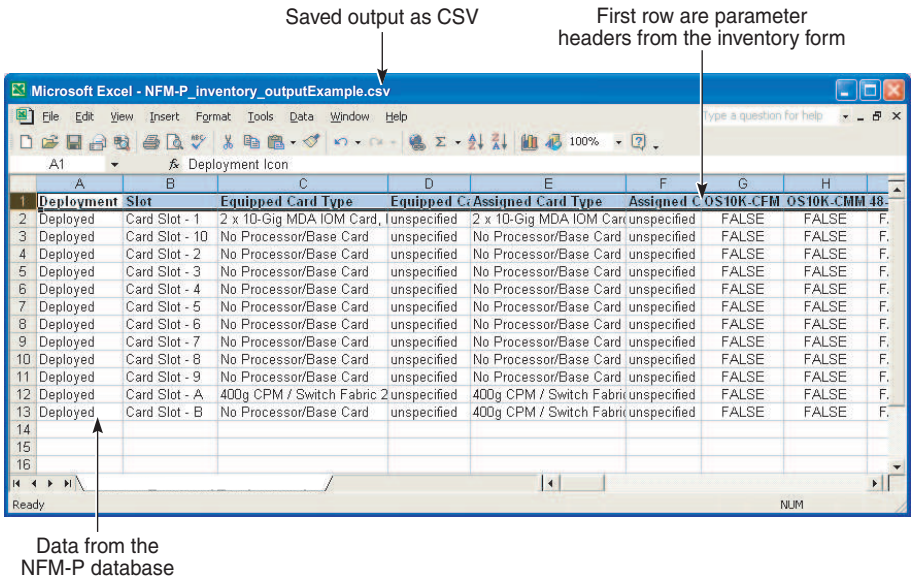

17.1.4 Exporting a network inventory file for adaptor modules and license details

Using the NFM-P Equipment Manager, you can export a zipped network inventory file that contains inventory files for NEs, shelves, cards, media adaptors, and license information for all NFM-P-managed NEs. See 17.12 “To export a network inventory file for adaptor modules and license details” (p. 697).

17.2 Sample inventory management workflow

17.2.1 Stages

1

Determine the part needs based on business requirements and network growth.

2

Based on the business requirements, use the NFM-P to generate one or more filtered inventory lists that contain only the objects for which spare parts may be required.

3

Save the lists of required spare parts in a format that is compatible with your back-office systems; for example, CSV or HTML. Move the lists to another system, for example, an order management or data management system, for processing.

4

Correlate inventory lists with data from other back-office systems.

17.3 Workflow to manage inventory

17.3.1 Stages

1

Determine the inventory management requirements.

2

Provide inventory list requirements to NOC operators.

3

As required, list and sort inventory data for a managed NE or the entire network; see 17.4 “To list and sort equipment information” (p. 689) .

4

As required, save an inventory list to a file; see 17.5 “To save an inventory list” (p. 690) .

5

As required, save search filters; see 1.33 “To save search filters” (p. 132) .

6

As required, generate an inventory of objects on one managed NE:

- CLEI codes; see 17.6 “To inventory the CLEI codes of NE objects” (p. 690)
- card software versions; see 17.7 “To inventory the card software versions of one NE” (p. 691)
- port types; see 17.8 “To inventory the port types of one NE” (p. 692)
- shelf data; see 17.9 “To inventory the shelf data for one NE” (p. 693)

7

As required, generate a network-wide inventory of all managed objects, see 17.10 “To generate a network-wide inventory of managed objects” (p. 693).

8

Generate an inventory list for NE SLA audits; see 17.11 “To collect inventory data for NE SLA audits” (p. 696) .

9

As required, export a zipped network inventory file for adaptor modules and license details; see 17.12 “To export a network inventory file for adaptor modules and license details” (p. 697).

10

Move inventory lists to other systems for further processing, as required.

17.4 To list and sort equipment information

17.4.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Perform one of the following to generate a list of equipment.

a. List the equipment on one NE.
   1. Choose Network Element (Network) from the object drop-down menu and click Search. A list of NEs is displayed.
   2. Select an NE and click Properties. The Network Element (Edit) form opens.
   3. Click on the Inventory tab and choose an object type from the object drop-down menu. A list of objects is displayed.

b. List the equipment in the entire network.
   1. Choose an object type from the object drop-down menu.
   2. Configure the filter criteria and click Search. A list of objects is displayed.

3

Perform one or more of the following to format the results.

a. To display the number of items in the list, right-click on the list heading and view the Count value.

b. To sort the list, click on a column heading. The column heading displays an arrow that indicates the sort order.

c. To move a column, drag and drop the column to a different position.

d. To remove a column:
   1. Right-click on the column heading and choose Column Display. The Column Display form opens.
   2. Select the columns to remove in the Displayed on Table list, and then click on the left arrow. The columns move to the Available for Table list.
   3. Click OK. The columns are removed from the table.

e. To sort multiple columns:
   1. Right-click on a column heading and choose Show Sorting. The Show Sorting form opens.

2. Select one or more properties in the Available for Sorting panel, and then click on the right arrow button. The properties move to the Used for Sorting panel.
3. Click Sort Ascending or Sort Descending, as required.
4. Close the Show Sorting form.

4

Save the inventory output, as required. see 17.5 “To save an inventory list” (p. 689) .

5

Close the form.

END OF STEPS

17.5 To save an inventory list

17.5.1 Steps

1

Right-click on a column heading of the inventory output and choose Save to File. The Save As file browser form opens.

2

Navigate to the directory in which you want save the file.

3

Configure the File Name and Files of Type parameters.

4

Click Save. The NFM-P saves the inventory list.

END OF STEPS

17.6 To inventory the CLEI codes of NE objects

17.6.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Perform one of the following to generate a list of network objects.

a. List the CLEI codes for objects on one NE.
   1. Choose Network Element (Network) from the object drop-down menu and click Search. A list of NEs is displayed.
   2. Select an NE and click Properties. The Network Element (Edit) form opens.
   3. Click on the Inventory tab and choose a network object. The form displays a list of network objects.

b. List the CLEI codes for all network objects of a specific type. Choose a network object.

3

Remove all columns from the list except for object ID, object type, and CLEI code. See 17.4 “To list and sort equipment information” (p. 689) for information.

4

Save the information to a file. See 17.5 “To save an inventory list” (p. 690) .

5

Close the forms.

END OF STEPS

17.7 To inventory the card software versions of one NE

17.7.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Choose Network Element (Network) from the object drop-down menu and click Search. A list of NEs is displayed.

3

Select an NE in the list and click Properties. The Network Element (Edit) form opens.

4

Click on the Inventory tab and choose Card (Physical Equipment) from the object drop-down menu. A list of cards is displayed.

5

Remove all columns from the list except for Slot ID and Software Version. See 17.4 “To list and sort equipment information” (p. 689) .

6

Save the inventory output. See 17.5 “To save an inventory list” (p. 690) .

7

Close the forms.

End of steps

17.8 To inventory the port types of one NE

17.8.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Choose Network Element (Network) from the object drop-down menu and click Search. A list of NEs is displayed.

3

Select an NE in the list and click Properties. The Network Element (Edit) form opens.

4

Click on the Inventory tab and choose Port (Physical Equipment) from the object drop-down menu. A list of ports is displayed.

5

To display more detailed port inventory information, expand the Port (Physical Equipment) object and choose from the available objects.

6

Sort the inventory output. See 17.4 “To list and sort equipment information” (p. 689) .

7

Save the inventory output. See 17.5 “To save an inventory list” (p. 690) .

8

Close the forms.

End of steps

17.9 To inventory the shelf data for one NE

Note: You can also generate an inventory of shelf data for all NEs in a managed network; see 17.10 “To generate a network-wide inventory of managed objects” (p. 693) for more information.

17.9.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Choose Network Element (Network) from the object drop-down menu and click Search. A list of NEs is displayed.

3

Select an NE in the list and click Properties. The Network Element (Edit) form opens.

4

Click on the Inventory tab and choose Shelf (Physical Equipment) from the object drop-down menu. A list of shelf objects is displayed.

5

Select a shelf in the list and click Properties. The Shelf (Edit) form opens.

6

Click on the Display tab. A picture of the selected shelf is displayed.

7

Use screen capture software to record the screen image, or otherwise copy the shelf information to a file for further processing.

8

Close the open forms, as required.

END OF STEPS

17.10 To generate a network-wide inventory of managed objects

17.10.1 Purpose

Perform this procedure to inventory all equipment of one object type, for example, all cards. The table below describes the columns to include in the inventory list for each equipment type.

Table 17-1 Inventory columns for managed objects

Equipment	Columns to include in inventory list
Cards	• Manufacture Date • Part Number • Serial Number • CLEI Code
CCMs	• Manufacture Date • Part Number • Serial Number • CLEI Code • CCM Index • CCM Equipped Type
Fan trays	• Operational State • Administrative State • Device State
Flash memory units	• Serial Number • Firmware Revision • Model Number • Capacity (sectors) • Amount Used (sectors)
Physical links	• Endpoint A Type • Endpoint A — Port • Endpoint B Type • Endpoint B — Port
Ports	• Name • CLI Name • MTU (bytes) • State
Power supply trays	• Administrative State • Operational State • AC Voltage Status • DC Voltage Status • Assigned Type

Table 17-1 Inventory columns for managed objects (continued)

Equipment	Columns to include in inventory list
Processors	• Operational State • Administrative State • CLEI Code • Manufacturer • Manufacture Date • Part Number • Serial Number
Shelves	• Operational State • Administrative State • CLEI Code • Manufacturer • Manufacture Date • CLLI Code • Part Number • Serial Number (Manufacturer Details)
Management ports	• System ID • Management IP Address • Location • Chassis Type • Sys Object ID • Software Version • Descriptor Version (software release) • Resource Group ID

17.10.2 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Choose Card (Physical Equipment) from the object drop-down menu.

3

Remove all columns from the list except the required columns shown in the table above. See 17.4 “To list and sort equipment information” (p. 689) for information.

4

Save the inventory output. See 17.5 “To save an inventory list” (p. 690) .

5

Close the form.

End of steps

17.11 To collect inventory data for NE SLA audits

Note: SLA inventory collection requirements differ. Ensure that you collect the required data for your SLA.

Inventory data collection can be resource-intensive; perform inventory collection during a maintenance period or a period of low network activity.

17.11.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

To collect the required SLA audit data, choose the required object type from the object drop-down menu.

a. For card SLA audits, choose Card (equipment).

b. For shelf SLA audits, choose Shelf (equipment).

c. For power supply tray SLA audits, choose Power Supply Tray (equipment).

3

Click Search. The SLA audit information is listed.

4

Typically, only a subset of data is required for SLA audits. Remove all columns from the list except for the required columns. See 17.4 “To list and sort equipment information” (p. 689) for information.

Required information typically includes the following:

- Identification information, such as Site ID, Site Name, Slot ID, and Slot Name
- SLA audit and ordering information, such as Manufacture Date, Part Number, Serial Number, CLEI Code, and Card Type

5

Save the inventory output. See 17.5 “To save an inventory list” (p. 690) .

6

Close the form.

END OF STEPS

17.12 To export a network inventory file for adaptor modules and license details

17.12.1 Purpose

Perform this procedure to export a zipped network inventory file that contains the following network inventory files:

- Network_MediaAdaptor.csv
- Network_LicenseInformation.csv
- Network_NetworkElement.csv
- Network_Shelf.csv
- Network_Card.csv

Note: Only one export request can be handled at a time; the Export Inventory button is disabled until the file export completes.

17.12.2 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Click Export Inventory. A dialogue box opens with a reminder that the network inventory file generation and export is a time consuming activity.

3

Click OK.

4

Navigate to where you want to save the exported inventory file and click Save. A download status pop-up displays when the inventory file is downloaded.

5

Click OK.

6

Close the form.

END OF STEPS

18 Card migration

Card migration management

18.1 Overview

18.1.1 NFM-P Card Migration Event Manager

You can use an NFM-P GUI utility called the Card Migration Event Manager to facilitate the transition to an IOM 3 on one or more 7450 ESS or 7750 SR NEs. The Card Migration Event Manager transfers the existing IOM and MDA configurations to new modules with minimal service interruption.

The NFM-P Card Migration Event Manager can do the following:

- Upgrade from an IOM 1 or IOM 2 to an IOM3 while retaining the same MDAs.
- Upgrade from one MDA type to a newer MDA type that has compatible features.
- Upgrade an IOM and the associated MDAs in one operation.

You can perform an immediate migration, or preconfigure one or more migration events for later execution. A migration event can include an NE reboot to put the new hardware configuration into effect immediately after the migration. You can upgrade multiple IOMs and MDAs in one operation to limit the number of required NE reboots to one.

During a migration event, the NFM-P identifies the configured objects that are bound to the IOM or MDA, deletes the objects, and then creates the objects on the new IOM or MDA. The old and new configurations are saved in the NFM-P database until the migration completes successfully. The NFM-P preserves the statistics data and alarm information that is associated with each object.

18.1.2 Restrictions

The following restrictions apply to NFM-P card migration event management:

- **General**
  - Only the NFM-P admin user, or a user with an assigned administrator scope of command role, can create, modify, or execute a migration event.
  - The migration functions are not available to NFM-P OSS clients.
  - You may need to re-enable SNMP on an NE after the NE reboots following a migration event.
  - The description of a SAP associated with the card cannot contain a numerical value.
- **IOM-specific**
  - An IOM downgrade, for example, a migration from an IOM 3 to an IOM 2, is not supported.
- **MDA-specific:**
  - You cannot migrate an empty MDA slot to an MDA.
  - MDA migration is limited to MDAs of the same physical transmission type; for example, migration from a SONET MDA to an Ethernet MDA is not supported.
  - You cannot upgrade an MDA that is integrated with an IOM, for example, an IMM.
  - For MDAs, you can migrate only to an MDA of similar capacity that has the same or a greater number of ports.

The NFM-P raises an alarm during a migration event if a target NE is unreachable. After a migration event, the NFM-P raises an alarm to indicate migration success or failure.

Table 18-1 Supported MDA migration types

Current MDA	MDAs supported by migration
10 x 1-Gig Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
5 x 1-Gig Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
1 x 10-Gig Ethernet	4 x 10Gig Extended Performance XFP 2 x 10Gig Extended Performance XFP 1 x 10Gig Extended Performance SFP
20 x 100 Ethernet Fx	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
20 x 10/100/1000 Ethernet Tx	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
2 x 10-Gig Ethernet XFP	2 x 10Gig Extended Performance XFP 4 x 10Gig Extended Performance XFP
20 x 10/100/1000 Ethernet SFP	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
1 x 10-gig Ethernet XFP	4 x 10GigExtended Performance XFP 2 x 10Gig Extended Performance XFP 1 x 10Gig Extended Performance XFP
5 x 10/100/1000	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
10 x 10/100/1000 Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
60 x 10/100 Ethernet	48 x 10/100/1000 Ethernet Extended Performance TX

18.2 Workflow to manage card migration

18.2.1 Stages

Note: When you attempt a migration on an NE that does not have a recent configuration backup in the NFM-P database, the NFM-P raises an alarm and the migration fails.

Note: A recent configuration backup is a backup that occurs after the latest configuration save on the NE.

Note: Use the NFM-P to back up the device configuration of each target NE. See Chapter 23, “NE backup and restore” for information about performing NE backups.

1

Create a card migration event and execute or save it, as required. See 18.3 “To create a card migration event” (p. 700) for more information.

2

Execute a saved card migration event, as required. See 18.4 “To execute a saved card migration event” (p. 703) for more information.

18.3 To create a card migration event

18.3.1 Steps

1

Choose Tools→Card Migration Event Manager from the NFM-P main menu. The Card Migration Event Manager form opens.

2

Click Create. The Card Migration Event (Create) form opens with the General tab displayed.

3

Configure the general parameters.

4

Click on the NE Migration Candidates tab and click Add. The Select Network Elements - Card Migration Event form opens.

5

Select one or more NEs and click OK. The Select Network Elements - Card Migration Event form closes and the NEs are listed on the Card Migration Event (Create) form.

6

Select one or more NEs and click Properties. The Network Element Migration Candidates form opens.

7

Configure the Auto Reboot parameter and click OK.

8

Click on the Card Migration Candidates tab and click Add. The Select Cards - Card Migration Event form opens.

9

Select one or more cards and click OK. The Select Cards - Card Migration Event form closes and the cards are listed on the Card Migration Event (Create) form.

10

Select one or more cards and click Properties. The Migration Details (Create) form opens.

11

Configure the New Type parameter in the IOM, MDA 1, or MDA 2 panels, as required, and click OK.

12

To execute the migration event immediately, perform the following steps. Otherwise, you can execute the migration later by performing 18.4 “To execute a saved card migration event” (p. 703) .

1. Use the NFM-P to back up the device configuration of each target NE.

   Note:

   When you attempt a migration on an NE that does not have a recent configuration backup in the NFM-P database, the NFM-P raises an alarm and the migration fails.

   A recent configuration backup is a backup that occurs after the latest configuration save on the NE.
2. Click Apply. The NFM-P saves the card migration event configuration.
3. Click on the General tab to display the Status indicator.
4. Click Initiate Migration and click Yes. The NFM-P starts to migrate the specified cards.
5. Monitor the card migration as it progresses by viewing the Status indicator on the form. The status can be one of the following:
   - Awaiting manual reboot to complete migration
   - Failed - Latest configuration not available
   - Failed - Unable to migrate configuration
   - Failed - Unable to reboot network element
   - Failed - Unable to transfer migrated configuration
   - Migration completed
   - Not Started
   - Rebooted network element
   - Started
   - Swap failed
   - Swap failed on some network elements

13

Save your changes and close the form.

END OF STEPS

18.4 To execute a saved card migration event

18.4.1 Steps

1

Use the NFM-P to back up the device configuration of each target NE.

Note: When you attempt a migration on an NE that does not have a recent configuration backup in the NFM-P database, the NFM-P raises an alarm and the migration fails. A recent configuration backup is a backup that occurs after the latest configuration save on the NE.

2

Choose Tools→Card Migration Event Manager from the NFM-P main menu. The Card Migration Event Manager form opens.

3

Select a card migration event and click Properties. The Card Migration Event (Edit) form opens with the General tab displayed.

4

Click Initiate Migration and click Yes. The NFM-P starts to migrate the specified cards.

5

Monitor the card migration as it progresses by viewing the Status indicator on the form. The status can be one of the following:

- Awaiting manual reboot to complete migration
- Failed - Latest configuration not available
- Failed - Unable to migrate configuration
- Failed - Unable to reboot network element
- Failed - Unable to transfer migrated configuration
- Migration completed
- Not Started
- Rebooted network element
- Started
- Swap failed
- Swap failed on some network elements

6

Save your changes and close the form.

END OF STEPS

19 TCA

TCA management using the NFM-P

19.1 Overview

19.1.1 Functional description

You can create policies to raise NFM-P alarms based on managed-object statistics. When a statistic counter of an object such as an interface reaches a specified threshold, the NFM-P alerts GUI operators using a threshold-crossing alarm, or TCA, based on policy specifications.

A TCA policy includes the following:

- a list of objects to monitor
- a traffic flow direction to monitor
- rules that define the rising and falling thresholds, and the alarm severity
- optional custom rules to monitor statistics based on user-defined formulas

Note: A TCA is not self-clearing. The NFM-P clears a TCA only when the TCA policy contains a falling-threshold rule in addition to a rising-threshold rule, and the alarm severity in the falling-threshold rule is set to cleared.

19.1.2 TCA operation

To enable TCA on an object, you must enable performance statistics collection on the object. See the *NSP NFM-P Statistics Management Guide* for information about enabling performance statistics collection.

The NFM-P compares the object counter values at each statistics collection to the threshold values in the associated TCA policy. When a value initially crosses a threshold, the NFM-P raises an alarm. Rather than raise new alarms for successive threshold-crossing events associated with the policy, the NFM-P adjusts the alarm severity based on the policy rules.

Note: If required, you can use a parameter on the TCA policy properties form to disable the alarm-severity demotion defined in a TCA policy. For example, an operator may want a critical alarm associated with an object to persist and be visible to other operators until the fault condition is cleared.

Using the NFM-P client GUI, you can associate a TCA policy with a managed object from the TCA policy configuration form, or from the object properties form. You can associate one TCA policy with multiple objects.

19.1.3 TCA policy rules

Depending on the object type, you can use rules to monitor the following:

- object utilization
- errored packets

- dropped packets
- custom values; see 19.1.4 “Custom profile TCAs” (p. 705)

A rule can apply to an absolute counter value or to a delta value, which is the difference between the current value and the previous value.

In a TCA policy rule, you can specify whether the rule raises alarms. Disabling the raising of alarms according to a policy rule may be required, for example, when the rule raises excessive alarms because of a fault that causes frequent threshold crossing. You can disable the raising of alarms while the fault is being resolved, and then re-enable the raising of alarms when the affected object returns to normal operation.

19.1.4 Custom profile TCAs

You can optionally create custom profile TCAs that define the TCA rules using one or more of the following in a mathematical formula:

- monitored object statistics counters
- monitored object properties
- numerical constants
- mathematical operators

Note: For monitored object property values in a formula, the current and previous values are the same. If a monitored object property value changes between collections, the new property value is used as the current and previous values.

After you create a custom profile TCA, you associate the custom profile TCA with a TCA policy to put the TCA into effect.

19.1.5 TCA configuration example

The following example describes the NFM-P configuration required to raise a TCA, adjust the alarm severity, and clear the alarm, based on the following utilization specifications:

- Raise a minor alarm if utilization is between 60% and 69%.
- Change the alarm severity to major if utilization rises to between 70% and 79%.
- Change the alarm severity to critical if utilization rises to 80% or higher.
- Change the alarm severity to major if utilization falls below 80%.
- Change the alarm severity to minor if utilization falls below 70%.
- Clear the alarm if utilization falls below 60%.

The following are the TCA rules and associated parameter values required to implement the example alarm behavior:

Rule to raise minor alarm for 60% or higher utilization:

- Alarm Severity—minor
- Threshold (%)—60
- Threshold Direction—Rising Above

Rule to change alarm severity to major for 70% or higher utilization:

- Alarm Severity—major
- Threshold (%)—70
- Threshold Direction—Rising Above

Rule to change alarm severity to critical for 80% or higher utilization:

- Alarm Severity—critical
- Threshold (%)—80
- Threshold Direction—Rising Above

Rule to change alarm severity to major if utilization falls below 80%:

- Alarm Severity—major
- Threshold (%)—80
- Threshold Direction—Falling Below

Rule to change alarm severity to minor if utilization falls below 70%:

- Alarm Severity—minor
- Threshold (%)—70
- Threshold Direction—Falling Below

Rule to clear alarm if utilization falls below 60%:

- Alarm Severity—cleared
- Threshold (%)—60
- Threshold Direction—Falling Below

19.2 Workflow to configure TCA

19.2.1 Purpose

The following is the sequence of high-level actions required to configure TCA and enable the TCA function on a monitored object.

19.2.2 Stages

1

Enable performance statistics collection on the object to monitor. Table 19-1, “Port statistics policy objects for TCA” (p. 707) lists the port statistics policy objects required to enable TCA on a network interface or physical link.

Table 19-1 Port statistics policy objects for TCA

Statistics policy type	MIB	MIB entry	Monitored class
NE MIB	IF-MIB	ifEntry	—
Specific MIB			equipment.PhysicalPort
Interface Additional Stats		ifXEntry	lag.Interface

2

If you require custom TCA rules, configure a custom profile TCA; see 19.3 “To configure a custom profile TCA” (p. 707) .

3

Configure a TCA policy; see 19.4 “To configure a TCA policy” (p. 709) .

1. Configure utilization, error, and drop TCA rules.
2. Optionally, associate a custom profile TCA with the policy.

4

Apply the TCA policy to an object; see 19.5 “To apply a TCA policy to objects using the object properties forms” (p. 711) .

5

If required, customize the TCA system preferences, such as specifying the default TCA severity or the maximum TCA limit; see the section on system preferences configuration procedures in the *NSP System Administrator Guide*.

19.3 To configure a custom profile TCA

19.3.1 Steps

1

Choose Tools→TCA Policies from the NFM-P main menu. The TCA Policies form opens.

2

Choose Custom Profile TCA (TCA Policy) and click Create. The Custom Profile TCA (Create) form opens.

3

Configure the general parameters.

4

Select a monitored object type.

5

Configure the Stats Type parameter.

6

Click Build Formula. The Build Formula form opens.

7

Choose a statistic counter from the Counter Field drop-down menu and click Add.

8

Add one or more values and operators to the formula, as required.

Note: To enclose an expression in parentheses, select the expression and click on the ( ) operator.

a. Add a statistic counter for the monitored object.
   1. Click on a mathematical operator.
   2. Choose a counter from the Counter Field drop-down menu and click Add.

b. Add a property of the monitored object.
   1. Click on a mathematical operator.
   2. Choose a property from the Class Field drop-down menu and click Add.

c. Add a constant value.
   1. Click on a mathematical operator.
   2. Enter a number in the Constant Value field and click Add.

9

Click OK to save your changes and close the Build Formula form.

10

Click OK to save your changes and close the Custom Profile TCA (Create) form.

11

Close the TCA Policies form.

END OF STEPS

19.4 To configure a TCA policy

19.4.1 Steps

1

Choose Tools→TCA Policies from the NFM-P main menu. The TCA Policies form opens.

2

Choose TCA Policy (TCA Policy) and click Create. The TCA Policy (Create) form opens.

3

Configure the general parameters.

4

Select a monitored object and click Apply.

Note: The panels displayed on the General tab vary, depending on the type of monitored object. If the monitored object type you need does not appear in the list, you may need to create a custom TCA profile first; see 19.3 “To configure a custom profile TCA” (p. 708).

5

To configure utilization TCA rules:

1. Configure the parameters in the Utilization panel.
2. Click Create in the Rules panel. The PercentageTCARule, tca.UtilizationTCA-0 (Create) form opens.
3. Configure the parameters and click OK.

   Note:

   You cannot create two rules that have the same Threshold (%) and Threshold Direction values.

6

To configure error TCA rules:

1. Configure the parameters in the Error panel.
2. Click Create in the Rules panel. The PercentageTCARule, tca.ErrorTCA-0 (Create) form opens.
3. Configure the parameters and click OK.

   Note:

   You cannot create two rules that have the same Threshold (%) and Threshold Direction values.

7

To configure drop TCA rules:

1. Configure the parameters in the Drop panel.
2. Click Create in the Rules subpanel. The PercentageTCARule, tca.DropTCA-0 (Create) form opens.
3. Configure the parameters and click OK.

   Note:

   You cannot create two rules that have the same Threshold (%) and Threshold Direction values.

8

To configure custom TCA rules:

1. In the Custom panel, click Create. The CustomTCA (Create) form opens.
2. Configure the parameters.

3. Select a Custom TCA Profile.
4. Click Create in the Rules panel. The CustomTCARule, tca.CustomTCA-0 (Create) form opens.
5. Configure the parameters and click OK.

9

Apply the TCA Policy to one or more monitored objects:

1. Click on the Monitored Objects tab.
2. Click Add. The Select form opens.
3. Select the required objects and click OK.

10

Click OK to save your changes and close the TCA Policy (Edit) form.

11

Close the TCA Policies form.

END OF STEPS

19.5 To apply a TCA policy to objects using the object properties forms

19.5.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Select an object type and click Search. A list of equipment objects is displayed.

3

Select an object and click Properties. The object properties form opens.

4

If the object is a Physical Link, assign a TCA policy to End Point A and End Point B. Go to Step 7 .

5

Click on the TCA tab and click Assign. The Select TCAPolicy form opens.

6

Select a TCA policy and click OK

7

Click OK to save your changes and close the object properties form.

8

Close the Manage Equipment form.

END OF STEPS

20 Bulk operations

Bulk operations using the NFM-P

20.1 Overview

20.1.1 Functional description

The NFM-P bulk operations function allows you to create bulk changes to modify a large amount of information. Bulk changes may be deployed to NEs or restricted to the NFM-P.

A bulk change can contain thousands of target objects. These objects, referred to as batch items, are grouped into batches for efficient bulk change execution. An operator can execute an entire bulk change or individual batches.

You should regenerate batches each time you execute a bulk change to ensure that each target object that matches the bulk change filter is modified.

If one or more of the batch items fail to change during a bulk change operation, the Batch Status and the Batch Status Summary parameters help identify the items. You can also use the Task Manager to view information about a batch.

The NFM-P displays a confirmation message before executing bulk change operations. You can configure user preferences to disable confirmation messages for bulk change operations; see 1.23 “To configure NFM-P user preferences” (p. 116) .

20.2 Workflow to manage bulk operations

20.2.1 Stages

1

Create a bulk change. See 20.3 “To create a bulk change” (p. 714) .

2

Generate batches for a bulk change. See 20.5 “To execute a bulk change” (p. 716) .

3

Execute one or more batches in a bulk change. See 20.5 “To execute a bulk change” (p. 716) .

4

As required, stop a bulk change. See 20.7 “To stop one or more bulk changes” (p. 718) .

5

View the bulk change execution results to verify that the bulk change was successful and to identify any items that failed to change. See 20.6 “To view executed batch information” (p. 717) .

20.3 To create a bulk change

20.3.1 Steps

1

Choose Tools→Bulk Operations from the NFM-P main menu. The Bulk Operations form opens.

2

Open the Create Bulk Change form by performing one of the following steps.

a. Click Create.

b. To create a bulk change by copying values from an existing bulk change:
   1. Choose an existing bulk change and click Properties. The Bulk Change (Edit) form opens.
   2. Click Copy.

3

Configure the general parameters and click Next. The Specify the filter form opens.

4

To specify a filter, perform one of the following.

a. Create a filter.
   1. Click on the Attribute drop-down menu and select an option.
   2. Click on the Function drop-down menu and select an option.
   3. Click on the Value drop-down menu and select or enter an option.
   4. Click Add. The options are displayed in the filter panel.
   5. Select AND or OR from the Operators drop-down menu. Repeat 1 to 4 to add more options.
   6. Click Save. The Save Filter form opens.
   7. Configure the general parameters.
   8. Click Save. The Save Filter form closes.

b. Load a filter.
   1. Click Saved Filters. The Saved Filters form opens.
   2. Select a saved filter and click Load. The Saved Filters form closes and the loaded filter is displayed in the filter panel.

5

Configure the filter criteria. Creating a filter that contains the attributes that you want to change can limit the number of objects generated for the bulk change.

6

Click Count to determine the number of network objects that are affected, based on the filter.

7

Click View to list the target objects. The Filtered List form opens.

8

Click Search. A list of target objects is displayed.

9

Click Next. The Specify the attributes form opens.

10

Double-click on the attribute or attribute group in the attribute tree panel that you want to include in the bulk change. The attribute or attribute group appears in the display panel.

If you included an attribute or attribute group that you want to exclude from the bulk change, you can remove it by clicking the Remove X icon in the display panel.

11

Perform one of the following for each listed attribute, as required.

a. Enter a value in the attribute field. The drop-down menu changes from Unchanged to Set, which means that the attribute changes to the new value when you execute the bulk change.

b. Choose Default from the drop-down menu to specify that the attribute is to change to the default value when you execute the bulk change.

c. Choose Clear from the drop-down menu. Depending on the object type, the attribute value is cleared or set to the default value when you execute the bulk change.

12

Click Next. The Change review form opens.

13

Configure the required parameters and click Finish.

14

Select the View the newly created Bulk Change check box to view the bulk change after closing the form, if required.

15

Click Close.

16

To create more Bulk Changes, repeat Step 2 to Step 15 .

END OF STEPS

20.4 To modify a bulk change

20.4.1 Steps

1

Choose Tools→Bulk Operations from the NFM-P main menu. The Bulk Operations form opens.

2

Choose a bulk change and click Properties. The Bulk Change (Edit) form opens with the General tab displayed.

3

Modify the required parameters.

4

Click on the Spans tab and specify a span, if required.

5

Save your changes and close the form.

END OF STEPS

20.5 To execute a bulk change

Note: You can only execute one bulk change at a time. All other bulk changes are queued.
You can only generate batches for one bulk change at a time.
You cannot generate batches when a bulk change is executing.

20.5.1 Steps

1

Choose Tools→Bulk Operations from the NFM-P main menu. The Bulk Operations form opens.

2

Choose a bulk change from the list and perform one of the following.

a. If the bulk change contains batches, click Execute to execute all of the batches.

Note: Changes made to the network after the batches are generated are not affected by the bulk change batches.

b. If the bulk change does not contain batches, or you want to regenerate the batches, perform the following.

1. Click Properties. The Bulk Change (Edit) form opens with the General tab displayed.

2. Click on the Batch Control tab.

   If the Generate Batches button is dimmed, set the Admin State parameter on the General tab to Enable.
3. Click Generate Batches. The batches are generated and listed on the form.
4. To execute only specific batches, choose one or more batches and click Execute Selected.
5. To execute all of the batches, click Execute All.

3

Click Yes. The NFM-P executes the bulk change.

4

Close the forms.

END OF STEPS

20.6 To view executed batch information

Note: You can also use the NFM-P Task Manager to view information about a batch.

20.6.1 Steps

1

Choose Tools→Bulk Operations from the NFM-P main menu. The Bulk Operations form opens.

2

Choose a bulk change and click Properties. The Bulk Change (Edit) form opens with the General tab displayed.

3

Click on the Batch Control tab.

4

Choose a batch and click Properties. The Bulk Change Batch form opens with the General tab displayed. View general information about the executed batch.

5

Click on the Batch Items tab and click Search. A list of items that were changed in the batch appears with the status of each batch item.

6

To view the parameters for each batch item, choose a batch item and click Properties.

7

Close the forms.

END OF STEPS

20.7 To stop one or more bulk changes

20.7.1 Steps

1

Choose Tools→Bulk Operations from the NFM-P main menu. The Bulk Operations form opens.

2

Perform one of the following:

a. Choose one or more bulk change to stop.

b. Choose one or more batches to stop by performing the following:

1. Choose a bulk change from the list and click Properties. The Bulk Change (Edit) form opens with the Batch Control tab displayed.
2. Choose one or more batches.

3

Click Stop and click Yes. The bulk change execution stops.

4

Close the forms.

END OF STEPS

21 Serial raw sockets for IP transport services

Creating serial raw sockets for IP transport services using the NFM-P

21.1 Serial socket and IP transport services

21.1.1 Overview

Raw socket for serial data transport is a method of transporting serial data over an IP network using L3-based services (IES or VPRN). A raw socket allows direct sending and receiving of IP packets between network devices without any protocol-specific transport layer formatting.

Through local-host and remote-host configurations on an NE, the serial socket IP transport sessions are established to carry serial data over an IP/MPLS network. Incoming serial data is packetized into L3 TCP/UDP packets which ingress IES or VPRN L3 SAPs and are forwarded across an IP/MPLS network. These packets belong to individual sessions between an RTU and SCADA management system/FEP, and are transported as an IP transport sub-service within an IES or VPRN context. The IP transport sub-service uses the local host and remote host details, and port numbers (session parameters) to establish TCP/UDP sessions.

Figure 21-1 Sample serial data transport configuration

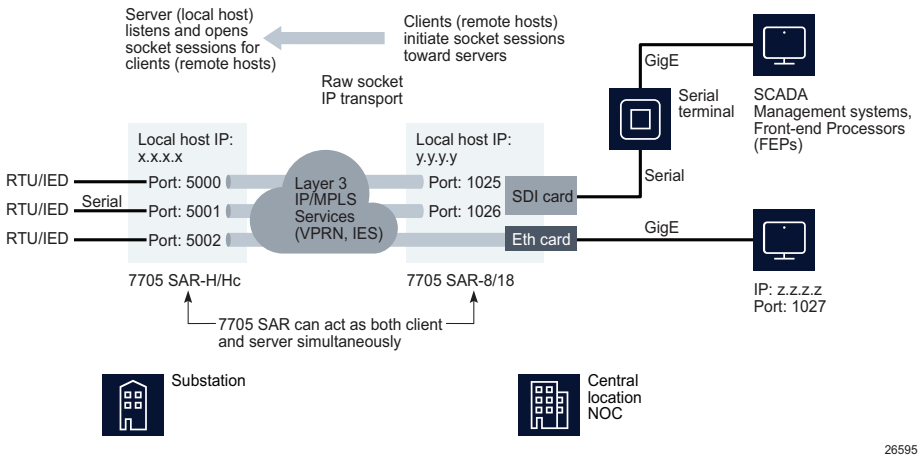

For example, SCADA data from RTUs can be transported to FEPs or SCADA controlling units. The serial port-facing RTU is configured as both server and client. Unsolicited messages are communicated by the RTU as a client to the FEPs. When the FEPs request data, the RTUs respond as a server. This transport method is an alternative to C-pipes where the transport medium is an L3 IP/MPLS service instead of L2. For more information about SCADA, see 15.2 “SCADA on the 7705 SAR” (p. 455).

This functionality is supported on the 7705 SAR-8, 7705 SAR-18, 7705 SAR-H and 7705 SAR-Hc, release 8.0 R4 or later, and the 7705 SAR-Hm, Release 15.0 R4 or later.

The benefits of serial sockets are:

- Lower bandwidth consumption - serial sockets only send data when needed.
- Multiple head end FEPs can poll RTUs at the same time which avoids serial data collisions from both FEPs at the RTU serial link. Serial sockets can be used to queue and manage data streams and respond back to each at once.
- A single RTU can send a single data stream to multiple FEPs for unsolicited polls that originate from the RTU.
- There is no requirement for terminal servers at the FEP locations to convert IP to serial data and back. SCADA controlling units can connect to the head end router via Ethernet to send and receive serial data to and from the RTU.

Serial socket and IP transport configuration setup

1. A 7705 SAR NE with a serial port is connected to local hosts (RTUs) through serial ports acting either as client, server, or both client and server.
2. Serial data over an IES or VPRN L3 interface is carried from the 7705 SAR NEs through another 7705 SAR over IP/MPLS cloud.
3. The remote hosts (SCADA management systems/FEPs) can be connected in one of two ways:
   - Through serial ports of a serial data card on a 7705 SAR-8 and 7705 SAR-18 connected to serial terminals, which is then connected to FEPs/SCADA management systems through GigE interfaces.
   - Through direct TCP/UDP sessions from RTU servers to FEPs/SCADA management systems through GigE interfaces with no serial terminal in between.

   The source and destination IP addresses and port numbers for the IP transport sessions are derived from the local/remote-host configurations associated with each serial port.

Figure 21-2 Sample IES/VPRN IP transport configuration

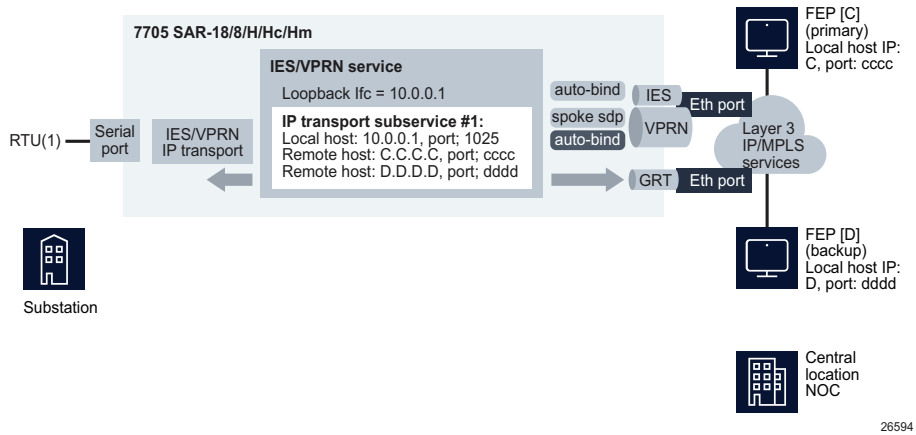

Serial raw socket squelch

A condition may occur where the end device connected to the serial port sends out a continuous stream of data, stopping the far end raw socket or controlling unit from receiving data from other end devices in the network. To resolve this condition, enable the squelch command on the raw socket at the port level (disabled by default) to suppress the socket from receiving any more data from the device connected to the port.

The squelch function is enabled by configuring a delay period that specifies the delay between when the start of a constant character stream begins (e.g. a new packet) and when to start the squelch function. When a constant character stream occurs for the delay timer value, the port is considered locked up, an alarm is raised indicating the lock-up, and the squelching function is triggered. The range of the delay between the alarm and the squelching function is from one to 120s.

To recover from squelch, there are two methods:

1. Manually clear the squelch function for the serial port and put the port in a normal state.
2. Automatically clear the squelch function after a configurable amount of time (one to 120 s) after the squelch function was instantiated to put the serial port back into a normal state.

Socket profile mismatch

A socket profile mismatch occurs when the local properties of a serial raw socket on a 7705 SAR-Hm node with an assigned socket profile no longer match the values defined in the profile. The Socket Profile Mismatch parameter on the RS-232 Socket tab of the port configuration form becomes enabled, and the ReApply button is available. To resolve the mismatch, click on the ReApply button in the Socket Profile panel. The mismatched parameters on the port are configured with the values from the profile.

21.2 Workflow to configure serial raw sockets for IP transport services

21.2.1 Stages

The following workflow describes the sequence of high-level tasks required to configure serial raw sockets for IP transport. This workflow assumes that the physical devices have been installed, commissioned, and discovered.

1

Create a socket profile and optionally, associate the Socket profile to a physical port on a 7705 SAR-Hm node. See 21.4 “To create a socket profile” (p. 723).

2

Configure the port level socket which includes serial link parameters and socket level configurations such as end-of-packet checking parameters and the inter-sessions delay for transmitting sessions data out of the serial link. See 21.5 “To configure a serial raw socket on a 7705 SAR” (p. 724).

3

As required, configure the remote-hosts for IP transport using the Local Host Entity Manager. See 21.6 “To configure a global entry using the Local Host Entity Manager” (p. 725) and 21.7 “To modify a global entry” (p. 726).

4

Configure IP transport services to associate the serial port within either an IES or VPRN SAP allowing TCP/UDP encapsulated serial data to be routed within the corresponding Layer 3 service. This includes a loopback interface for IP transport local-hosts (servers), IP transport local-host and remote-hosts configuration, and the associated serial port SAP. See 21.8 “To configure IP transport on an IES site” (p. 727) or 21.9 “To configure IP transport on a VPRN site” (p. 729).

5

If required, you can use the Check TCP Connection tool to test the health of a TCP connection from the client to the server. See 21.8 “To configure IP transport on an IES site” (p. 727) or 21.9 “To configure IP transport on a VPRN site” (p. 729).

21.3 Serial IP transport using the Local Host Entity Manager

21.3.1 Overview

Serial transport over raw sockets can be either between two 7705 SAR NEs (one at an RTU and one at the controlling unit), or between a 7705 SAR NE and another NE such as a 7750 SR. These types of local-hosts do not require a 7705 SAR NE or any IP transport services to be configured. These would typically be 3rd party devices such as FEPs, SCADA controlling units, other terminal servers, or other devices not managed by the NFM-P.

For IP transport, the remote-host can be one of the following:

1. For a 7705 SAR NE, the remote-host can be another local-host IP that belongs to a different site in the same service.
2. For a non-7705 SAR NE, the remote-host can be another IP address that is globally reachable to the NE (via routing), but does not belong to the same service.

To facilitate adding remote-hosts to an IP transport, the Local Host Entity Manager is used. The Local Host Entity Manager lists the following:

- All local-hosts from all services/sites managed by the NFM-P, called discovered entries.
- A list of user-created global host entries called global entries.

When an NE is managed by the NFM-P, it checks the NE for all configured IP transports. The local-host details (IP Address and Port Number parameters) from all of the IP transports are populated in the Local Host Entity Manager. The remote-host details are not populated.

When the local-host entries are populated in the Local Host Entity Manager, you can click the Add Remote Hosts button on the IP Transport Remote Host Properties form which will list all possible entries from the Local Host Entity Manager that are valid remote-hosts for a particular local-host.

Serial raw socket and IP transport procedures

21.4 To create a socket profile

21.4.1 Purpose

You can use the Tools→Socket Profile menu on the NFM-P GUI to create a reusable template to configure serial port parameters for raw sockets such as parity, character size, start-stop bits, squelch delay, and EOP (end-of-packet) length, timeout, and special characters.

21.4.2 Steps

1

Choose Tools→Socket Profile from the NFM-P main menu. The Socket Profile form opens.

2

Click Create, or select an existing socket profile and click Properties. The Socket Profile (Create | Edit) form opens.

3

Configure the parameters on the General tab as required.

4

Click on the Socket Specifics tab and configure the parameters as required.

5

Save your changes and close the form.

6

Optionally, associate the socket profile to a physical port on a 7705 SAR-Hm node:

1. In the navigation tree equipment view, navigate to a port on a 7705 SAR-Hm node (Shelf→Card Slot→Daughter Card→[Supported] Port) and choose Properties. The Physical Port (Edit) form opens.
2. On the RS-232 Socket tab, select a socket profile to associate with the 7705 SAR-Hm port and click OK. The parameters in the Serial RS-232 and Serial Socket panels automatically update with the parameter settings defined in the socket profile.

   Note: If the socket profile properties are changed after you associate the socket profile to the port, the Socket Profile Mismatch parameter located on the RS-232 Socket tab, Socket Profile panel, is enabled to indicate the socket profile parameters have been modified without the changes being applied to the port. Also, if the serial port/raw socket parameters are changed without changing the socket profile properties, a mismatch will occur between the socket profile and the actual configuration. Clicking on the ReApply button updates the parameter settings for the port to match the socket profile.

3. Save your changes and close the form.

End of steps

21.5 To configure a serial raw socket on a 7705 SAR

21.5.1 Before you begin

A raw socket IP transport interface can be configured for each RS-232 serial port on an NE. IP transport interfaces are only supported on RS-232 serial ports with the Encap Type parameter configured as RAW. This allows the serial port to receive TCP connections or UDP session packets from multiple remote-hosts when operating as a server, or to create new multiple sessions to remote-hosts in order to send and receive serial data when operating as a client.

Configuration of serial raw sockets is supported on the 7705 SAR-8, 7705 SAR-18, 7705 SAR-H, and 7705 SAR-Hc, Release 8.0 R4 or later.

There are two main configurations required for a serial raw socket to be operational and support sending/receiving serial data in either server mode, client mode, or both:

1. Port level socket configuration - This includes serial link parameters such as baud rate, start/stop, bits, etc. Socket level configurations are required such as end-of-packet checking parameters and the inter-sessions delay for transmitting sessions data out of the serial link.
2. IP transport services configuration – Creating an IP transport service is required to associate the serial port within either an IES or VPRN so that TCP/UDP encapsulated serial data can be routed within the corresponding Layer 3 service. This configuration includes IP transport local-host items (local servers) and remote host configurations, TCP timers and sessions control, etc. See 21.8 “To configure IP transport on an IES site” (p. 727) and 21.9 “To configure IP transport on a VPRN site” (p. 729) for more information.

21.5.2 Steps

1

Configure the applicable MDA for the 7705 SAR NE. See 15.78 “To configure an MDA” (p. 536) for more information.

The 7705 SAR-8 and 7705 SAR-18 must be configured with a 12-port Serial Data v2 or v3 MDA.

The 7705 SAR-H must be configure with a 2 × Serial Data + 2 × Channelized DS1/E1 MDA.

The 7705 SAR-Hc has fixed MDAs that cannot be modified.

2

In the navigation tree Equipment view, right-click on the port and choose Create Channel/Socket. The Serial Channel (Create) form opens.

3

Click on the States tab and configure the Administrative State parameter as Up.

4

Click Apply and confirm your action.

5

Click on the General tab and enable the Serial Socket Enabled parameter.

Note: When you enable the Serial Socket Enabled parameter, the Channel tab is no longer available and the RS-232/Socket tab appears. The NFM-P will automatically set the Device Mode parameter as Asynchronous and the Gender parameter as DCE, on the RS-232/Socket tab.

6

Select a Socket Profile ID in the Socket Profile panel and click Apply.

7

Click on the RS-232/Socket tab and configure the required parameters.

8

Save your changes and close the form.

You can click on the IP Transports tab to view the IP Transport to which this serial raw socket is associated.

END OF STEPS

21.6 To configure a global entry using the Local Host Entity Manager

21.6.1 Before you begin

Consider the following regarding global and discovered entries:

- You can modify the Name and Description parameters for both global and discovered entries.
- You can modify the IP Address and Port Number parameters for global entries, but not for discovered entries. The modification of the IP Address and Port Number parameters of a global entry is propagated to all matching remote-hosts in the network. These matching remote-hosts can be viewed on the IP Transport tab of the Global Entry properties form.
- Discovered entries cannot be deleted. They will automatically disappear when the corresponding local-host is deleted.
- Whenever the IP Address and Port Number of an IP transport local-host is modified using CLI or the NFM-P, the old discovered entry with the old IP Address and Port Number is deleted and a new discovered entry is added automatically, with a modified IP Address and Port Number.
- All discovered entries from an NE are automatically deleted if the node is unmanaged/deleted from the NFM-P.

21.6.2 Steps

1

Choose Tools→Local Host Entity Manager from the NFM-P main menu. The Local Host Entity Manager form opens.

2

Click Create. The Network Wide Local Host (Create) form opens.

3

Configure the required parameters on the General tab and click Apply. The global entry is listed on the Local Host Entity Manager form.

Once you have added remote-hosts to an IP transport, you can click on the IP Transports tab to view the associations of the IP transports and their remote-hosts. Associations can be seen for both global and discovered entries. Associations are dynamic; the association is automatically removed if an entry is deleted as a remote-host for its IP Transport.

4

Close the forms.

END OF STEPS

21.7 To modify a global entry

21.7.1 Before you begin

Global entries are user-created entries in the Local Host Entity Manager. These entries (IP Address and Port Number) are generally those of either third party servers, or NEs that are un-managed in the NFM-P. The NFM-P allows modification of global entries. When the IP Address and Port Number parameters of a global entry are modified, these parameters are modified correspondingly across all matching remote-hosts in the network. This would allow seamless migration among third part servers.

Consider the following:

- This procedure is applicable only to global entries. The IP Address and Port Number parameters of discovered entries cannot be modified.
- Modification of the IP Address and Port Number parameters of a global entry is propagated to all matching remote-hosts in the network. These “matching” remote-hosts can be viewed on the IP Transports tab of the Global Entry properties form.
- Modification of the Name and Description parameters remains a local change in the global/discovered entry and they are not propagated across matching remote-hosts in the network.
- Deletion of a global entry will only delete the entry from the NFM-P. It will not delete matching remote-hosts in the network.

21.7.2 Steps

1

Choose Tools→Local Host Entity Manager from the NFM-P main menu. The Local Host Entity Manager form opens.

2

Select a global entry in the list and click Properties. The Network Wide Local Host (Edit) form opens.

3

On the General tab, configure the IP Address and Port Number parameters.

4

Click OK to confirm your actions.

5

Close the form.

END OF STEPS

21.8 To configure IP transport on an IES site

21.8.1 Purpose

IP transport service configurations are required to associate a serial port within either an IES or VPRN SAP, allowing TCP/UDP encapsulated serial data to be routed within the corresponding Layer 3 service.

Before you begin, ensure that you configure the port level serial socket on the serial data card of the applicable 7705 SAR NE. See 21.5 “To configure a serial raw socket on a 7705 SAR” (p. 724) for more information.

The 7705 SAR-Hm uses serial ports instead of raw sockets. The serial ports are 1/3/1 and 1/3/2. The 7705 SAR-Hm supports IP transport with a serial port or GNSS. One IP transport with GNSS can be configured per 7705 SAR-Hm. Multiple IP transports with serial ports and GNSS can exist within the same service.

For information about configuring IP transport on a VPRN site, see 21.9 “To configure IP transport on a VPRN site” (p. 729).

Note: When an IES IP transport is operationally down with an existing IP Transport Down Reason, a second IP Transport Down Reason is not automatically updated in the NFM-P. When this occurs, the IP transport requires a manual resynchronization.

21.8.2 Steps

1

Create an IES site. See 78.5 “To configure an IES site” (p. 2431).

2

Click on the IP Transports tab and click Create. The IES IP Transport (Create) form opens.

3

Configure the general parameters as required.

4

Select a transport port in the Transport Port panel.

Note: You must configure a serial socket on the 7705 SAR SDI MDA port in order to choose a transport port. See 21.5 “To configure a serial raw socket on a 7705 SAR” (p. 724) for more information.
For a 7705 SAR-Hm NE, choose one of the serial ports, 1/3/1 and 1/3/2, or choose GNSS.

5

Configure the IP address for the local host in the Local Host panel.

6

Configure the required parameters in the Session Details and TCP panels.

7

Click on the Remote Host tab and click Create. The IP Transport Remote Host (Create) form opens.

8

Configure the required parameters and click OK.

9

Click on Add Remote Hosts. The Select Network Wide Local Host form opens. All applicable remote-host entries from the Local Host Entity Manager are listed.

10

Select one or more entries that you want to add as remote-hosts. Click Apply and then OK.

The remote-host IDs are auto-assigned by the NFM-P.

11

Click OK to save your changes.

12

If required, use the Check TCP Connection tool to test the health of a TCP connection from the client to the server.

1. Click on the Remote Host tab.
2. Select a remote-host in the list and click Properties. The IP Transport Remote Host (Edit) form opens with the General tab displayed.
3. Configure the Check TCP Connection parameter in the TCP Test panel and click OK. The Result parameter shows if the TCP test failed or succeeded.

Note: When you perform a TCP connection test on a remote-host, the result is not updated automatically. You must perform a manual resynchronization of the remote-host in order to update the result.

13

Save your changes and close the forms.

END OF STEPS

21.9 To configure IP transport on a VPRN site

21.9.1 Purpose

IP transport services configurations are required to associate a serial port within either an IES or VPRN SAP allowing TCP/UDP encapsulated serial data to be routed within the corresponding Layer 3 service.

Before you begin, ensure that you configure the port level serial socket on the serial data card of the applicable 7705 SAR NE. See 21.5 “To configure a serial raw socket on a 7705 SAR” (p. 724).

The 7705 SAR-Hm uses serial ports instead of raw sockets. The serial ports are 1/3/1 and 1/3/2. The 7705 SAR-Hm supports IP transport with a serial port or GNSS. One IP transport with GNSS can be configured per 7705 SAR-Hm. Multiple IP transports with serial ports and GNSS can exist within the same service.

See 21.8 “To configure IP transport on an IES site” (p. 727) for information about configuring IP transport on an IES site.

Note: When a VPRN IP transport is operationally down with an existing IP Transport Down Reason, a second IP Transport Down Reason is not automatically updated in the NFM-P. When this occurs, the IP transport requires a manual resynchronization.

21.9.2 Steps

1

Create a VPRN site. See 79.11 “To configure a VPRN site” (p. 2545).

2

Click on the IP Transports tab and click Create. The VPRN IP Transport (Create) form opens.

3

Configure the general parameters as required.

4

Select a transport port in the Transport Port panel.

Note: You must configure a serial socket on the 7705 SAR serial data MDA port in order to choose a transport port. See 21.5 “To configure a serial raw socket on a 7705 SAR” (p. 724) for more information.
For a 7705 SAR-Hm NE, choose one of the serial ports, 1/3/1 and 1/3/2, or choose GNSS.

5

Configure the IP address for the local host in the Local Host panel.

6

Configure the required parameters in the Session Details and TCP panels.

7

Click on the Remote Host tab and click Create. The IP Transport Remote Host (Create) form opens.

8

Configure the required parameters and click OK.

9

Click on Add Remote Hosts. The Select Network Wide Local Host form opens. All applicable remote-host entries from the Local Host Entity Manager are listed.

10

Select one or more entries that you want to add as remote-hosts. Click Apply and then OK.

The remote-host IDs are auto-assigned by the NFM-P.

11

Click OK to save your changes.

12

If required, use the Check TCP Connection tool to test the health of a TCP connection from the client to the server.

1. Click on the Remote Host tab.

2. Select a remote-host in the list and click Properties. The IP Transport Remote Host (Edit) form opens with the General tab displayed.
3. Configure the Check TCP Connection parameter in the TCP Test panel and click OK. The Result parameter shows if the TCP test failed or succeeded.

Note: When you perform a TCP connection test on a remote-host, the result is not updated automatically. You must perform a manual resynchronization of the remote-host in order to update the result.

13

Save your changes and close the forms.

END OF STEPS

Part III: NE maintenance

Overview

Purpose

This part provides information about NE maintenance using the NFM-P.

Contents

22 NE maintenance overview

Maintaining NEs using the NFM-P

22.1 Overview

22.1.1 NE maintenance

The NFM-P has policies and functions that facilitate the maintenance of managed NEs. You can use the NFM-P to do the following:

- configure how and when to deploy NE configuration changes—see Chapter 25, “NE deployment”
- back up device configurations on demand, or using a schedule—see Chapter 23, “NE backup and restore”
- restore device configurations—see Chapter 23, “NE backup and restore”
- perform device configuration rollbacks—see Chapter 24, “NE configuration rollback”
- upgrade device software on demand, or using a schedule—see Chapter 24, “NE configuration rollback”
- monitor and troubleshoot operations in progress
- browse NE file systems—see 22.5 “To view an NE file system using an FTP file browser” (p. 739) and 22.2 “To view an NE file system using an SSH file browser” (p. 736)

You can also view the accounting statistics collection status of an NE using 22.3 “To view the accounting statistics collection status of an NE” (p. 737) , and view the SNMP trap metrics of an NE using 22.6 “To view NE trap metrics information” (p. 740) .

22.1.2 NE file system browsing

An NFM-P operator can browse the file system of a managed NE to list the contents of the compact flash devices. You can browse the file system of a 7210 SAS, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, or OmniSwitch using simple FTP or a CLI session using SSH. The NFM-P GUI is used to browse the different types of files.

Note: Some 7210 SAS devices are equipped with a USB flash memory (uf1). In the equipment tree and on Alarm Info forms, the NFM-P displays only Flash Memory. The display does not distinguish between compact flash memory (cf) and USB flash memory (uf).

Browsing an NE file system using the NFM-P is a convenient way to confirm that operations such as the following occur as planned by verifying the sizes and time stamps of local NE files:

- NE configuration saves
- NE software image transfers and upgrades
- NE configuration restores
- NE accounting-statistics collection

FTP file browsing on an NE requires FTP user account access on the NE. SSH file browsing requires console user-account access and the configuration of SSH security on the NE. See

Chapter 8, “Device commissioning and management” for information about enabling FTP and console access for an NE user and configuring SSH on an NE.

Note: A 7705 SAR may become temporarily unreachable while enabling SSH and starting the device SSH server.

See 22.5 “To view an NE file system using an FTP file browser” (p. 739) for information about browsing an NE file system using an FTP file browser, and 22.2 “To view an NE file system using an SSH file browser” (p. 735) for information about browsing an NE file system using an SSH file browser.

22.2 To view an NE file system using an SSH file browser

22.2.1 Prerequisites

Perform this procedure to browse and list the contents of a managed NE using a secure file browser. You need console and SSH user-account privileges on the NE for access to the NE file system, and an SSH server must be configured on the NE. See Chapter 8, “Device commissioning and management” for information about enabling console or SSH access for an NE user and configuring the SSH server on an NE.

22.2.2 Steps

1

Initiate an SSH file browser session using one of the following methods:

a. Use the NFM-P main menu.
   1. Choose Tools→Network Elements→NE Sessions→SSH File Browser. The SSH File Browser form opens.
   2. Enter the IP address of the NE that you need to browse in the field at the top of the form.
   3. Click Connect. The Enter the Username and Password form opens.

b. Use the contextual menu for an NE.
   1. Select an NE icon in the NFM-P network navigation tree or topology map.
   2. Right-click on the NE icon and choose NE Sessions→File Browser. The SSH File Browser form opens, then displays the Enter the Username and Password form.

c. Use the NE alarm contextual menu.
   1. Select an NE alarm in the NFM-P alarm window.
   2. Right-click on the alarm item and choose NE Sessions→File Browser. The SSH File Browser form opens, then displays the Enter the Username and Password form.

Note: When you use the NFM-P main menu to open a file-browser session, you are not restricted to the original NE; you can use the same form to connect to other NEs. This is useful when you need to browse several NEs in succession.

2

Enter the user name and password of a user account with FTP and SSH access privileges on the NE and click OK. If the NE accepts the credentials, the form lists the contents of the NE.

<DIR> in the Type column indicates a directory. The file path to the current directory is displayed in the Path field.

On an NE with redundant CPMs, the form lists the contents of the cf3 device on the active CPM. You can browse the cf3 device on the standby CPM by specifying cf3-B:\ in the Path field.

3

Navigate the file system as required. Perform one of the following actions to open a directory and list the contents:

a. Double-click on the directory row in the list.

b. Select the directory row and press Ctrl-O.

c. Type the path to the directory in the Path field and click Go.

4

If you opened the browser using the NFM-P main menu or from the Software Upgrade form, you can browse another NE file system using the same form, if required.

1. Click Disconnect to end the browsing session.
2. Enter the IP address of the NE that you need to browse in the field at the top of the form and click Connect.

5

Close the form.

END OF STEPS

22.3 To view the accounting statistics collection status of an NE

22.3.1 Steps

1

Choose Tools→Statistics→Accounting Retrieval Status from the NFM-P menu. The Accounting Retrieval Status form opens with a list of managed NEs displayed.

2

Select an NE in the list and click Properties.

3

View the statistics collection information for the NE.

4

Close the forms.

END OF STEPS

22.4 To configure an event log policy

22.4.1 Purpose

Perform this procedure to specify the amount of time that log records are kept in the NFM-P database and enable or disable the saving of log records in the database.

Event log policies control the retention time and administrative state of each log type across the managed network.

Note: You must have a user account with access to the “event” permission to perform this procedure. This permission is included by default in the “NFM-P Management and Operations” scope of command role.

Note: You can also edit an event log policy by clicking Edit Policy in the Events tab after selecting a log type.

22.4.2 Steps

1

Choose Tools→Events→Event Policies from the NFM-P main menu. The Manage Event Policies form opens.

2

Select an entry and click Properties. The Event Policy (Edit) form opens.

3

Configure the required parameters.

Note: The Administrative State parameter enables or disables saving the event type in the NFM-P database.

4

Click OK. The changes to the policy are saved.

END OF STEPS

22.5 To view an NE file system using an FTP file browser

22.5.1 Prerequisites

Perform this procedure to browse and list the files on a managed NE. You need FTP user-account privileges on the NE for access to the NE file system. See Chapter 8, “Device commissioning and management” for information about enabling FTP access for an NE user.

22.5.2 Steps

1

Initiate an FTP file browser session using one of the following methods:

a. Use the NFM-P main menu.

1. Choose Tools→Network Elements→NE Sessions→FTP File Browser. The FTP File Browser form opens.
2. Enter the IP address of the NE that you need to browse in the field at the top of the form.
3. Click Connect. The Enter the Username and Password form opens.

b. Use the contextual menu for an NE.

1. Select an NE icon in the NFM-P network navigation tree or topology map.
2. Right-click on the NE icon and choose NE Sessions→File Browser. The FTP File Browser form opens, then displays the Enter the Username and Password form.

c. Use the NE alarm contextual menu.

1. Select an NE alarm in the NFM-P alarm window.
2. Right-click on the alarm item and choose NE Sessions→File Browser. The FTP File Browser form opens, then displays the Enter the Username and Password form.

Note: When you use the NFM-P main menu to open a file-browser session, you are not restricted to the original NE; you can use the same form to connect to other NEs. This is useful when you need to browse several NEs in succession.

2

Enter the user name and password of a user account with FTP access privileges on the NE and click OK. If the NE accepts the credentials, the form lists the contents of the NE.

<DIR> in the Type column indicates a directory. The file path to the current directory is displayed in the Path field.

On an NE with redundant CPMs, the form lists the contents of the cf3 device on the active CPM. You can browse the cf3 device on the standby CPM by specifying cf3-B:\ in the Path field.

3

Navigate the file system as required. Perform one of the following actions to open a directory and list the contents.

a. Double-click on the directory row in the list.

b. Select the directory row and press Ctrl-O.

c. Type the path to the directory in the Path field and click Go.

4

If you opened the browser using the NFM-P main menu, you can browse another NE file system using the same form, if required.

1. Click Disconnect to end the browsing session.
2. Enter the IP address of the NE that you need to browse in the field at the top of the form and click Connect.

5

Close the form.

END OF STEPS

22.6 To view NE trap metrics information

22.6.1 Steps

1

Choose Tools→Statistics→Trap Metrics Information from the NFM-P main menu. The Trap Metrics Information form opens.

The form lists the NEs that send the most traps during the most recent collection interval. The collection interval is indicated by the Start Collection Period and End Collection Period values.

2

View the trap metrics information for one or more NEs.

3

Click Search to refresh the information, as required.

4

Close the form.

END OF STEPS

23 NE backup and restore

NE backup and restore overview

23.1 NE backups and restores

23.1.1 Backup policies

An NFM-P backup policy specifies the conditions under which the NFM-P performs device configuration backups. The NFM-P stores the backups, which can quickly be restored in the event of an NE failure. Each managed NE that is not associated with a policy is assigned to a default policy, but you can create multiple additional policies and assign NEs to the policies according to your requirements.

A backup policy specifies settings that include the following:

- the backup schedule
- the files that a backup is to retrieve
- the type of file compression to use, if any
- the age and number of backup files that the NFM-P retains

An NFM-P operator that has the network element software management scope of command role can perform device configuration save, backup, and restore operations.

Note: An NFM-P user that has the lawful intercept management scope of command role can back up and restore an LI configuration. Backup data is saved only when the LI Local Save Allowed parameter on the Polling tab of the NE properties form is selected.

Note: An on-demand backup takes priority over scheduled backup operations, and is processed as soon as the NFM-P has the required resources.

You can export backup files from the NFM-P to a GUI client file system, and can import NE backup file sets to the NFM-P. In a redundant NFM-P system, the backup files are synchronized between the primary and standby main servers.

You can configure the NFM-P to automatically delete device backup files after the associated NE is unmanaged; see the chapter on NFM-P database management in the *NSP System Administrator Guide* for more information.

When a device configuration requires replacement, for example, when it is corrupt, you can restore a previous configuration. Unless otherwise specified, the NFM-P restores the most recent configuration backup. See 23.5 “To perform an on-demand backup, restore, or configuration save” (p. 748) and 23.8 “To restore a device configuration other than the most recent” (p. 752) for information about restoring device backups.

A restore works implicitly with a backup policy; there is no explicit restore policy. An operator must manually restore a device backup.

Before you can delete a backup policy, you must remove from the policy all NEs that are assigned to the policy.

GNE backups

The NFM-P NE backup function is supported for GNEs, depending on the device type and release. Backing up a GNE configuration requires a valid NFM-P GNE driver that supports the function.

Note: Some GNE devices may have specific backup requirements. Before you attempt to perform a GNE backup operation, see the appropriate domain user guide for information about driver capabilities and driver-specific information, including backup requirements.

GNE backups are configured and performed using the policy-based mechanism that is used for natively managed devices; the type of backup policy required for GNEs is Generic NE Node. In a Generic NE Node backup policy, you must specify the transfer protocol and user credentials, and a directory on the main server that is used for transferring the backup files from the GNEs.

Note: GNE backups have the following special requirements.

- FTP, SFTP, or TFTP must be enabled on each main server, depending on which protocol you use.
- If a firewall is used, the required FTP, SFTP, or TFTP ports must be open.
- The FTP, SFTP,or TFTP user that you specify in a backup policy must be a valid FTP, SFTP, or TFTP user on each main server station, and must belong to the sam user group on each main server station.
- The nsp user requires read, write, and execute access to the directory that you specify in the policy.

23.2 Backup policy configuration example

23.2.1 Example details

The following example describes an NFM-P backup policy configuration using appropriate parameter values for most scenarios. The Backup Policy form contains the parameters in the example; see 23.4 “To configure a backup policy” (p. 746) for the configuration information.

The sample backup policy specifies that the NFM-P obtains the backup files by FTP from the NE once every hour regardless of configuration activity, and after every 25 configuration changes. The policy also specifies that the NFM-P backs up the device configuration file, license file, and boot options file, or BOF, only when a newer file is present, and uses gzip file compression. The NFM-P is to retain at most 30 backup versions, and purge versions that are more than 30 days old.

The following figure illustrates the activities of the NFM-P backup and restore process.

Figure 23-1 NFM-P backup process

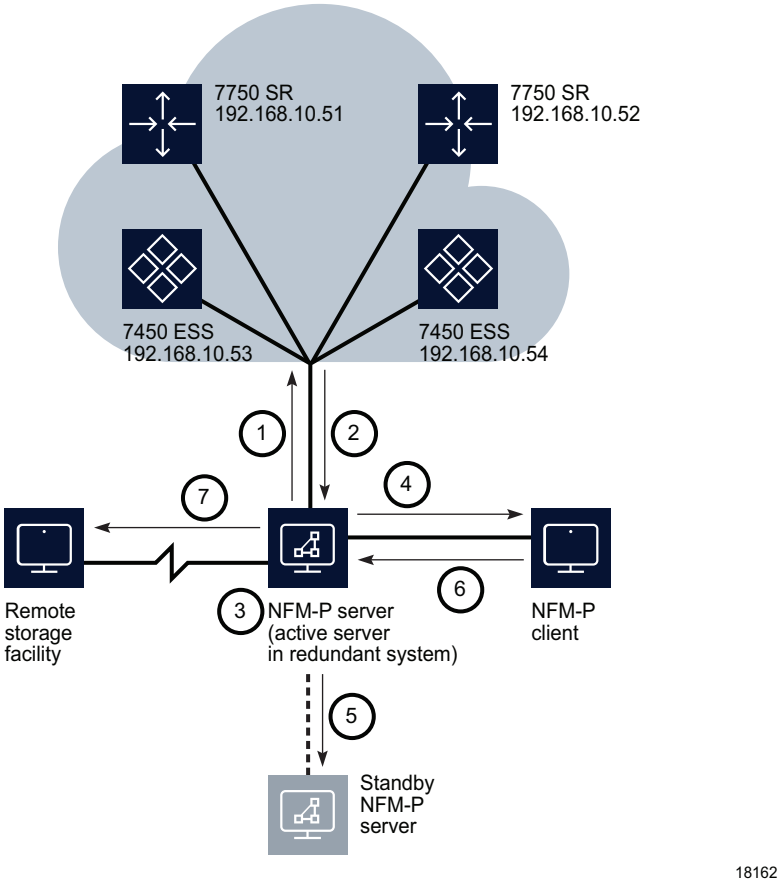

The labels in Figure 23-1, “NFM-P backup process” (p. 743) correspond to events in the following sequence:

1. At the interval specified, which is every hour in the example, the NFM-P issues a request to all NEs for a backup.
2. The NEs send the BOF, configuration files, and license file to the NFM-P.
3. The NFM-P stores the new versions of the received files.
4. A GUI client operator views the backup status on the Backup/Restore Status tab of the Backup Policy form.
5. If the NFM-P system is redundant, the active and standby main servers synchronize the backup files.
6. An NFM-P operator uses the Backup/Restore form to perform on-demand and scheduled backups, restores, and configuration saves, as required.

7. A third-party application periodically sends a copy of the backup files from the NFM-P to a remote storage location for safekeeping.

The policy parameters and values are the following:

- Enable Backup—enabled
- Auto Reboot After Successful Restore—enabled
- Scheduled Backup Scheme—Every Scheduled Interval
- Scheduled Backup Interval—1 hour
- Scheduled Backup Sync Time — 00:00
- Scheduled Backup Threshold (operations)—5
- Auto Backup Scheme—Every Nth NFM-P Server Initiated Save
- Auto Backup Threshold (operations)—50
- CLI Config File Mode—New Version Only
- CLI Config Save Details—disabled
- CLI Debug Save Config File Mode—disabled
- Boot Option File Mode—New Version Only
- File Compression—GZIP
- Backup 7705 Radio Database—disabled
- Auto-Purge Scheme—By Age But Retain A Minimum Number Of Backups
- Number of Backups—30
- Maximum Backup Age (days)—30

23.3 Workflow to perform NE backups and restores

23.3.1 Stages

1

For secure backups and restores, verify that SSH2 is configured and enabled on each NE, and that the NFM-P mediation policy for the NEs is configured to use secure FTP.

2

Configure device backup policies; see 23.4 “To configure a backup policy” (p. 746) .

3

Perform an on-demand backup, restore, or configuration save; see 23.5 “To perform an on-demand backup, restore, or configuration save” (p. 748) .

4

Perform an on-demand OmniSwitch backup or configuration save; see 23.6 “To perform an on-demand OmniSwitch backup or configuration save” (p. 749) .

5

Restore an OmniSwitch device backup; see 23.7 “To restore an OmniSwitch configuration” (p. 751) .

6

Restore a device backup other than the most recent; see 23.8 “To restore a device configuration other than the most recent” (p. 752) .

7

View the backup, restore, or configuration save status; see 23.9 “To view the backup, restore, or configuration save status of an NE” (p. 753) .

8

Import a device backup file set to the NFM-P; see 23.11 “To import a device configuration backup” (p. 755) .

9

Export a device backup file set from the NFM-P; see 23.10 “To export a device configuration backup” (p. 754) .

NE backup and restore procedures

23.4 To configure a backup policy

23.4.1 Steps

Note: All NFM-P-managed NEs that are not assigned to a backup policy are automatically assigned to a default backup policy upon device discovery.

Note: Beginning in NSP 23.11, the backup policy with settings “CLI Config File Mode” and “Boot Option File Mode” set to “New Version only” will only allow the second backup or consecutive backup operation to be successful if there is a change in the config or bof file. Optional files like rollback, tech-support and certificate backup will not happen even if selected in the policy unless there is a “New Version” in BOF or Config file. This behavior is the same for both manual backups and scheduled backup operations.

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Policy tab.

3

Click Create, or select an existing policy and click Properties. The Backup Policy (Create | Edit) form opens.

4

Configure the Enable Backup parameter. If you disable the Enable Backup parameter, go to Step 9 .

5

If you are creating a policy, configure the Policy Type parameter.

Note: For 7210 SAS, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, or Wavence NEs, choose the SR Based Node option.

Note: For a GNE, choose the Generic NE Node option.

6

If the Policy Type is set to SR Based Node, you can perform a reboot after the configuration is restored on the NE by specifying the Auto-Reboot After Successful Restore parameter.

Note: When you use an NFM-P GUI client to restore a configuration, and the Auto-Reboot After Successful Restore parameter is unselected, the BOF is overwritten by a subsequent “bof save” command on the NE.

Note: When a reboot does not immediately follow a restore, use a CLI to view the BOF and ensure that it has not changed since the restore.

Note: For Wavence NEs, you must enable the Auto-Reboot After Successful Restore parameter in order for the restored configuration to be activated on the NE. Otherwise, the restored configuration is not activated on the NE during a later reboot.

7

Configure the parameters in the Backup Triggering, Backup Settings, and Backup Purging panels. You can schedule backups based on a time interval, or based on the number of configuration saves initiated by the NFM-P.

Backup purge parameters allow you to specify the number of backup files that are retained. The purge settings eliminate the need for manual backup file deletion. The purge criteria can be the number of files, the age of the files, or both.

Note: The Backup roll back files and Backup tech-support files parameters appear and are enabled by default for SR Based Node type. Once Backup is successful on NFM-P for a node, these files as well as node bof, config files and other files, are saved on NFMP-server under NE Backup location (/opt/nsp/nfmp/nebackup).

Note: If the Policy Type parameter is set to SR Based Node, in order to enable the CLI Debug Save Config File Mode parameter, you must specify the location of the debug configuration files in the NFM-P configuration. See the chapter on NFM-P component configuration in the *NSP System Administrator Guide* for information about setting the parameter.

Note: For an OmniSwitch, the NFM-P can back up only files in the certified directory. If you need to back up configuration files in the working directory, you must ensure that the files in the certified and working directories are identical. See 26.23 “To certify or synchronize OmniSwitch software” (p. 812) for information about issuing a Certify and Synchro command to synchronize the files.
For supporting OmniSwitch NEs, you can configure the Restore Directory parameter to specify the directory on the node that the backup files are restored to. To specify a user-defined directory, enter a text string that matches the name of the user-defined directory on the node.

Note: If the Backup 7705 Radio Database parameter is enabled, the Wavence MPT radio databases are backed up with the configuration file, depending on the CLI Config File Mode setting. Only the 7705 SAR-8 and 7705 SAR-18 support the backup of the Wavence MPT radio databases.

8

If the Policy Type is Generic NE Node, configure the Transfer Protocol parameter, and then the FTP or SFTP user credentials in the Generic NE Backup Settings panel.

Note: You must ensure that the following conditions are met before you attempt to use the backup policy.

- FTP, SFTP, or TFTP is enabled on each main server, depending on which protocol you use.
- If a firewall is used, the required FTP, SFTP, or TFTP ports are open.
- The FTP, SFTP, or TFTP user that you specify is a valid FTP, SFTP, or TFTP user on each main server station, and belongs to the sam user group on each main server station.
- The nsp user has read, write, and execute access to the Root Directory that you specify.

9

Click Apply.

10

Click on the Backup/Restore Policy Assignment tab and assign one or more NEs to the policy, as required.

11

Close the form.

END OF STEPS

23.5 To perform an on-demand backup, restore, or configuration save

23.5.1 Prerequisites

Note: To perform an on-demand backup or configuration save on an OmniSwitch, see 23.6 “To perform an on-demand OmniSwitch backup or configuration save” (p. 749) . To restore an OmniSwitch configuration, see 23.7 “To restore an OmniSwitch configuration” (p. 751) .

When you perform an on-demand backup, the device configuration is backed up based on the backup policy to which the NE is assigned.

A restore operation uses the most recent backup file unless otherwise specified; see 23.8 “To restore a device configuration other than the most recent” (p. 752) for information about restoring a backup file other than the most recent.

The following must be true in order for you to perform this procedure:

- You have an NFM-P user account with an administrator or network element software management scope of command role, or a scope of command role with write access to the mediation package.
- FTP or secure FTP is configured in the NE mediation policy.
- BOF persistence is enabled on each NE that supports BOF persistence.

23.5.2 Steps

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Status tab. The managed NEs are listed.

3

Select an NE and click Backup, Restore, or Save Config, as required.

4

Click Yes. The operation begins, and the current operation state is indicated in the appropriate column.

5

You can resynchronize an NE with the NFM-P, if required, by clicking Resync.

6

Close the form.

END OF STEPS

23.6 To perform an on-demand OmniSwitch backup or configuration save

23.6.1 Steps

Note: The following limitations apply to an OmniSwitch virtual chassis node, depending on how the node is managed:

- If the NE is managed using the out-of-band management IP address, and is a single-shelf node, vcboot.cfg and vcsetup.cfg files are backed up from primary.
- If the NE is managed using out-of-band management IP address and is a multi-shelf node, the following are backed up:

–vcboot.cfg file from primary

–vcsetup.cfg file from the primary and all secondary NEs

If any of the NE is unreachable from the NFM-P server, the backup operation fails.

- If the NE is managed using a VLAN-based management IP address or an in-band management IP address, the vcboot.cfg file is backed up from primary.

1

On the Equipment tree, expand an OmniSwitch object, right-click on a shelf object, and choose Properties. The Shelf (Edit) form opens.

2

Click on the Software Control Module tab and set the Command to Apply parameter to Certified.

3

Click Resync.

Note: You must resynchronize the NE to ensure the /flash/working and /flash/certified directory contents are the same, even though the Certified Status displays Certified.

4

Save your changes and close the form.

5

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

6

Click on the Backup/Restore Policy tab. The backup policies are listed.

7

Select a policy and Click Properties. The Backup Policy (Edit) form opens.

8

Click on the Backup/Restore Status tab. The tab lists NEs of the type specified by the Policy Type parameter setting on the General tab.

9

Select an NE and click Backup or Save Config, depending on the operation that you need to perform.

10

The operation begins; the Last Operation Details column displays the operation status.

11

When the operation is complete, close the forms.

END OF STEPS

23.7 To restore an OmniSwitch configuration

Note: You can import the locally available backup file for the same NE to the NFM-P which can be restored. Ensure the backup policy remains the same as was in place during the backup procedure. This will prevent the NFM-P from generating an error while importing the file.

The restored configuration is typically stored in a working directory. Supporting OmniSwitch devices allow you to configure a user-defined directory for the restored files; see 23.4 “To configure a backup policy” (p. 746).

The NFM-P indicates when the file is successfully transferred to the appropriate NEs. However, the file will not become active until the device(s) have been reloaded from the specified working or user-defined directory.

23.7.1 Steps

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Status tab. The NEs for which a backup is available are displayed.

3

Select an NE and click Restore. The restore begins, and the Last Operation Details column displays the operation status.

4

On the Equipment tree, expand the OmniSwitch icon, right-click on an OmniSwitch shelf object, and choose Properties. The Shelf (Edit) form opens.

5

Click on the Software Control Module tab.

6

Set the Command to Apply parameter to Reload.

7

If required, configure the Delayed Activation Timer and Redundancy Time (seconds) parameters.

8

Configure the Image Files Directory parameter.

9

Save your changes and close the form.

10

After reboot, on the Shelf (Edit) form, click on the Software Control Module tab and verify that the Command to Apply parameter is set to Certified.

11

When the operation is complete, close the forms.

END OF STEPS

23.8 To restore a device configuration other than the most recent

23.8.1 Purpose

You can choose to restore an older version of the device configuration to meet special network requirements.

CAUTION

Service Disruption

Older backups may not have the most recent network information.

Restoring an older configuration may be service-affecting.

23.8.2 Steps

1

Back up the current device configuration as a safeguard; see 23.4 “To configure a backup policy” (p. 746) .

2

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P menu. The Backup/Restore form opens.

3

Click on the Backup/Restore Status tab. The managed NEs are listed.

4

Select an NE and click Properties. The Backup/Restore Status form for the NE opens.

5

Click on the Backups tab, which lists the available configuration backups for the NE, ordered from the oldest to the most recent.

6

Select a backup and click Restore. The NFM-P begins to restore the device configuration.

7

Click Resync to ensure that the latest network information is available to the NFM-P.

8

Close the form.

END OF STEPS

23.9 To view the backup, restore, or configuration save status of an NE

23.9.1 Steps

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Status tab. Depending on the operation type, the Backup State, Restore State, or Config Save State column displays the current operation status for each listed NE, which is one of the following:

- Not Attempted—The operation has not begun.
- Saving Config—The NE is saving the configuration.
- Transferring Files—A file transfer is in progress.
- Success—The operation is complete and successful.
- Failure—The operation is complete but unsuccessful.
- CPM Sync and Pending Reboot—The device configuration is restored and the NE is synchronizing the CPMs before a reboot.
- CPM Sync and Pending Reboot Standby—The NFM-P awaits the reboot of the standby CPM.
- Standby Reboot and Pending Redundant Switch-over—The NFM-P awaits the switchover to the standby CPM.
- MPR reboot pending—The NFM-P awaits the reboot of a Wavence NE.

Note: During a backup, if an NE is unresponsive to the NFM-P because SNMP on the NE is disabled, the Backup State column entry for the NE does not immediately display the

correct value of Failed. This latency is caused by the inability of the NFM-P to communicate with the unresponsive NE. In such a situation, the Backup State column displays the Saving Config status until three 10-minute SNMP polling periods, or 30 minutes, have elapsed, after which the Backup State changes to Failed if SNMP remains disabled.

3

To view more information about the operations performed on an NE, select the NE and click Properties. The NE Backup/Restore Status form opens. The Last Operation Details indicator describes the most recent operation type.

4

Click on the Backups tab to display a list of the recent device backups, if required.

5

Click on the Configuration Saves tab to view information about recent configuration saves, if required.

6

Close the forms.

END OF STEPS

23.10 To export a device configuration backup

23.10.1 Steps

Perform this procedure to export a backup file set from the NFM-P to a GUI client file system.

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Status tab.

3

Select the NE for which you are exporting a backup and click Properties. The NE Backup/Restore Status form opens.

4

Click on the Backups tab. A list of backups for the NE is displayed.

5

Select a backup and click Export. A file browser form opens.

6

Specify the local directory that is to contain the exported file set and click OK. The backup file set is saved in the directory.

7

Close the forms.

END OF STEPS

23.11 To import a device configuration backup

23.11.1 Steps

Perform this procedure to import a backup file set from a GUI client file system to the NFM-P.

1

Choose Administration→NE Maintenance→Backup/Restore from the NFM-P main menu. The Backup/Restore form opens.

2

Click on the Backup/Restore Status tab.

3

Select the NE for which you are importing a backup and click Properties. The NE Backup/Restore Status form opens.

4

Click Import. A file browser form opens.

5

Specify the directory that contains the backup file set and click OK. The NFM-P imports and stores the backup file set.

6

Close the forms.

END OF STEPS

24 NE configuration rollback

NE configuration rollback overview

24.1 Comparing configuration files

24.1.1 Using the compare function

A compare function allows you to display NE configuration changes between selected checkpoints, a selected checkpoint and the active configuration, a selected checkpoint and the rescue file, or the active configuration and the rescue file. This is useful for troubleshooting NE configurations.

Note: In order to use the compare function, valid SSH or Telnet credentials must be specified in the applicable mediation policy.

When you choose one of the compare NE configuration rollback file options, the Checkpoint Compare window opens listing differences between the selected configuration files. The window title displays the names of the NE configuration files that are compared. For example, the title “Index 0 and rescue” indicates that the first NE configuration checkpoint file listed on the Checkpoint Files tab is compared with the NE configuration rescue file.

Differences in the configuration files are noted as follows:

- A plus sign (+) indicates that the configuration is in the first NE configuration file listed in the window title but not in the second NE configuration file. In this case, the configuration is in the checkpoint file index 0.
- A minus sign (–) indicates that the configuration is in the second NE configuration listed file in the window title but not in the first NE configuration file. In this case, the configuration is in the rescue file.
- Configurations that are in both NE configuration files are not listed in the Checkpoint Compare window.

24.2 Workflow to configure NE configuration rollback

24.2.1 Stages

1

Configure the storage location for the NE configuration rollback rescue file and checkpoint files; see 24.3 “To configure NE configuration rollback file storage” (p. 759) .

2

Create the NE configuration rollback rescue file; see 24.4 “To create an NE configuration rollback rescue file” (p. 760) .

3

As required, perform one or both of the following:

a. Create the NE configuration rollback checkpoint files manually; see 24.5 “To create NE configuration rollback checkpoint files” (p. 760) .

b. Configure scheduled checkpoint file creation; see 24.6 “To configure scheduled checkpoint file creation” (p. 762) .

4

As required, compare NE configuration rollback files with each other, or with the current NE configuration file; see 24.7 “To compare NE configuration rollback files” (p. 763) .

5

As required, view the configuration files to verify the content; see 24.9 “To view NE configuration files” (p. 764) .

6

As required, perform one or both of the following:

a. Revert to a previous NE configuration manually; see 24.8 “To revert to a previous NE configuration” (p. 764) .

b. Configure scheduled reversion to an NE configuration rollback checkpoint file; see 24.6 “To configure scheduled checkpoint file creation” (p. 762) .

NE configuration rollback procedures

24.3 To configure NE configuration rollback file storage

24.3.1 Purpose

Perform this procedure to configure the storage location for the NE configuration rollback rescue file or checkpoint files. To configure NE configuration rollback file storage locations, FTP must be configured on the device.

24.3.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Click on the NE Rollback tab. The NE Rollback list form opens.

3

Choose a device and click Properties. The NE Rollback (Edit) form opens.

4

To configure the storage location for the NE configuration rollback rescue file:

1. In the Rescue File panel, configure the Rescue File Location Type parameter.
   - If you specified FTP, in the Rescue File — Remote Location panel, configure the required parameters.
   - If you specified CFLASH, in the Rescue File — CFLASH panel, configure the required parameters.
2. Save your changes.

5

To configure the storage location for the NE configuration rollback checkpoint files:

1. In the Checkpoint File Storage panel, configure the required parameters.
2. Configure the Rollback Location Type parameter.
   - If you specified FTP, in the Checkpoint File Storage — Remote Location panel, configure the required parameters.
   - If you specified CFLASH, in the Checkpoint File Storage — CFLASH panel, configure the required parameters.
3. Save your changes. The Checkpoint Files tab lists the checkpoint files.

6

Close the forms.

END OF STEPS

24.4 To create an NE configuration rollback rescue file

24.4.1 Prerequisite

The rollback rescue file storage location must be configured before you can create the NE configuration rollback rescue file; see 24.3 “To configure NE configuration rollback file storage” (p. 759) .

24.4.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Click on the NE Rollback tab. The NE Rollback list form opens.

3

Choose a device from the list and click Properties. The NE Rollback (Edit) form opens.

4

Click Create Rescue File and confirm the action. The rescue file is saved to the specified file location. The Delete Rescue File, View Rescue File, and Revert to Rescue File buttons are enabled.

Note: If a rescue file exists, the Rescue File Exists parameter is enabled, and the dialog box prompts you to confirm that you need to overwrite the existing rescue file.

5

Close the forms.

END OF STEPS

24.5 To create NE configuration rollback checkpoint files

24.5.1 Prerequisite

The checkpoint file storage location must be configured before you can create NE configuration rollback checkpoint files; see 24.3 “To configure NE configuration rollback file storage” (p. 759) .

You can also automate NE configuration rollback checkpoint file creation; see 24.6 “To configure scheduled checkpoint file creation” (p. 762) .

24.5.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

To create a checkpoint file:

1. Click on the NE Rollback tab.
2. Choose a device and click Properties. The NE Rollback (Edit) form opens.
3. Click on the Checkpoint Files tab.
4. Click Create and confirm. The Create Checkpoint form opens.
5. If required, add a comment to the checkpoint file. Click OK. The comment dialog box closes.
6. Click OK. The Create Checkpoint form opens and the checkpoint is added to the list form on Checkpoint Files tab.

3

Note: You can enable redundant rollback synchronization only when the following conditions are met.

- The rollback file location is CFLASH.
- The device has a standby CPM.

To enable redundant rollback synchronization between the active and standby CPMs:

1. Click on the General tab.
2. Perform one of the following:
   - To configure automated redundant rollback synchronization, select the Redundant Rollback Synchronization check box.
   - To perform manual redundant rollback synchronization, click Rollback-Sync Now.

4

Save your changes and close the form.

END OF STEPS

24.6 To configure scheduled checkpoint file creation

24.6.1 Additional resources

See Chapter 5, “NFM-P-based schedules” for information about creating schedules. You can also create NE configuration rollback checkpoint files manually; see 24.5 “To create NE configuration rollback checkpoint files” (p. 760) .

24.6.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Choose a rollback policy and click Properties. The Rollback Policy (Edit) form opens.

3

In the Checkpoint File Creation panel, click Create. The Checkpoint Create Scheduled Task (Create) form opens.

4

Configure the required parameters.

5

Click Select beside the ID parameter. The Select Schedule - Checkpoint Create Scheduled Task form opens.

6

Perform one of the following:

a. Click Create to create a schedule. The NFM-P Schedule (Create) form opens.

   1. Configure the required parameters.
   2. Click OK. The schedule is saved.
   3. Choose the schedule in the Select Schedule list and click OK.

b. Choose an existing schedule and click OK.

7

Click OK in the Checkpoint Create Scheduled Task (Create) form. The policy is updated with the selected scheduled task.

END OF STEPS

24.7 To compare NE configuration rollback files

24.7.1 Purpose

Perform this procedure to compare NE configuration rollback files with each other or with the current NE configuration file.

24.7.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Click on the NE Rollback tab.

3

Choose a device and click Properties. The NE Rollback (Edit) form opens.

4

Click on the Checkpoint Files tab.

5

Perform one of the following:

a. To compare the NE configurations of two checkpoint files:
   1. Choose two checkpoint files.
   2. Click Compare File and choose Checkpoint Vs Checkpoint.

b. To compare the NE configurations of a checkpoint file and the active configuration:
   1. Choose a checkpoint file.
   2. Click Compare File and choose Checkpoint Vs Active-Configuration.

c. To compare the NE configurations of a checkpoint file and the rescue file:
   1. Choose a checkpoint file.
   2. Click Compare File and choose Checkpoint Vs Rescue.

d. To compare the NE configuration of the current configuration and the rescue file, click Compare File and choose Active-Configuration Vs Rescue.

The Checkpoint Compare window opens listing differences between the selected configuration files; see 24.1 “Comparing configuration files” (p. 757) for information.

END OF STEPS

24.8 To revert to a previous NE configuration

24.8.1 Purpose

Use this procedure to revert to a previous device configuration without rebooting an NE. You can automate NE configuration reversion for NE configuration rollback checkpoint files; see 24.6 “To configure scheduled checkpoint file creation” (p. 762) .

24.8.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Click on the NE Rollback tab.

3

Choose a device and click Properties. The NE Rollback (Edit) form opens.

4

Perform one of the following:

a. To revert to the NE configuration rollback rescue file:
   1. Click Revert to Rescue File. A dialog box appears.
   2. Click OK to acknowledge the warning and perform the reversion.

b. To revert to a NE configuration rollback checkpoint file:
   1. Click on the Checkpoint Files tab.
   2. Choose a checkpoint file and click Revert Checkpoint File. A dialog box appears.
   3. Click OK to acknowledge the warning and perform the reversion.

END OF STEPS

24.9 To view NE configuration files

24.9.1 Purpose

Use this procedure to view NE configuration rollback rescue files and checkpoint files, or the current NE configuration file.

24.9.2 Steps

1

Choose Administration→NE Maintenance→NE Configuration Rollback from the NFM-P main menu. The NE Configuration Rollback form opens.

2

Click on the NE Rollback tab.

3

Choose a device and click Properties. The NE Rollback (Edit) form opens.

4

Perform one of the following:

a. To view the NE configuration rollback rescue file:
   1. Click View Rescue File. The Rescue View window opens and the NFM-P loads the file data.
   2. View the file contents and close the form.

b. To view the NE configuration rollback checkpoint files or the active configuration file:
   1. Click on the Checkpoint Files tab.
   2. Click View File and choose an option. The View window for the chosen option opens.
   3. View the file contents and close the form.

5

Close the NE Configuration Rollback form.

END OF STEPS

25 NE deployment

Using the NFM-P to deploy NEs

25.1 NE deployment overview

25.1.1 Overview

When you use the NFM-P to apply an NE configuration change, for example, by clicking OK or Apply after setting a service parameter, the NFM-P deploys the change to the NE according to the NFM-P deployment policy. The NFM-P deployment policy also specifies the conditions under which each managed NE performs a configuration save.

The information in a deployment policy includes the following:

- number and frequency of NFM-P deployment retries
- NE configuration save settings, such as the following:
  - save frequency
  - level of configuration detail to save
  - delay between consecutive saves

In a lab or testing environment, it is sometimes necessary to disable NFM-P deployment. See the *NSP NFM-P XML API Developer Guide*for information about disabling NFM-P deployment.

Note: The Deployment tab on an object properties form lists the failed deployments and deployments that are in progress. In addition, a deployment icon is displayed on the configuration form beside the parameter associated with the failed or attempted deployment, beside the NE object in the navigation tree, in list tables, and in map info tables.

25.2 Sample deployment policy configuration

25.2.1 Example details and settings

The following example describes the configuration and operation of an NFM-P deployment policy. See 25.4 “To configure the NFM-P deployment policy” (p. 768) for deployment policy configuration information.

The example policy settings are as follows:

- Auto Save Scheme—Every Nth Successful Deployment
  Initiate a device configuration save after the number of configuration changes specified by Auto Save Threshold.
- Auto Save Threshold—3
  Initiate a configuration save after every three successful deployments, if Auto Save Scheme is set to Every Nth Successful Deployment.
- Scheduled Save Scheme—None
  Do not perform a scheduled configuration save on any NE.

- Scheduled Save Interval—1 hour
  If scheduled saves are enabled, perform a save every hour.
- Save Details—disabled
  Save only non-default parameter values.
- Configuration Save Interval—30 seconds
  Delay consecutive configuration save requests for an NE that are less than 30 seconds apart.
- Interval Repeat Limit—5
  Wait up to five Configuration Save Interval periods before performing a configuration save.
- Retry Scheme—Retry Number Of Times
  Retry each deployment the number of times specified by Retry Interval.
- Retry Interval—5 minutes
  Wait five minutes before retrying a failed deployment.
- Retry Threshold—3
  Retry each failed deployment three times.

25.3 Workflow to configure and manage NE deployment

25.3.1 Stages

1

Configure the NFM-P deployment policy to specify how and when the NFM-P deploys configuration changes to the managed NEs; see 25.4 “To configure the NFM-P deployment policy” (p. 768) .

2

As required, view the deployment status and troubleshoot failed configuration deployments; see 25.5 “To view and manage failed deployments” (p. 769) .

25.4 To configure the NFM-P deployment policy

25.4.1 Steps

1

Choose Administration→NE Maintenance→Deployment from the NFM-P main menu. The Deployment form opens.

2

Click on the Deployment Policy tab.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

25.5 To view and manage failed deployments

25.5.1 Failed deployment scenarios

The NFM-P continues to retry a deployment after a failed or incomplete attempt, based on the NFM-P deployment policy.

When a deployment error occurs, a number of problems can result. For example:

- The NFM-P database may lose synchronization with the device database.
- Configuration changes requested using the client GUI may conflict with configuration changes, retries, and recovery applications developed by an OSS system, or by an operator using a CLI.

When a deployment fails, a Problems Encountered form displays error information about the failure.

Note: The Request Id and Task Name values on the form can be used as troubleshooting references in the NFM-P Task Manager.

The Problems Encountered form is also displayed for non-deployment errors.

25.5.2 Steps

1

If a Problems Encountered form opens and you want to navigate directly to the object of the deployment failure, click View Affected Object. The object properties form opens.

The current failed deployments are listed on the Deployment tab, from which you can select an entry and click Properties to view the deployment information.

2

Perform one of the following.

a. Open the list of failed deployments from a Problems Encountered form; click Deployment.

b. Open the list of failed deployments by choosing Administration→NE Maintenance→Deployment from the NFM-P main menu.

The Deployment form opens.

3

Review the deployment information. The State value indicates the cause of the deployment failure.

4

Select a deployment and click Properties. The properties form for the deployment opens. The objects of the deployment failure are shown in the Objects list.

5

Select an Objects list entry and click Properties. The Object Change form that describes the attempted configuration change opens.

6

Select an entry in the Attributes list and click Properties. The Attribute Change form opens and displays the following object attribute information:

- NE attribute to be modified
- previous attribute value
- new attribute value that the deployment fails to assign

7

Close the forms.

8

Perform one or more of the following, depending on the result of the investigation into the deployment failure.

a. Click Suspend Retries to override the deployment policy setting and prevent further deployment retries.

b. Click Resume Retries to override a previous Suspend Retries action performed on the deployment.

c. Click Clear to clear the deployment.

Note: Clearing a failed deployment may result in a loss of data synchronization between the NFM-P database and the NE. Nokia recommends that you resynchronize the NE objects associated with a failed deployment after you clear the deployment.

d. Click Force Submit to force the NFM-P to immediately reattempt the deployment.

9

Click Refresh to update the failed deployment list.

10

Close the form.

END OF STEPS

26 NE software upgrades

NE software upgrade overview

26.1 Software upgrades

26.1.1 Software upgrade requirements

CAUTION

Service Degradation

Ensure that you regularly remove from the NFM-P the device software images that are no longer required, for example, by deleting the images.

An accumulation of device software images can dramatically increase the length of an operation such as an NFM-P database backup, restore, or reinstantiation.

A software upgrade requires a software upgrade policy that specifies the device family, software image and image backup locations, and the actions to perform; for example, image download, activation, or ISSU. Depending on the device family, the activation function supports the reboot and reboot upgrade options. See 26.3 “Reboot and reboot upgrade” (p. 773) in this section for more information.

Using software upgrade policies

Using a software upgrade policy, an NFM-P operator can independently perform the image download, upgrade, and activation tasks. For example, you can configure a policy to perform the time-consuming software image downloads only, and then schedule the image activations as a separate task. When the images are downloaded in advance, the NFM-P can perform more activations in a maintenance period.

Nokia recommends that you plan the image download tasks independent of upgrade and activation tasks. The download tasks can be time-consuming depending on the software image size, network bandwidth, and number of NEs the image is being downloaded to.

Note: The NFM-P limits the software image download to 10 NEs at a time.

Software compatibility and upgrades

During a software upgrade, the NFM-P verifies that the new software is compatible with the device and that the required files are present; otherwise, the upgrade is not attempted. You can use the NFM-P to roll back a software upgrade in the event of an upgrade failure.

You cannot upgrade an NE to a chronologically older release, regardless of the release identifier; for example, the R1 revision of a release may predate the R9 revision of the previous release. See the *NSP NFM-P Network Element Compatibility Guide* for information about the supported device upgrade paths.

VSR-a software (SROS) upgrade is not supported in NFM-P.

Nokia recommends that you import all of the images that are bundled for a release. All 7x50 NEs use the same bundle. When downloading the software, the NFM-P automatically chooses the required files for a specific NE variant.

For SR NEs, the NFM-P downloads a complete list of software bundles regardless of chassis type.

For IXR NEs, the NFM-P downloads only the required list of image files based on the chassis type.

GNE software upgrades

The NFM-P NE software upgrade function is supported for GNEs, depending on the device type and release. Upgrading a GNE requires a valid NFM-P GNE driver that supports the function.

Note: Some GNE devices may have specific upgrade requirements. Before you attempt to perform a GNE upgrade, see the appropriate domain user guide for information about driver capabilities and driver-specific information.

GNE software upgrades are configured and performed using the policy-based mechanism that is used for natively managed devices; the type of upgrade policy required for GNEs is Generic NE Node. In a Generic NE Node upgrade policy, you must specify FTP or SFTP user credentials and a directory on the main server that is used for transferring the software image files to and from the GNEs.

Note: GNE software upgrades have the following special requirements.

- FTP or SFTP must be enabled on each main server, depending on which protocol you use.
- If a firewall is used, the required FTP or SFTP ports must be open.
- The FTP or SFTP user that you specify in a software upgrade policy must be a valid FTP or SFTP user on each main server station, and must belong to the sam user group on each main server station.
- The nsp user requires read, write, and execute access to the directory that you specify in the policy.

26.2 ISSUs

26.2.1 Upgrading in-service NEs

CAUTION

Service Disruption

Before you attempt an ISSU, see the appropriate device release notice and the *NSP NFM-P Network Element Compatibility Guide* for information about the following:

- device releases that support ISSU
- supported upgrade paths

The software upgrade information in the device documentation takes precedence over the information in this chapter.

You can use the NFM-P to perform an in-service software upgrade, or ISSU, on a managed NE that has dual CPMs. An ISSU allows an NE to provide uninterrupted service during the upgrade process.

A device software upgrade requires a CPM restart, which causes a temporary NE outage. When an NE has dual CPMs, however, one CPM can remain active while the other restarts using the upgraded software. The alternate CPM restarts mean that the NE remains fully in service during an upgrade. If the upgrade of a CPM fails, the CPM reports the failure and an alarm is raised.

Note: ISSUs are restricted to device maintenance releases only. See the device documentation for the supported release transitions.

Some 7210 SAS chassis types do not support ISSUs; see the device documentation for more information.

26.3 Reboot and reboot upgrade

26.3.1 Manual or automatic reboots

You can perform a manual or an automatic NE reboot during an on-demand software upgrade, and can configure an automatic reboot in a software upgrade policy. See 26.5 “To configure a software upgrade policy” (p. 776) for information about configuring an automatic reboot in a software upgrade policy, or 26.9 “To perform an ISSU or on-demand software upgrade” (p. 783) for information about performing a manual reboot during an on-demand software upgrade.

The NFM-P supports the 7210 SAS and 7705 SAR reboot upgrade option, which upgrades the system firmware. You can manually perform a reboot upgrade on an NE, or configure an automatic reboot upgrade in a software upgrade policy. See 26.5 “To configure a software upgrade policy” (p. 776) for information about configuring an automatic reboot upgrade, and 26.9 “To perform an ISSU or on-demand software upgrade” (p. 783) for information about performing a manual reboot upgrade during an on-demand software upgrade.

26.3.2 Preventing unwanted NE reboots

Only an operator with the admin or Network Element Software Management scope of command role can perform an NE software upgrade. The Network Element Software Management role includes a permission that you can enable to allow the configuration of NE reboots in a software upgrade policy.

If the mediation.SoftwareUpgradePolicy.property_autoRebootType permission is enabled in the Network Element Software Management scope of command role of an operator, the operator can enable or disable the NE reboot function in a software upgrade policy, and perform an NE reboot as part of a software upgrade. NE upgrade operators for whom the permission is disabled cannot configure an NE reboot in an upgrade policy or perform an NE reboot during a software upgrade.

26.3.3 Log Last Reboot

The Log Last Reboot indicator on the Shelf properties form shows the time of the previous chassis reboot and the user that performed the reboot..

NE software upgrade workflow and procedures

26.4 Workflow to manage NE software upgrades

26.4.1 Purpose

The following is the sequence of high-level actions required to perform and manage NE software upgrades for the following device types:

- 210 WBX
- 7210 SAS
- 7250 IXR
- 7450 ESS
- 7705 SAR
- 7705 SAR-Hm
- 7750 SR
- 7950 XRS
- OmniSwitch

26.4.2 Stages

1

Perform the following pre-upgrade checks.

1. It is the user responsibility to keep the valid checksum file before import.
2. Verify that the device supports the new software.
3. Verify that there is sufficient space on the NE compact flash drives for the new software image files and a backup of the current image.
4. For NEs that have redundant CPMs, use a CLI to verify that the boot environments are synchronized.

2

Configure a software upgrade policy; see 26.5 “To configure a software upgrade policy” (p. 776) .

3

Import the required device software image to the NFM-P; see 26.6 “To import device software files to the NFM-P” (p. 779) .

4

Back up the device configuration.

5

Schedule an upgrade of one or more NEs, as required; see 26.7 “To schedule an NE software upgrade” (p. 781) for node upgrades and 26.10 “To schedule an extension shelf software upgrade” (p. 787) for extension shelf upgrades.

6

Turn up, shut down, or delete a scheduled upgrade task; see 26.8 “To manage scheduled software upgrades” (p. 782) .

7

Perform an on-demand upgrade or an ISSU, as required.

See 26.9 “To perform an ISSU or on-demand software upgrade” (p. 783) for 7210 SAS, 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, or 7950 XRS upgrade information.

See 26.12 “To upgrade 7705 SAR-Hm radio card firmware” (p. 790) for 7705 SAR-Hm upgrade information.

See 26.21 “To perform an OmniSwitch on-demand software upgrade or ISSU” (p. 803) and 26.22 “To perform an OS 6400, OS 6850E, or OS 6855 ISSU” (p. 809) for OmniSwitch upgrade information.

See 26.11 “To perform an extension shelf on-demand software upgrade” (p. 788) for extension shelf upgrade information.

8

Monitor the upgrade progress; see 26.17 “To monitor the status of a software upgrade” (p. 799) .

9

Activate a software image on an NE, if required; see 26.18 “To activate a device software image” (p. 799) .

10

After the upgrade, verify the boot environment synchronization for NEs that have dual CPMs.

11

For an OS 6250SME or OS 6450, perform a software license upgrade to enable the Ethernet Metro functions; see 26.20 “To upgrade OS 6250SME and OS 6450 NE software licenses for an Ethernet (Metro) role” (p. 801) .

12

Certify or synchronize the OmniSwitch software on an NE; see 26.23 “To certify or synchronize OmniSwitch software” (p. 812) .

13

If required, reboot the NE.

14

Resynchronize the NE.

15

As required, perform upgrade verification tests.

16

Remove outdated software images from the device and from the NFM-P.

NOTE:: Ensure that you regularly remove from the NFM-P device software images that are no longer required, for example, by deleting the images or by exporting to a remote file system. An accumulation of device software images can greatly increase the length of an NFM-P database operation such as a backup, restore, or reinstantiation.

26.5 To configure a software upgrade policy

26.5.1 Purpose

Perform this procedure to configure a policy that you can use to perform an on-demand device software upgrade, ISSU, or a scheduled device software upgrade.

Note: For information about device software downgrades, contact Nokia technical support.

26.5.2 Steps

1

Ensure that the following conditions are present.

- An FTP account is configured and available on each device.
- The device configuration is backed up.

2

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

3

Click Create. The Software Upgrade Policy (Create) form opens.

4

If you are creating a policy, configure the Policy Type parameter. Perform one of the following:

a. For a 210 WBX, 7210 SAS, 7250 IXR, 7450 ESS, 7705 SAR, 7705 SAR-Hm, or 7950 XRS, choose the SR Based Node option.

b. For a GNE, choose the Generic NE Node option.

c. For an 1830 VWM, choose 1830 VWM OSU Node and go to Step 7.

5

If the Policy Type is SR Based Node, configure the Software Download, Activate Image, and Software Upgrade parameters to specify the actions that the NFM-P is to perform using the

policy. Configure the parameters in the Extension Shelf Settings panel to specify the actions performed on extension shelves that are connected to the node. Configure the Firmware Settings, see 26.12 “To upgrade 7705 SAR-Hm radio card firmware” (p. 790).

Note: The Activate Image parameter does not apply to the 210 WBX.

To check the disk space, perform one of the following:

- Enable the Validate Disk Space parameter in the Disk Space Validation Settings panel. Configure the Expected free disk space after upgrade (%) parameter to determine the expected free disk space as a percentage of the total disk space. The validation checks the available space in the target file system against the size of the image to be uploaded. If the NE file system does not have sufficient capacity, the NFM-P raises an alarm and does not complete the software upgrade for the specific NE. If the NE has two CPMs, the NFM-P checks for free space for both the active and standby CPMs. For 7705 SAR NEs, the disk space validation is only supported on the 7705 SAR-8 and 7705 SAR-18.
- On the supported NE, click Software Images, click Import and click Inspect Disk Space. This inspects the node disk space on the selected NE before performing an upgrade. The success or failure log message is displayed in the log viewer.

If the Software Download parameter is selected, the NFM-P downloads a software image to the NE when performing an activation, regardless of whether the image is already present on the NE.

If the Software Upgrade parameter is selected, perform one of the following to configure the Software Upgrade Options parameter.

Note: Depending on the permissions in your scope of command profile, the Software Upgrade Options parameter may not be configurable.

a. Select Reboot if you want the NE to reboot after the upgrade.

b. To specify an out-of-service upgrade for a 7210 SAS or 7705 SAR that includes an NE reboot with firmware upgrade, select Reboot with Firmware Upgrade.

   Note: Depending on the device configuration, a 7210 SAS or 7705 SAR determines whether a firmware upgrade is required, and may override the parameter setting.

c. To specify an ISSU for a device with dual CPMs, select ISSU (In Service Software Upgrade).

   Note: ISSUs are restricted to device maintenance releases only. See the device documentation for the supported ISSU release transitions.

d. To specify a hypervisor upgrade for a 210 WBX that includes an NE reboot after the hypervisor upgrade, select 210 WBX Reboot with Hypervisor Upgrade.

   Note: The option 210 WBX Reboot with Hypervisor Upgrade is applicable only to 210 WBX NEs. The options Reboot, ISSU (In Service Software Upgrade), and Reboot with Firmware Upgrade do not apply to the 210 WBX. However, the 210 WBX NE can be rebooted from the shelf properties form.

6

If the Policy Type is Generic NE Node, configure the Transfer Protocol parameter and the parameters in the FTP Settings or SFTP Settings panel.

Note: You must ensure that the following conditions are met before you attempt to use the software upgrade policy.

- FTP or SFTP is enabled on each main server, depending on which protocol you use.
- If a firewall is used, the required FTP or SFTP ports are open.
- The FTP or SFTP user that you specify is a valid FTP or SFTP user on each main server station, and belongs to the sam user group on each main server station.
- The nsp user has read, write, and execute access to the Root Directory that you specify.

7

If the Policy Type is 1830 VWM OSU Node, perform the following:

1. Configure the SFTP user ID and password in the 1830 Based Setting panel.
2. Configure the Server IP parameter by providing the remote server IP address required to import 1830 VWM software from a remote location. See 26.24 “To perform an 1830 VWM on-demand software upgrade” (p. 812) for more information about performing an 1830 VWM software upgrade.
3. Configure the remaining parameters.
4. Go to Step 11.

8

Click Apply.

9

Depending on the Transfer Protocol parameter value you specified, configure the FTP or SFTP parameters.

10

Configure the remaining parameters.

Note: The default value for the CFlash Image Root Path and CFlash Backup Root Path parameters specifies compact flash drive cf3, which may not exist on some devices. To check the drive availability and capacity on an NE, click FTP Browser or SSH File Browser; see 22.5 “To view an NE file system using an FTP file browser” (p. 739) and 22.2 “To view an NE file system using an SSH file browser” (p. 736) for browser usage information.
Some 7210 SAS devices are equipped with a USB flash memory (uf1). The NFM-P display does not distinguish between compact flash memory (cf) and USB flash memory (uf).

To avoid file transfer failures, Nokia recommends that you equip 7210 SAS-X NEs with a USB flash drive, and configure the CFlash Backup Root Path parameter with a file path on the USB drive (uf1:).

11

Click Apply.

12

Click on the Software Upgrade Policy Assignment tab and assign one or more NEs to the policy, as required.

13

Close the form.

END OF STEPS

26.6 To import device software files to the NFM-P

26.6.1 Purpose

Perform this procedure to import a set of device software image files or software description files to the NFM-P for use during an NE software upgrade.

Note: The Wavence SM and Wavence SA require software description files, which identify the files that require an upgrade. A software description file has a .DSC file extension.

Note: For releases prior to NFM-P 21.11, NFM-P supports 210 WBX image import only when the files onie-installer-x86_64 and md5sums.txt are present in the package. The md5sum.txt must be renamed to md5sums.txt. No additional files or folder must be present.

The NFM-P supports 210 WBX image import with all the files in Nauge-210-wbx-x package except MIBs folder. Before importing the image MIBs folder must be removed from the package.

26.6.2 Steps

1

Make the new device software files available to the NFM-P.

1. If the device software files are compressed in an archive, for example, a TiMOS ZIP file, extract the archive files to an empty directory. The directory structure in a device software archive depends on the device type and release, and must have one of the following structures:
   - flat—all files are in one directory, and there are no subdirectories

- tiered—the directory contains a boot loader file, and a subdirectory contains the software image files

For example, if you extract a set of 7750 SR files to */ExtractDir*, the directory structure is tiered:

- */ExtractDir/cflash* contains the boot.ldr file.
- */ExtractDir/cflash/R.r.Rn* (where *R.r.Rn* is the software release identifier) contains the software image files, both.tim and cpm.tim

Note:

The 7450 ESS and 7750 SR use a common software image with a Target Product name of Alcatel-SR/ESS-7XXX.

2. Copy the files to a location that is accessible to the NFM-P.

Note:

Nokia recommends that all OmniSwitch software image files, including optional and boot files, be made available.

The Wavence SM requires only the software description files.

2

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

3

Click on the Software Images tab.

4

Click on the tab that corresponds to the type of NE to be upgraded.

5

Click on the Inspect Disk Space.

6

Click Import. One of the following happens:

a. The Open form opens.

b. If upgrading 7705 SAR-Hm firmware, the Import Firmware Image form opens. After you set the Image Version and click Select, the Open form opens.

7

Navigate to the directory that contains the software files and click Open.

Note: If the directory structure is tiered, you must navigate to the directory that contains the boot loader file, then to the subdirectory that contains the software files. For the example in Step 1 , this directory is the *cflash/R.r.RN* directory for the software image files (where *R.r.RN* is the release identifier).

Some 7210 SAS devices are equipped with a USB flash memory (uf1). The NFM-P display does not distinguish between compact flash memory (cf) and USB flash memory (uf).

The NFM-P verifies the file set, imports the files, and lists an entry for the imported image.

8

If the directory does not contain only the required files, a message is displayed. Perform the following steps.

1. Click OK.
2. Copy or move files, as required, to ensure that the directory contains only the files required for the upgrade.
3. Go to Step 6 .

9

Close the Software Upgrade form.

END OF STEPS

26.7 To schedule an NE software upgrade

26.7.1 Purpose

Perform this procedure to schedule a device software upgrade on one or more managed NEs according to a software upgrade policy and an NFM-P schedule. See Chapter 5, “NFM-P-based schedules” for information about NFM-P schedules.

Note: A new NFM-P scheduled task is shut down by default and must be turned up before it can be executed, with one exception: Scheduled software upgrade tasks associated with all NE types are automatically enabled by default.

You must perform 26.6 “To import device software files to the NFM-P” (p. 779) to import the required device software image before you perform this procedure.

26.7.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Select the required software upgrade policy.

Note: The NFM-P performs the upgrade according to the configuration in the software upgrade policy that you select.

3

Click on the Software Images tab.

4

Select a software image and click Schedule Upgrades. The Select Sites form opens.

5

Select one or more NEs and click OK. The Select Schedule form opens.

Note: You cannot use a schedule in which the Ongoing parameter is enabled.

6

Select a schedule and click OK.

7

Click Yes. The NFM-P schedules the upgrade.

8

Close the form.

END OF STEPS

26.8 To manage scheduled software upgrades

Note: You cannot reuse a completed scheduled task.
The NFM-P does not delete a scheduled task after it runs; you must delete it manually.

26.8.1 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Upgrade Status tab.

3

Click Scheduled Task. The Scheduled Task form opens.

4

Select a scheduled software upgrade entry and click Properties. The Software Upgrade Scheduled Task form opens.

1. Administratively enable or disable the scheduled software upgrade, if required, by configuring the Administrative State parameter.
2. Click Properties in the Schedule panel to view the schedule information, if required.
3. Click Properties in the Task panel to view the NFM-P task information, if required.
4. Save your changes and close the form.

5

To turn up, shut down, or execute the task, click Task Action and select the appropriate option.

6

Delete the scheduled task, if required.

1. Click Task Action and choose Shut Down.
2. Click Delete. The scheduled task is deleted.

7

Close the forms.

END OF STEPS

26.9 To perform an ISSU or on-demand software upgrade

26.9.1 Purpose

Perform this procedure to upgrade the device software on one or more NEs.

The following must be true before you attempt a device software upgrade.

- You have an NFM-P user account with an administrator or network element software management scope of command role, or a scope of command role with write access to the mediation package.
- The FTP or SFTP credentials are configured in the NE mediation policy.

CAUTION

Service Disruption

A device may require a firmware upgrade before a software upgrade.

To avoid a service outage, ensure that the device firmware version supports the software release. See the device software Release Notes for information about firmware and software compatibility, and about the firmware upgrade procedure.

CAUTION

Service Disruption

Nokia recommends that you open a physical console session on the NE that you need to upgrade so that you can monitor the upgrade and recover in the event of an upgrade failure.

Read the software upgrade information in the device documentation before you perform a software upgrade.

Before you attempt an ISSU, see the appropriate device release notice and the NSP NFM-P Network Element Compatibility Guide for information about the supported upgrade paths.

You cannot upgrade an NE to a chronologically older release, regardless of the release identifier; for example, the R1 revision of a release may predate the R9 revision of the previous release. See the NSP NFM-P Network Element Compatibility Guide for information about the supported device upgrade paths.

Note: ISSUs are restricted to device maintenance releases only. See the device documentation for the supported release transitions.

Some 7210 SAS chassis types do not support ISSUs; see the device documentation for more information.

Before you downgrade a device, you must unmanage and delete the device from the NFM-P. Contact Nokia technical support for information about device downgrades.

If the NFM-P raises a NodeVersionMismatch alarm, manually unmanage and remanage the device after performing the upgrade or downgrade.

The NFM-P does not support an upgrade of the 7705 SAR-H from Release 5.0 to Release 6.1. You must first unmanage the device and then perform the upgrade using a CLI, after which you can again use the NFM-P to manage the NE.

26.9.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Images tab.

3

Select a software image and click Import. The Open window appears.

4

Navigate to the directory that contains the software image, select the image, and click Open.

Note: The directory must contain only the files required for the upgrade.
The 7450 ESS, 7750 SR, and 7950 XRS use a common software image with a Product Name of Nokia-SR/ESS-7XXX.

If you are upgrading a 7705 SAR-8 or 7705 SAR-18, two images are required: One image contains the boot.ldr and both.tim files, and the second contains the MWA files in the pkgs/MPT folder. If the image folder contains the MPT files, the MPT Image Software Version is displayed on the Software Image form. If the image folder does not contain MPT packages, the MPT Image Software Version shows N/A.

5

If the directory contains only the required files, the NFM-P imports the files, and an entry for the image is listed on the Software Upgrade form. Go to Step 7 .

6

If the directory does not contain only the required files, a message is displayed.

Perform the following steps.

1. Click OK.
2. Copy or move files, as required, to ensure that the directory contains only the files required for the upgrade.
3. Go to Step 3 .

7

Click Upgrade Sites. A list of NEs is displayed. Only the devices that support the software image are listed.

8

Select one or more NEs and click OK. The software upgrade begins.

9

Click on the Software Upgrade Status tab to view the upgrade progress.

10

CAUTION

Service Disruption

Rebooting an NE that is in service is service-affecting.

Ensure that you perform the reboot only during a scheduled maintenance period.

If the Software Upgrade parameter in the software upgrade policy is not selected, you must activate the image on each NE and then reboot each NE.

Note: Some device software upgrades, for example, ISSUs, do not require a reboot. See the device documentation for more information.
When you perform an ISSU, you can manually soft reset or hard reboot the IOMs or IMMs after the upgrade. A soft reset results in minimal downtime, but has limited support. See

26.14 “To perform an IOM, IMM, or XCM soft reset” (p. 793) to perform a manual soft reset, or 26.15 “To perform an IOM or XCM hard reboot” (p. 794) to perform a manual hard reboot.
When the IOMs are not manually soft reset or hard rebooted, the device performs a soft reset, if supported, after two hours; otherwise, the device performs a hard reboot after two hours.

1. Select the NEs that are to be upgraded.
2. Click on the Software Images tab. A list of software images on the selected NEs is displayed.
3. Select the required software image and click Activate Image. A confirmation message is displayed.
4. Click Yes. The software image is activated.

   When the NFM-P activates an NE software image, it does the following:
   - backs up the original boot.ldr at the location specified by the CFlash Backup Root Path parameter
   - synchronizes the CPMs on NEs that have redundant CPMs
   - updates the BOF with the new software image location
5. Click on the Software Upgrade Status tab to monitor the activation progress. The Last Operation column displays each operation as it completes, and the Upgrade Status column displays the overall activation status.
6. Wait until the image activation completes.
7. In the navigation tree equipment view, navigate to the shelf object. The path is Network→*NE*→Shelf.
8. Right-click the shelf icon and choose Reboot or Reboot Upgrade. A dialog box appears.

   Note:

   Only the 7210 SAS and 7705 SAR support the Reboot Upgrade option.
9. Click Yes. The NE reboots.

11

Click on the Software Upgrade Status tab to monitor the upgrade progress. The Upgrade State column displays the current upgrade status.

12

Use the NFM-P FTP or SSH file browser to verify that the transferred files and configuration are on each upgraded NE. See 22.5 “To view an NE file system using an FTP file browser” (p. 739) for information about using the FTP file browser, and 22.2 “To view an NE file system using an SSH file browser” (p. 736) for information about using the SSH file browser.

13

Resynchronize each upgraded NE.

END OF STEPS

26.10 To schedule an extension shelf software upgrade

26.10.1 Purpose

Perform this procedure to schedule a device software upgrade on one or more extension shelves, according to a software upgrade policy and an NFM-P schedule. See Chapter 5, “NFM-P-based schedules” for information about NFM-P schedules.

Note: A new NFM-P scheduled task is shut down by default and must be turned up before it can be executed, with one exception: scheduled software upgrade tasks associated with all NE types are automatically enabled by default.

You must perform 26.6 “To import device software files to the NFM-P” (p. 779) to import the required device software image before you perform this procedure.

The NFM-P automatically creates software repositories for use with extension shelf upgrades. To manually create a software repository, see 26.13 “To create a software repository” (p. 793).

26.10.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Select the required software upgrade policy.

Note: The NFM-P performs the upgrade according to the configuration in the software upgrade policy that you select.

3

Click on the Software Images tab, and then on the SR Extended Shelf Software Images tab.

4

Select a software image and click Schedule Upgrades. The Select NE form opens.

5

Select an NE and click OK. The Select Extended Shelf form opens.

6

Select one or more extended shelves and click OK. The Select Software Repository form opens.

7

A software repository is automatically created for the upgrade. To manually specify a software repository, click on the Properties button and select a software repository. Click OK. The Select Schedule form opens.

8

Select a schedule and click OK.

Note: You cannot use a schedule in which the Ongoing parameter is enabled.

9

Click Yes. The NFM-P schedules the upgrade.

10

Close the form.

END OF STEPS

26.11 To perform an extension shelf on-demand software upgrade

26.11.1 Purpose

Perform this procedure to upgrade the device software on one or more extension shelves.

The following must be true before you attempt a device software upgrade.

- You have an NFM-P user account with an administrator or network element software management scope of command role, or a scope of command role with write access to the mediation package.
- The FTP or SFTP credentials are configured in the NE mediation policy.

CAUTION

Service Disruption

A device may require a firmware upgrade before a software upgrade.

To avoid a service outage, ensure that the device firmware version supports the software release. See the device software Release Notes for information about firmware and software compatibility, and about the firmware upgrade procedure.

CAUTION

Service Disruption

Nokia recommends that you open a physical console session on the NE that you need to upgrade so that you can monitor the upgrade and recover in the event of an upgrade failure.

Read the software upgrade information in the device documentation before you perform a software upgrade.

You cannot upgrade an NE to a chronologically older release, regardless of the release identifier; for example, the R1 revision of a release may predate the R9 revision of the previous release. See the NSP NFM-P Network Element Compatibility Guide for information about the supported device upgrade paths.

Note: Before you downgrade a device, you must unmanage and delete the device from the NFM-P. Contact Nokia technical support for information about device downgrades.

The NFM-P automatically creates software repositories for use with extension shelf upgrades. To manually create a software repository, see 26.13 “To create a software repository” (p. 793).

26.11.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Images tab, then on the SR Extended Shelf Software Images tab.

3

Select a software image and click Import. The Open window appears.

4

Navigate to the directory that contains the software image, select the image, and click Open.

Note: The directory must contain only the files required for the upgrade.

5

If the directory contains only the required files, the NFM-P imports the files, and an entry for the image is listed on the Software Upgrade form. Go to Step 7 .

6

If the directory does not contain only the required files, a message is displayed.

Perform the following steps.

1. Click OK.

2. Copy or move files, as required, to ensure that the directory contains only the files required for the upgrade.
3. Go to Step 3 .

7

Click Upgrade Shelves. A list of NEs is displayed. Only the devices that support the software image are listed.

8

Select an NE and click OK. The Select Extended Shelf form opens.

9

Select one or more extended shelves and click OK. The Select Software Repository form opens.

10

A software repository is automatically created for the upgrade. To manually specify a software repository, click on the Properties button and select a software repository. Click OK. The software upgrade begins.

11

Click on the Software Upgrade Status tab to view the upgrade progress. The Upgrade State column displays the current upgrade status.

12

Use the NFM-P FTP or SSH file browser to verify that the transferred files and configuration are on each upgraded NE. See 22.5 “To view an NE file system using an FTP file browser” (p. 739) for information about using the FTP file browser, and 22.2 “To view an NE file system using an SSH file browser” (p. 736) for information about using the SSH file browser.

13

Resynchronize each upgraded NE.

END OF STEPS

26.12 To upgrade 7705 SAR-Hm radio card firmware

26.12.1 Purpose

The Sierra Wireless cellular radio card firmware is not bundled with the SAR-Hm SROS image. Instead, an operator must download a specific RHEL firmware file (Linux Version) from Sierra Wireless to the NSP host server.

For the 7705 SAR-Hm, the firmware is a zip file. For 7705 SAR-Hmc, the firmware is a tim file.

The 7705 SAR-Hm SROS software must be upgraded to Release 15.0 R6 or later before a firmware upgrade is possible.

The 7705 SAR-Hmc supports both ADP and non-ADP radio card firmware upgrade.

When ADP upgrades the 7705 SAR-Hm software, ADP also upgrades the 7705 SAR-Hm firmware if a new firmware version is available for the new 7705 SAR-Hm release.

26.12.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click Create. The Software Upgrade Policy (Create) form opens.

3

Set the Policy Type parameter to the SR Based Node option and configure the parameters.

For Non-ADP, enable the Firmware Download and Firmware Activate parameters and set the Firmware Upgrade Options parameter to Reboot.

For ADP, enable the Firmware Download and Firmware Activate parameters and set the Firmware Upgrade Options parameter to None.

Note: Depending on the permissions in your scope of command profile, the Software Upgrade Options parameter may not be configurable.

4

Click Apply.

5

Click on the Software Upgrade Policy Assignment tab and assign a 7705 SAR-Hm to the policy.

6

Close the form.

7

On the Software Upgrade form, click on the Software Images tab.

8

Click on the SAR-Hm/Hmc Firmware Images tab and click Import. The Import Firmware Image form opens.

9

Specify the Image Version and click Select.

For the 7705 SAR-Hm, use a format that is aligned with the Sierra Wireless web portal, [VENDOR NAME]<space>[FIRMWARE_SUFFIX]<space>[PRI]. For example: GENERIC 02.24.03.00 002.026_000

For the 7705 SAR-Hmc, use a format such as this example: 7705_SAR-Hmc_FW_CAT4_GS_BYPASS_0.3.3.9_V1.7.tim

The Open window appears.

10

Navigate to the directory that contains the firmware image, select the image, and click Open.

Note: The directory must contain only the firmware file.

11

If the directory contains only the required files, the NFM-P imports the files, and an entry for the image is listed on the Software Upgrade form. Go to Step 13 .

12

If the directory does not contain only the required files, a message is displayed.

Perform the following steps.

1. Click OK.
2. Copy or move files, as required, to ensure that the directory contains only the files required for the upgrade.
3. Go to Step 9 .

13

Select the imported image and click Upgrade Sites. The Select NE For Firmware Upgrade page opens.

14

Select an NE and click Next.

15

Set the Upgrade Sim parameter and click Finish. A confirmation pop-up opens.

16

Click Yes to begin the firmware upgrade.

17

To view the upgrade progress, click on the Software Upgrade Status tab and then on the Firmware Upgrade Status tab. For additional information, select the NE and click Properties.

When the firmware version is upgraded on the node, the NFM-P raises an alarm.

END OF STEPS

26.13 To create a software repository

26.13.1 Purpose

Perform this procedure to create a software repository for use with an extension shelf. A software repository defines a location where system upgrade files for the extension shelf are stored.

26.13.2 Steps

1

Choose Administration→NE Maintenance→Software Repository from the main menu.

2

Click on the Create button.

3

Configure the Primary Location parameter, and any other required parameters.

4

Save and close the form.

END OF STEPS

26.14 To perform an IOM, IMM, or XCM soft reset

26.14.1 Purpose and prerequisities

Perform this procedure to perform a manual soft reset on an IOM, IMM or XCM.

Soft reset is supported only under the following conditions:

- The IOM or XCM is operationally up.
- The IOM, IMM, or XCM is supported. The IOM, IMM, or XCM supports soft reset when the Soft Reset menu item is selectable after you right click on the object. The Soft Reset menu is greyed out if it is not supported.
- The MDAs or XMAs are provisioned, and are Ethernet, but not HSMDA.

MDAs that do not support soft reset are hard rebooted during a soft reset operation.

For more information about soft reset support, see the appropriate node documentation.

26.14.2 Steps

1

On the Equipment tree, locate a card slot object by expanding Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the object and choose Soft Reset. A Warning form opens.

3

Click View Dependencies to view the dependency information, if required.

4

Select the check box and click Yes. The NFM-P resets the IO card; the operational state of the IO card displays the soft reset progress.

END OF STEPS

26.15 To perform an IOM or XCM hard reboot

26.15.1 Steps

1

On the Equipment tree, locate an IOM or XCM card slot object by expanding Network→*NE*→Shelf→Card Slot *n*.

2

Right-click on the object and choose Properties. The Card Slot (Edit) form opens.

3

Click on the IO Card tab and click Reboot. A Warning form opens.

4

Click View Dependencies to view the dependency information, if required.

5

Select the check box and click Yes. The NFM-P resets the IO card; the operational state of the IO card displays the hard reboot status.

END OF STEPS

26.16 To upgrade the ISA-AA MDA software

26.16.1 Purpose and prerequisites

Perform this procedure to upgrade only the ISA-AA MDA software on an NE, for example, when the new software includes new AA protocol signatures.

Note: You cannot use the procedure to upgrade between major releases; only minor-release ISA-AA upgrades within the same major release are supported.

An ISA-AA MDA software upgrade is an in-service upgrade that does not require an NE reboot. Although the Software Upgrade parameter in the associated software upgrade policy must be selected for the image activation and upgrade to occur, the Software Upgrade Options in the policy do not apply to an ISA-AA MDA software upgrade.

The 7450 ESS and 7750 SR support ISA-AA MDA upgrades.

The following must be true before you attempt an ISA-AA software upgrade.

- You have an NFM-P user account with an administrator or network element software management scope of command role or a scope of command role with write access to the mediation package.
- The FTP or secure FTP credentials are configured in the NE mediation policy.

CAUTION

Service Disruption

Nokia recommends that you open a physical console session on the device that you need to upgrade. The console session allows you to monitor the upgrade and recover in the event of an upgrade failure.

Read the software upgrade information in the device documentation before you perform a software upgrade.

Before you attempt an ISA-AA MDA software upgrade, see the appropriate device release notice and the NSP NFM-P Network Element Compatibility Guide for information about the supported upgrade paths. See "NE software upgrade overview" (p. 771) for general ISSU information.

26.16.2 Steps

1

Perform the following steps.

1. Verify that the device supports the new ISA-AA software.
2. Extract the isa-aa.tim and md5sums.txt files from the Nokia software package to a directory that is reachable by the NFM-P. Ensure that there are no other files in the directory.
3. Manually verify the software image file checksum against the file checksum listed in the md5sums.txt file.
4. Verify that the device file system has space for the software image file.

5. For NEs with redundant CPMs, verify that the boot environments are synchronized by using the appropriate CLI command.
6. Configure the appropriate software upgrade policy, as described in 26.5 “To configure a software upgrade policy” (p. 776) ; apply the policy to the NEs that you need to upgrade.

2

Back up the device configuration.

3

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

4

Click on the Software Images tab and click Import. The Open window appears.

5

Navigate to the directory that contains the software image, select the directory, and click Open.

Note: The directory must contain only the isa-aa.tim and md5sums.txt files.

6

The NFM-P verifies the isa-aa.tim checksum and imports the file to the NFM-P database. An entry for the image is listed on the Software Upgrade form.

7

Select the required ISA-AA image.

8

Click Upgrade Sites. A list of NEs is displayed. Only the devices that support the software image are listed.

9

Select one or more NEs.

10

Click OK. A dialog box appears.

11

Ensure that the conditions described in the dialog box are met, and then click Yes. The software upgrade begins.

12

Click on the Software Upgrade Status tab to monitor the upgrade progress.

The following Upgrade State value is displayed:

- Transferring Image Files—The NFM-P imports the image file from the specified directory to the NFM-P database.

Note: Each upgrade stage also has one or more specific failure values.

If the Software Upgrade parameter in the software upgrade policy is selected, the NFM-P activates the new software image and displays the following sequence of Upgrade State values during the activation:

- Backing Up ISA-AA File—The NFM-P makes a backup copy of the isa-aa.tim file in the primary-image location specified by the BOF.
- Updating ISA-AA File—The new isa-aa.tim file is downloaded to the primary-image location specified by the BOF, and the appropriate ISA-AA upgrade command is issued on each NE.
- Pending ISA MDA Reboot—A reboot of the ISA-AA MDA is required.

13

If the Software Upgrade parameter in the software upgrade policy is unselected, you must activate the software image.

Note: Alternatively, you can select the required ISA-AA image on the Software Images tab and click Activate Image.

Perform the following steps.

1. Edit the following software upgrade policy settings, as described in 26.5 “To configure a software upgrade policy” (p. 776) :
   - Deselect the Software Download parameter.
   - Select the Software Upgrade parameter.
2. Select the required software image and click Upgrade Sites. The Select Sites form opens.
3. Use the form to select the required NEs, and then click OK. A dialog box appears.
4. Ensure that the conditions described in the dialog box are met, and then click Yes. The software image upgrade begins.
5. Click on the Software Upgrade Status tab to monitor the upgrade progress. The Upgrade State displays the following sequence of values:
   - Backing Up ISA-AA File—The NFM-P makes a backup copy of the isa-aa.tim file in the primary-image location specified by the BOF.
   - Updating ISA-AA File—The new isa-aa.tim file is downloaded to the primary-image location specified by the BOF, and the appropriate ISA-AA upgrade command is issued on each NE.
   - Pending ISA MDA Reboot—A reboot of the ISA-AA MDA is required.

14

Perform the following steps on each upgraded ISA-AA MDA to shut down and reboot the MDA.

1. In the navigation tree equipment view, navigate to the upgraded ISA-AA MDA. The path is Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*.
2. Right-click on the Daughter Card Slot object and choose Shut Down. A dialog box appears.
3. Click View Dependencies. A dialog box displays the number of objects that shutting down the ISA-AA MDA may affect.
4. Click OK.
5. When you are certain that shutting down the MDA has no unintended effects, select the check box and click Yes. The NFM-P shuts down the ISA-AA MDA.
6. Right-click on the Daughter Card Slot object and choose Reboot. A dialog box appears.
7. Click View Dependencies. A dialog box displays the number of objects that rebooting the ISA-AA MDA may affect.
8. Click OK.
9. When you are certain that rebooting the MDA has no unintended effects, select the check box and click Yes. The NFM-P reboots the ISA-AA MDA.
10. After the MDA reboots, right-click on the Daughter Card Slot object and choose Turn Up. A dialog box appears.
11. Click Yes. The NFM-P turns up the ISA-AA MDA.

15

Click on the Software Upgrade Status tab to monitor the upgrade progress.

The Upgrade State column displays the current upgrade status, which is one of the following:

- Success—The new image is successfully loaded on each ISA-AA MDA.
- Failed to upgrade ISA—The upgrade is a failure on at least one ISA-AA MDA.

16

If required, use the NFM-P FTP or SSH file browser to verify that the transferred file and configuration are on each upgraded NE. See 22.5 “To view an NE file system using an FTP file browser” (p. 739) for information about using the FTP file browser, and 22.2 “To view an NE file system using an SSH file browser” (p. 736) for information about using the SSH file browser.

17

Update each ISA-AA policy, as required, to enable the new protocol signatures. See Chapter 87, “Application assurance” for information about configuring AA policies.

END OF STEPS

26.17 To monitor the status of a software upgrade

26.17.1 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Upgrade Status tab.

3

Click Search. A list of NEs is displayed.

4

Select an NE and click Properties. The Software Upgrade Status (View) form opens.

5

Click on the Software Upgrade tab to view status information about the upgrade.

6

Close the forms.

END OF STEPS

26.18 To activate a device software image

26.18.1 Purpose

Perform this procedure to activate a device software image on an NE.

Note: If the BOF update fails, the original boot.ldr file is put in place to align with the BOF specification.

The NFM-P ensures that the software image is present on the NE and valid for the device before it updates the BOF.

26.18.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Upgrade Status tab.

3

Select an NE.

4

Click on the Software Images tab. The software images on the NE are listed.

5

Select a software image and click Activate Image; a confirmation message is displayed.

6

Click Yes. The software image is activated.

When the NFM-P activates an NE software image, it does the following:

- backs up the original boot.ldr at the location specified by the CFlash Backup Root Path parameter
- synchronizes the CPMs on NEs that have redundant CPMs
- updates the BOF with the new software image location

7

Click on the Software Upgrade Status tab to monitor the activation progress. The Last Operation column displays each operation as it completes, and the Upgrade Status column displays the overall activation status.

8

Close the form.

END OF STEPS

26.19 To export a device software image from the NFM-P to a GUI client file system

26.19.1 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Images tab.

3

Click on the appropriate tab for the device type, select a software image and click Export. A file browser form opens.

4

Navigate to the directory that is to contain the exported software image and click OK. The software image is saved as files in the specified directory.

5

Close the form.

END OF STEPS

26.20 To upgrade OS 6250SME and OS 6450 NE software licenses for an Ethernet (Metro) role

26.20.1 Purpose

Perform this software license upgrade to allow OS 6250SME and OS 6450 NEs to perform in the Ethernet (Metro) role and support the OmniSwitch Ethernet (Metro) features such as ERP, SAA/OAM, stacking/IPM Enterprise VLANs, Dying Gasp, Link OAM, CPE Test-Head profiles, and group profiles. This procedure is not applicable OS 6250 Ethernet (Metro) nodes.

After you upgrade the OS 6250SME and OS 6450 NEs for an Ethernet (Metro) role, perform 26.21 “To perform an OmniSwitch on-demand software upgrade or ISSU” (p. 803) to perform any future software license upgrades on these NEs.

26.20.2 OmniSwitch restrictions and prerequisites

CAUTION

Service Disruption

This procedure is service-affecting because the OS 6250SME and OS 6450 NEs will reboot.

Ensure that the license upgrade activity occurs during a maintenance window.

The following prerequisites are required before you can perform an upgrade:

- You need a new license key for each OS 6250 and OS 6450 OmniSwitch to be upgraded. Contact your Nokia sales representative if you need a new software license.
- You must copy the new license key to a text file and position the file in a folder in the node file system so the file can be imported into the NFM-P. Record the path to the folder where the license key text file is located.

26.20.3 Steps

1

Choose Administration→NE Maintenance→License Upgrade from the NFM-P main menu. The License Upgrade form opens with the License Upgrade Policy tab selected.

2

Optionally, click on the License Upgrade Status tab to confirm the software version and license policy that is currently applied to the NE being upgraded. Otherwise, go to Step 3 .

1. Select an NE and click Properties. The License Upgrade Status - [Policy Name] [Router ID] [View] form appears.
2. Confirm the software version that is installed and the license policy that is applied to the NE.
3. Click Cancel to close the form.
4. Click on the License Upgrade Policy tab.

3

Perform one of the following:

a. To create a new AOS license policy, go to Step 4 .

b. To modify the default AOS default license policy, go to Step 7 .

4

Click Create. The License Upgrade Policy [Create] form opens.

5

Configure the required parameters:

- Policy ID
- Auto-Assign ID
- Name
- Root Path

6

Click OK to save the changes and close the License Upgrade Policy [Create] form. The new AOS default license policy appears on the License Upgrade form. Go to Step 10 .

7

Select the AOS default license policy on the License Upgrade form, and click Properties. The License Upgrade Policy - AOS Default Policy [Edit] form opens.

8

As required, modify the Root Path parameter to specify the license text file location.

9

Click OK to save the changes and close the License Upgrade Policy [Edit] form. The modified AOS Default License policy appears on the License Upgrade form.

10

Click on the Licenses tab on the License Upgrade form.

11

Click Import. The Open form appears.

12

Locate the license key text file in the appropriate folder in the node file system and click Open. The license_key.txt file appears on the Licenses tab on the License Upgrade form.

13

Optionally, review the license key information before the installation by double-clicking on the license_key.txt file. The License Info form appears. Click Cancel to close the form and return to the License Upgrade form.

14

Select a license to install and click Install License. The Select Sites-Select Sites form appears.

15

Select one or more NEs to upgrade and click OK.

16

Click Yes to proceed.

17

Optionally, click on the License Upgrade Status tab on the License Upgrade form to confirm whether the upgrade on the NEs was successful. The status is displayed in the Upgrade State column.

18

Click Close to close the License Upgrade form.

END OF STEPS

26.21 To perform an OmniSwitch on-demand software upgrade or ISSU

26.21.1 Purpose

Perform this procedure to upgrade the device software on an OmniSwitch.

On-demand software upgrades are supported on all OmniSwitch variants.

Not all OmniSwitch variants support ISSU. This procedure describes how to perform an ISSU on the OS 6860, OS 6860E, OS 6860N, OS 6900, OS 9700E, OS 9800E, and OS 10K. For

information about performing an ISSU on an OS 6400, OS 6850E, or OS 6855, see 26.22 “To perform an OS 6400, OS 6850E, or OS 6855 ISSU” (p. 809) .

You can perform the following types of software upgrades:

- Image file
- Boot files
- CPLD/FPGA files (OS 6250 and OS 6450, Release 6.6.3-453R01 and later only)

26.21.2 OmniSwitch restrictions

CAUTION

Service Disruption

An OmniSwitch may require a firmware upgrade before a device software upgrade.

To avoid a service outage, ensure that the device firmware version supports the software upgrade. See the device software Release Notes for information about firmware and software version compatibility.

CAUTION

Service Disruption

Nokia recommends that you establish a physical console session on the device that you need to upgrade.

This console session will allow you to monitor the upgrade and recover the device in the event of an upgrade failure.

The following conditions must be true before you can perform a device software upgrade on an OmniSwitch:

- You have an NFM-P user account with an administrator or network element software management scope of command role or a scope of command role with write access to the mediation package.
- FTP is configured in the NE mediation policy.

The following are the OmniSwitch software upgrade restrictions and requirements.

- You can perform an ISSU on the OS 6860, OS 6860E, OS 6860N, and OS 6900, when stacked in virtual chassis mode.
- The OS 9700E, OS 9800E, and OS 10K NE platforms must be fully synchronized and certified.
- Target images must be loaded to the /flash/issu directory.
- Sufficient flash memory must be available for upgrade images.
- OS 9700E and OS 9800E NEs running an ‘R##’ build, such as 6.4.1.123.R01 do not support ISSU patches. The NE must first be upgraded to an ‘S##’ build such as 6.4.1.123.S01. This does not apply to OS 10K NE platforms.
- The directory structure that stores the image and configuration files is divided into two parts:

- The certified directory contains files that have been certified by an authorized user as the default files for the switch.
- The working directory contains files that may or may not be modified from the certified directory. The working directory is a holding place for new files. Files in the working directory must be tested before you commit them to the certified directory.

Note: To perform an ISSU, use the software upgrade window only.

The OS 6400, OS 6850E, OS 9700E, OS 9800E, and OS 10K support standard software upgrades in addition to ISSU. On an OmniSwitch, the introduction of new images requires a system reload which disrupts all data traffic during the reload process. Data traffic loss is limited to L3 base routing instance traffic; no loss of Layer 2 data traffic should occur.

26.21.3 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Perform one of the following:

a. If you need to upgrade the image files or the image and boot files, go to Step 3 .

b. If you need to upgrade the boot files, go to Step 17 .

c. If you need to upgrade the FPGA files, go to Step 30 .

3

Select the appropriate software upgrade policy.

Note: The NFM-P performs the upgrade according to the configuration in the software upgrade policy that is assigned to the NE.

4

Click on the Software Images and AOS Software tabs.

5

Select a software image.

If you are performing an ISSU on an OS 6860, OS 6860E, OS 6860N, or OS 6900, in virtual chassis mode, the ISSU package is identified in the list on the AOS Software tab by the word “present” in the AOS Software ISSU file column. The NFM-P will not validate the ISSU software package; it assumes that the ISSU software package is valid and contains the required files and image. During an ISSU, the NFM-P upgrades the secondary shelves and automatically switches one secondary shelf to become the primary shelf, then upgrades the former primary shelf, which becomes a secondary shelf.

6

Click Transfer to Sites. A list of NEs opens. The list is filtered to display only the device type that is appropriate for the selected software image.

7

Select one or more NEs.

8

Click OK. The selected software image file is uploaded to the working directory of the selected NEs.

9

Click on the Software Upgrade Status tab to view the status of the upgrade as it progresses. Wait until the files have been successfully transferred before going to Step 10 .

10

Click Reload working Sites. A list of NEs opens.

11

Select one or more NEs.

12

CAUTION

Service Disruption

Rebooting an NE that is in service is service-affecting.

Ensure that the reboot activity occurs during a maintenance window.

Click OK. The selected NEs reboots using the new software image that was uploaded to the working directory.

Note: Nokia recommends monitoring the switch to ensure that the reboot completes successfully.

13

Click Certify Sites. A list of NEs opens.

Note: Only software that is thoroughly validated as viable and reliable software should be copied to the certified directory. After you copy the software to the certified directory, you cannot recover a previous version of the image or configuration files.

14

Select one or more NEs.

15

Click OK. The software image stored in the NE working directory is copied to the certified directory. The working directory and the certified directory are synchronized so that the same files are in both directories.

16

Perform one of the following:

a. If you need to upgrade the boot files, go to Step 24 .

b. If you need to upgrade the FPGA files, go to Step 31 .

c. If you are upgrading only the image files, go to Step 35 .

17

Select the appropriate software upgrade policy.

Note: The NFM-P performs the upgrade according to the configuration in the software upgrade policy that is assigned to the NE.

18

Click on the Software Images and AOS Software tabs.

19

Select a software image.

20

Click Transfer to Sites. A list of NEs opens. The list is filtered to display only the device type that is appropriate for the selected software image.

21

Select one or more NEs.

22

Click OK. The boot files are uploaded to the root directory of the selected NEs.

23

Click on the Software Upgrade Status tab to view the status of the upgrade as it progresses. Wait until the files have been successfully transferred before going to Step 24 .

24

Click Upgrade Boot files. A list of NEs opens. The list is filtered to display only the device type that is appropriate for the selected files.

25

Select one or more NEs.

26

Click OK. The boot files are upgraded on the selected NEs.

27

Click on the Software Upgrade Status tab to view the status of the upgrade as it progresses. Wait until the files have been successfully transferred before going to Step 28 .

28

Click Delete boot files.

29

Perform one of the following:

a. If you need to upgrade the FPGA files, go to Step 31 .

b. If you do not need to upgrade the FPGA files, go to Step 35 .

30

Click on the Software Images and AOS Software tabs.

31

Select a software image.

32

Click Upgrade FPGA files. A list of NEs opens. The list is filtered to display only the device type that is appropriate for the selected file type.

33

Select one or more NEs.

34

Click OK. The selected FPGA files are uploaded to the root directory of the selected NEs.

35

Click on the Software Upgrade Status tab to monitor the upgrade progress. The Upgrade State column displays the current upgrade status.

After the FPGA files are successfully transferred, the selected NEs are rebooted.

END OF STEPS

26.22 To perform an OS 6400, OS 6850E, or OS 6855 ISSU

26.22.1 Prerequisites

This procedure applies to the following OmniSwitch NEs, in stacked configuration:

- OS 6400, Release 6.4.5 R02 or later
- OS 6850E, Release 6.4.5 R02 or later
- OS 6855, Release 6.4.5 R02 or later

ISSU cannot be used to upgrade from one major release to another.

26.22.2 Steps

1

In the navigation tree equipment view, expand the icon of the OmniSwitch to upgrade, right-click on an OmniSwitch shelf object in the equipment view, and click Properties. The Shelf (Edit) form opens.

2

Click on the Software Control Module tab.

3

Click Resynch and click Yes to clear the dialog box.

4

Ensure that one of the following is true:

- The Certify Status is Need Certify and the Synchronization Status is Not Synchronized.
- The Certify Status is Certified and the Synchronization Status is Synchronized.

5

Click OK to apply the changes. A dialog box appears. Click Yes. The Shelf (Edit) form closes.

6

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

7

Select the ISSU software upgrade policy for an AOS Based Node.

Note: For stacked NEs, the ISSU directory should be created automatically under /flash/issu. If the ISSU directory does not exist, it can be created by setting the File Transfer Type parameter to Secure on the NE being upgraded. See 9.17 “To configure device mediation” (p. 301) for information.

8

Click on the Software Images and AOS Software tabs.

9

Select a software image.

Note: ISSU software images should not contain Kencrypt image files. If the ISSU software image contains a Kencrypt image file, it will transfer the image to the working directory and perform an on-demand software upgrade.
For OS 6400, ISSU software images should not contain Gdiag image files. If the ISSU software image contains a Gdiag image file, it will transfer the image to the working directory and perform an on-demand software upgrade.

10

Click Transfer to Sites. A list of NEs opens. The list is filtered to display only the device type that is appropriate for the selected software image.

11

Select one or more NEs.

12

Click OK in the dialog boxes that appear and click Yes. The boot files are uploaded to the root directory of the selected NEs.

13

Click on the Software Upgrade Status tab to view the status of the upgrade as it progresses. Once the files have been successfully transferred, continue to Step 14 .

14

Click on the Software Images tab.

15

Click Certify Sites. A list of NEs opens.

Note: Only software that is thoroughly validated as viable and reliable software should be copied to the certified directory. After you copy the software to the certified directory, you cannot recover a previous version of the image or configuration files.

16

Select one or more NEs.

17

Click OK. A dialog box appears.

18

Click Yes. The software image stored in the NE working directory is copied to the certified directory. The working directory and the certified directory are synchronized so that the same files are in both directories.

19

Click on the Software Upgrade Status tab to view the status of the upgrade as it progresses. Once the sites have been successfully certified, continue to Step 20 .

Note: After the desired sites have been certified, the Command to Apply parameter should be set to Flash Synchro on each affected shelf; see 15.41 “To manage an OmniSwitch running configuration” (p. 505) for configuration information.

20

Click on the Software Images tab.

21

Click Reload Sites. A list of NEs opens.

22

Select one or more NEs and click OK. A dialog box appears.

23

CAUTION

Service Disruption

Rebooting an NE that is in service is service-affecting.

Ensure that the reboot occurs only during a scheduled maintenance period.

Click Yes. The selected NEs reboot using the new software image that was uploaded to the working directory.

24

Click on the Software Upgrade Status tab to view the upgrade status.

END OF STEPS

26.23 To certify or synchronize OmniSwitch software

26.23.1 Steps

1

Choose Administration→NE Maintenance→OMNI Software Maintenance from the NFM-P main menu. The OMNI Software Maintenance form opens.

2

Select one or more NEs and perform one of the following:

a. Click Certify. Go to Step 5 .

b. Click Certify and Synchronize. Go to Step 5 .

c. Click Flash Synchronization. Go to Step 5 .

d. Click Properties. The Software Control Module (Edit) form opens.

3

Configure the Command to Apply parameter.

4

If the Command to Apply parameter is set to Reload, configure the Delayed Activation Timer parameter.

5

Save your changes and close the forms.

END OF STEPS

26.24 To perform an 1830 VWM on-demand software upgrade

26.24.1 Before you begin

Configure a software upgrade policy or use the default “VWM OSU Default Policy”.

See 26.5 “To configure a software upgrade policy” (p. 776) for more information about configuring a software upgrade policy.

26.24.2 Steps

1

Choose Administration→NE Maintenance→Software Upgrade from the NFM-P main menu. The Software Upgrade form opens.

2

Click on the Software Images tab and then on the 1830 VWM OSU Software Images tab.

3

Note: NFM-P requires only CATALOG and osu.isd files to be present in the specified folder to import the image. Copy the CATALOG and osu.isd files to a separate folder or remove all other files from the folder before you begin with import operation.

Perform one of the following:

a. Import image files from the NFM-P main server file system:

1. Click Import. The Select Import VWM Software form opens.
2. Select a software image and click OK. The NFM-P saves the image in a directory below /opt/nsp/nfmp/nesoftware/SAM/VWMOSU/.

   When the import is complete, an entry for the image file appears in the list.

b. Import 1830 VWM software from a remote location.

Note: The remote system address is the server IP address that is configured in the software upgrade policy. See Step 7 of the 26.5 “To configure a software upgrade policy” (p. 776) for more information.

1. Click Remote Path. A Remote Path notification box appears.
2. Click Yes. The Software Directory Path form opens.
3. Enter the remote software path, which is an absolute path in the following format:

   /directory/device-software_version

   For example:

   /root/Desktop/1830VWMOSU-8.5-0

When the remote path import is complete, an entry for the image file is listed on the form.

4

Choose an appropriate 1830 VWM software image and click Audit. The Confirm Audit Settings dialog box opens.

5

Configure the Audit Settings parameter, choose an 1830 VWM device and click OK.

6

Click Load and choose the 1830 VWM device.

7

Click Activate and choose the 1830 VWM device.

8

Click Revert and choose an 1830 VWM device to revert to a previously committed software release, if required.

9

Click Commit and choose the 1830 VWM device.

10

Save your changes and close the forms.

END OF STEPS

Part IV: Network management

Overview

Purpose

This part provides information about network management using the NFM-P.

Contents

27 NE routing and forwarding

27.1 NE routing and forwarding

27.1.1 NE routing and forwarding objects

You can configure and manage the following routing and forwarding objects on managed NEs using the NFM-P navigation tree routing view:

- routing instances
- VRF instances
- web portal routing instances
- IS-IS instances
- OSPFv2 and OSPFv3 instances
- static routes
- L3 interfaces
- network domains

You can also configure the routing protocols assigned to routing and forwarding objects such as BGP, L2TP, MPLS, and LDP. The supported protocols are device-dependent and are displayed in the navigation tree routing view as child objects of the routing instance. See Chapter 28, “Routing protocol configuration” for more information about configuring routing protocols.

The following figure shows an example of the various routing objects that can be displayed in the routing view of the navigation tree.

Figure 27-1 Routing and forwarding objects

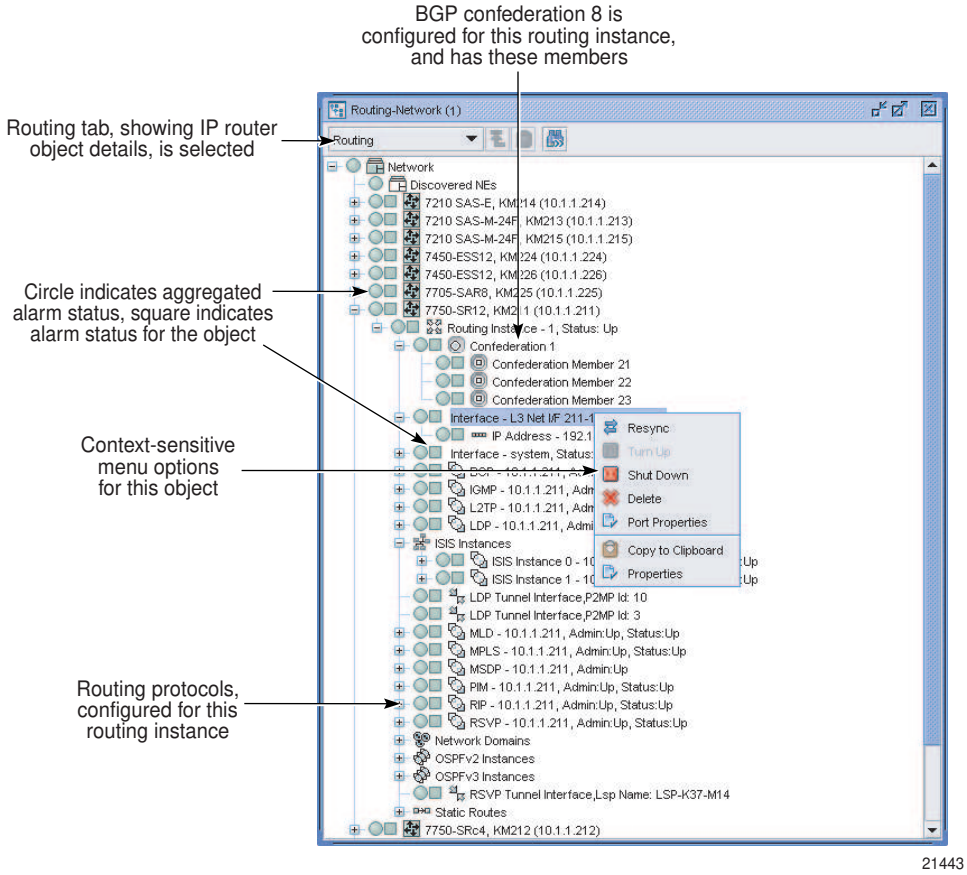

27.1.2 Routing instances

Routing instances are virtual routers. A routing instance in routing navigation tree is a representation of all of the characteristics/objects assigned to the router such as the interfaces, protocols, static routes, network domains, and other instance types (IS-IS, OSPF) being managed by the NFM-P. The supported objects are displayed in the navigation tree as child objects of the routing instance.

Routing instances contain their own dedicated interfaces and routing tables that are used to deploy multiple logical devices in one physical chassis.

27.1.3 DHCP relay and snooping on OmniSwitch routing instances

A DHCP relay agent is a BOOTP relay agent that relays DHCP messages between DHCP clients and DHCP servers on different IP networks. The Omniswitch supports global or per-VLAN BOOTP/DHCP relay service and per-VLAN UDP port relay services.

DHCP Option 82 and DHCP snooping provide security for a DHCP relay service. The DHCP Option 82 switch-level feature allows the relay agent to insert identifying information into client-originated DHCP packets before the packets are forwarded to the DHCP server. DHCP Snooping, at the switch or VLAN level, improves network security by filtering DHCP messages received from devices outside the network and building and maintaining a binding table to track access information for each device. DHCP relay Option 82 and snooping cannot be run on the switch at the same time.

DHCP relay and snooping is configured at the default routing instance level, see 27.12 “To configure UDP relay, DHCP snooping, and DHCP Option 82 on OmniSwitch routing instances” (p. 850) . In addition, per-VLAN DHCP snooping can be enabled on a VLAN site. See 75.7 “To create a standard VLAN service on OmniSwitch devices” (p. 2075) , 75.8 “To create an OmniSwitch L2 VPN TLS VLAN service” (p. 2076) , and 75.9 “To create an OmniSwitch BTV VLAN service” (p. 2078) for more information.

It is necessary to configure ports connected to DHCP servers within the network and/or firewall as trusted ports so that necessary DHCP traffic to/from the server is not blocked. Configuring the port mode as trusted also identifies the device connected to that port as a trusted device within the network. See 16.56 “To configure OmniSwitch Ethernet ports” (p. 647) for more information.

27.1.4 Web portal routing instances

You can use a web port routing instance to configure the WPP protocol running between a BRAS/BNG and a Web portal server. WPP is used for web portal authentication of WLAN users (DHCP Host). A web portal routing instances can be configured on a base routing instance or a VPRN routing instance. See 28.19 “WPP” (p. 907) in Chapter 28, “Routing protocol configuration” for more information.

27.1.5 VRF instances

Virtual routing and forwarding (VRF) allows multiple instances of a routing table to co-exist within the same router at the same time. This increases functionality by allowing network paths to be segmented without using multiple devices. Because traffic is automatically segregated, VRF also increases network security and can eliminate the need for encryption and authentication. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other.

27.1.6 IS-IS and OSPF instances

You can use IS-IS and OSPF non-forwarding instances to separate a very large network into smaller administrative entities. Instead of configuring a large number of filters, non-forwarding instances can be used to filter routes, thereby instantiating policies. Non-forwarding instances can be used to reduce the amount of routing information advertised throughout all components of a network. Routing information associated with a particular instance can be announced where required, instead of being advertised to the whole network.

27.1.7 Static routes

Static routing uses a manually-configured routing entry, rather than information from a dynamic routing protocol to forward traffic and never change unless you explicitly update them. Additionally, all traffic destined for a static address is routed through the same router and are automatically imported into the routing table when a router comes online.

Static routing is useful for networks with customers whose traffic must always flow through the same routers. Static routing avoids the bandwidth cost and route import latency of dynamic routing.

27.1.8 L3 network interfaces

An L3 network interface is a logical IP object that is defined on a physical port, such as an Ethernet port. An L3 interface:

- associates an IP address and subnet mask with a physical port or channel
- has a physical port or channel cabled to another device
- requires QoS policy configuration
- requires routing protocol configuration

The physical connection of one device to another device is through a port or channel. However, the L3 network interface determines its IP connectivity. The L3 network interface passes both routing information and IP traffic.

You can configure the following L3 network interface types on routing instances with the NFM-P:

- ICMP interface
- system interface
- management interface
- IGMP interface
- PIM interface
- MLD interface
- RSVP tunnel interface
- LDP tunnel interface
- multi-homing interface
- IS-IS interface

27.1.9 DoS protection on L3 network interfaces

In a subscriber aggregation network, an NE typically receives few control-plane packets from a specific subscriber. If one or more subscribers generate excessive control-plane traffic, DoS protection policies can help to ensure that NEs do not become overburdened by these unwanted packets. Global DoS protection controls the arrival rate for unprovisioned link-layer protocol packets from CE devices.

You can use the NFM-P to create DoS protection policies and apply them to L3 network interfaces. A DoS protection policy limits the number of control-plane protocol packets that are received each second from a subscriber host, and optionally logs a violation notification if a policy limit is exceeded. The interface drops the excessive packets before they are queued or processed by the NE. You can use the NE System Security form to view the violations for a specific NE.

You can apply DoS protection policies to control the following on L3 network interfaces:

- the control-plane packet arrival rate per subscriber host
- the overall control-plane packet arrival rate for the interface
- whether an NE sends a notification trap if a policy limit is exceeded

An NE that supports DoS protection automatically assigns a default DoS protection policy to each L3 network interface. The default policy limits only the overall packet arrival rate for the interface, and cannot be deleted or modified.

You can also apply DoS protection policies to certain L2 and L3 access interfaces. See the appropriate service chapter for information about applying DoS protection policies to access interfaces.

You can configure a global NE DoS protection policy and distribute it to NEs by choosing Administration→Security→NE DoS Protection from the NFM-P main menu. See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for more information about configuring and applying NE DoS protection policies.

27.1.10 DDoS protection on network and access interfaces

You can configure a global NE DDoS protection policy and distribute it to NEs by choosing Administration→Security→NE DDoS Protection from the NFM-P main menu. You can configure a DDoS protection policy on an L2 and L3 network interface. See the procedure to configure an NE DDoS protection policy in the *NSP System Administrator Guide* for information about configuring a DDoS protection policy.

You can also apply DDoS protection policies to certain L2 and L3 access interfaces and SAPs. See the appropriate service chapter for information about applying DDoS protection policies to access interfaces.

27.1.11 System interfaces

The system interface is associated with a network entity, such as a specific device, not a specific interface.The system interface is also referred to as the loopback interface.

The system interface is used to preserve connectivity when an interface fails or is removed. A system interface must have an IP address and a 32-bit subnet mask. The system interface is used as the device identifier by higher-level protocols such as OSPF and BGP. The system interface is associated during the configuration of the following entities:

- the termination point of service tunnels
- the hops when configuring MPLS paths and LSPs
- the addresses on a target device for BGP and LDP peering

27.1.12 NE routing policies

Policy-based routing (PBR) gives you a flexible means of routing packets over devices using the NFM-P by allowing you to configure a defined policy that override default routing protocol decisions on the router for traffic flows, and lessening reliance on routes derived from routing protocols. PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.

You can set up PBR as a way to route packets based on configured policies. For example, you can implement routing policies to allow or deny paths based on the identity of a particular end system, an application protocol, or the size of packets.

There are no default routing policies on the NFM-P. Each policy must be created and applied to a routing object, a routing protocol, or the forwarding table. Each set of rules that is associated with controlling routes are called routing policy statement entries on the client GUI. See Chapter 54, “Routing policies” for more information on routing policies.

27.1.13 Self-generated traffic and QoS

The NFM-P supports QoS configuration for self-generated traffic (SGT) on routing instances and VPRN sites. NEs produce SGT for various applications, such as Telnet, SNMP, SSH, and others. For each application, you can configure the DSCP or dot1p value for the traffic generated by that application. You can also map DSCP values to forwarding classes. See 27.15 “To configure QoS for self-generated traffic on a routing instance” (p. 854) and 79.13 “To configure QoS for self-generated traffic on a VPRN site” (p. 2546).

27.1.14 Network domains

You can use the navigation tree Routing view to configure and manage network domains and to associate network interfaces or service tunnels with the network domain. For supported devices, there is a default Network Domain that is assigned to the associate Routing Instance. The default network domain cannot be deleted but the properties associated with it can be changed.

By default, all network interfaces in a routing instance belong to the default network domain. You can associate an interface to any user defined network domain. The loopback and system interfaces cannot be associated with user defined network domains.

Network domains help determine which network ports are eligible to transport traffic of individual SDPs. This information is used for the SAP-ingress queue allocation which is applied to VPLS SAPs. No SAP-ingress queues are allocated if a given port does not belong to the network domain used in a VPLS.

Note: A maximum of four network domains are supported in any VPLS.

If an SDP is used for E-PIPE, I-PIPE or A-PIPE bindings, the network domain configuration is not considered.

Network domains are not applicable to loopback and system interfaces.

An SDP can be assigned to only one network domain. If no user defined network domain is created, an SDP will be assigned to the default network domain. All SAPs in VPLS will have a queue reaching all fwd complexes serving interfaces belonging to the same network domains as the SDPs. You can assign or remove a network domain association of the interface or SDP without deleting the respective object.

27.1.15 TCP MSS adjustment for PPPoE sessions

The NFM-P supports ISA-based TCP MSS adjustment for PPPoE sessions. The functionality allows a BNG to adjust TCP MSS values, preventing PPP clients from sending large packets which could be dropped in an access network.

In the NFM-P, TCP SYN-flagged packets are matched by an IP filter policy configured with a filter action of TCP MSS Adjust, and forwarded through an internal LAG SAP to an ISA-BB provisioned on the routing instance. Base and VPRN routing instances are associated with an ISA-BB through an ISA NAT group containing multiple ISA-BBs. The ISA-BB adjusts the TCP MSS value if:

- the MSS option is not present, in which case MSS field is inserted into the packet with the MSS value configured on the routing instance.
- the MSS value is larger than the value configured on the routing instance, in which case the packet MSS value is changed to the configured value.

27.1.16 Workflow to configure TCP MSS adjustment

1

Configure an ISA NAT group to be used for MSS adjustment; see 30.3 “To configure an ISA-NAT group” (p. 1080).

2

Associate the ISA-NAT group with a base routing instance or VPRN routing instance and configure the MSS value on the routing instance; see 27.2 “To configure a routing instance or a VRF instance” (p. 826) or 79.36 “To configure TCP MSS adjustment on a VPRN site” (p. 2586).

3

Create an IPv4 or IPv6 filter policy with a filter action of TCP MSS Adjust to match TCP packets requiring MSS adjustment; see 51.5 “To configure an ACL IP filter policy” (p. 1671) or 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

4

Apply the IPv4 or IPv6 filter policy to an SLA profile or L3 network interface; see 64.5 “To configure an SLA profile” (p. 1845) or 27.18 “To configure L3 network interfaces” (p. 863).

27.1.17 Workflow to configure NE routing and forwarding

1

Configure a base routing instance or VRF instance on a device; see 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

2

Enable the routing protocols to be supported on the device or routing instance. The supported protocols are device dependent. See the “Routing protocol configuration workflow and procedures” (p. 909) in Chapter 28, “Routing protocol configuration” for more information about configuring routing protocols.

3

Perform the following actions on the base routing instances as required:

- configure a local DHCPv4 or DHCPv6 server; see 27.5 “To configure a local DHCPv4 server on a routing instance” (p. 839) or 27.6 “To configure a local DHCPv6 server on a routing instance” (p. 842).
- perform a Force Partner Down action on a local DHCP server failover; see 27.7 “To perform a Force Partner Down action on a local DHCP server failover” (p. 845).
- configure a RADIUS server; see 27.8 “To configure a RADIUS server on a routing instance” (p. 846).
- configure a RADIUS proxy server; see 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847).

- configure UDP relay/DHCP snooping on an OmniSwitch; see 27.12 “To configure UDP relay, DHCP snooping, and DHCP Option 82 on OmniSwitch routing instances” (p. 850).
- configure static routes; see 27.13 “To configure a static route on a routing instance” (p. 852).
- configure QoS for self-generated traffic; see 27.15 “To configure QoS for self-generated traffic on a routing instance” (p. 854).
- configure LSP entries with indirect static routes; see 27.16 “To configure LSP entries with indirect static routes” (p. 855).
- configure a BGP confederation; see 28.30 “To configure a BGP confederation” (p. 916) .
- configure an LDP peer or LDP targeted peer; see 28.54 “To configure an LDP peer” (p. 946) or 28.53 “To configure an LDP targeted peer” (p. 945) .
- configure ECMP for LDP routing; see 28.55 “To configure ECMP for LDP routing” (p. 948) .
- configure an IS-IS link group, or an IS-IS NET address; see 28.60 “To configure IS-IS on a routing instance” (p. 953) or 28.61 “To configure an IS-IS link group on a routing instance” (p. 956) .
- configure Web Portal routing instances; see 28.134 “To create a web portal routing instance” (p. 1051) .
- configure NAT; see 30.9 “To configure NAT on a routing instance” (p. 1086) .

4

Configure L3 network interfaces as required:

a. For the following L3 network interfaces; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863) .

- ICMP interface
- system interface
- management interface
- IGMP interface
- control tunnel
- PIM interface
- MLD interface

Perform the following steps:

1. Associate a network port with the L3 interface.
2. Enable the multicast routing protocols on the interfaces as required. See Chapter 28, “Routing protocol configuration” for more information about multicast configuration and parameters.
3. Enable bridging on the interfaces as required. The NFM-P supports bridging on OmniSwitch devices. See Chapter 28, “Routing protocol configuration” for more information about bridging.
4. Assign a DoS protection policy to the interface, as required; see 27.18 “To configure L3 network interfaces” (p. 863) .
5. If you enabled OSPF on the device, add a L3 network interface or create an OSPF area range as required; see 28.69 “To add a Layer 3 interface to an OSPF router” (p. 965) or 28.70 “To create an OSPF area range” (p. 967) .

6. Assign an IES or network interface to a virtual router as a VRRP instance. See Chapter 37, “VRRP” for more information about VRRP configuration and parameters.

b. RSVP tunnel interfaces; see 27.2 “To configure a routing instance or a VRF instance” (p. 826)

c. Multi-homing interfaces; see 27.2 “To configure a routing instance or a VRF instance” (p. 826)

d. LDP tunnel interfaces; see 28.52 “To configure an LDP interface” (p. 944) .

e. IS-IS interfaces; see 28.63 “To configure an IS-IS interface” (p. 958) .

f. PIM interfaces; see 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013) .

g. IGMP interfaces; see 28.106 “To configure an IGMP interface” (p. 1019) .

h. MLD interfaces; see 28.123 “To configure an MLD interface on a base routing instance or VPRN routing instance” (p. 1035) .

5

Configure CPM virtual routing instances as required:

- Configure a CPM virtual routing instance; see 27.3 “To configure a CPM virtual routing instance ” (p. 837).
- Configure network interfaces on the CPM virtual routing instance as required; see 27.19 “To create a network interface on a CPM virtual routing instance” (p. 865) and 27.20 “To configure network interfaces on a CPM virtual routing instance” (p. 866).
- Configure static routes on the CPM virtual routing instance as required; see 27.13 “To configure a static route on a routing instance” (p. 852)

6

Cable the network ports on the devices to the network ports on other devices.

7

Configure the appropriate routing policies as required; see Chapter 54, “Routing policies” .

8

Create a network domain as required; see 27.21 “To create a network domain” (p. 867) .

a. Associate a network interface or service tunnel with a network domain; see 27.22 “To associate a network interface or service tunnel with a network domain” (p. 868) .

b. Delete a network interface or remove a service tunnel from a network domain; see 27.23 “To remove a network interface or service tunnel from a network domain” (p. 869)

9

View or list the following NE routing and forwarding related objects, as required:

- routing instances and child objects; see 27.24 “To list and view routing instances and child objects” (p. 869) .
- DHCP leases or prefixes assigned to routing instances; see 27.25 “To view and clear DHCP leases or prefixes assigned to a routing instance” (p. 870) .
- DHCPv6 log events; see 27.26 “To view DHCPv6 log events” (p. 872) .
- MVPN Extranet objects for a NE; see 27.28 “To list MVPN Extranet objects for a NE” (p. 873) .
- IOM/IMM label, next hop, and outgoing interface information for BGP, LDP and RSVP tunnels; see 27.29 “To display show router fp-tunnel information for a routing instance” (p. 874) .
- LDP session information; see 28.56 “To view the LDP session information” (p. 949) .
- LSN subscriber host statistics data; see 30.16 “To plot LSN subscriber host statistics” (p. 1097) .
- policy variable usage in a routing policy statement; see 54.20 “To view policy variable usage in a routing policy statement” (p. 1766) .
- routing policy usage; see 54.21 “To view routing policy usage” (p. 1767) .
- routing policies from a CLI session; see 54.22 “To show a routing policy CLI configuration in the client GUI” (p. 1768) .

27.2 To configure a routing instance or a VRF instance

27.2.1 Purpose

This procedure describes the base configuration of a routing instance or VRF instance.

Note: The functionalities described in this procedure are device-dependent. Not all functionalities described in this procedure are supported by all devices.

27.2.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Perform one of the following:

a. If you are configuring a VRF instance on the following OmniSwitch NEs, go to Step 3 .

- OS 6850E
- OS 6855-U24X
- OS 6860
- OS 6860E
- OS 6860N
- OS 6865
- OS 6900
- OS 9700E
- OS 9800E
- OS 10K

Note: SNMPv3 must be used to manage supported OmniSwitch devices when configuring a VRF routing instance.

b. If you are configuring a routing instance on all other devices, go to Step 5 .

3

Right-click on an OmniSwitch routing instance icon and choose Create VRF Instance. The Routing Instance (Create) form opens.

4

Configure the VRF Name parameter and click Apply. The Routing Instance (Edit) form opens. Go to Step 6 .

5

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

6

Edit or view the base routing policies associated with the routing instance if required:

Perform the following steps:

1. Click Edit Routing Policies. The Routing Policy Manager - Site (Edit) form opens.
2. Configure the parameters if required.
3. Click Show Policy. A Routing Policy Show Policy form opens and a CLI session is initiated.

   The Show Policy button only displays if a routing policy has been distributed to the device.
4. View the policy as required.
5. Close the form. The Routing Policy Manager - Site (Edit) form reappears.
6. Click OK if required or close the form. The Routing Instance (Edit) form reappears.

7

Configure the parameters on the General tab as required.

1. Select a reassembly group, if required.
2. If the routing instance is part of a TCP MSS adjustment configuration, select an ISA-BB group and configure the Segment Size parameter.
3. If the routing instance is part of a home LAN extension configuration, place a check mark in the Enable Home LAN Extension check box.

4. If the routing instance is part of an ISA service chaining configuration, place a check mark in the Enable ISA Service Chaining check box.

8

If GRT route leaking is required for VPRN, configure the parameters on the Leak Export Policies tab.

9

Click on the Protocols tab or Multicast tab to modify the default routing or multicast protocols assigned to the routing instance.

1. Configure the parameters as required. See Chapter 28, “Routing protocol configuration” for more information about routing protocol configuration using the GUI.
2. Click Apply to save the changes.

10

Click on the OAM tab and perform the following as required:

a. Assign a test generation option to the routing instance by configuring the Test Generation Options parameters on the Configuration tab as required.

b. Configure a TWAMP Light Reflector on the routing instance:

Note: A TWAMP Light reflector is required to conduct a TWAMP Light Test Session. See 92.11 “To configure a TWAMP Light reflector” (p. 3153) for more information about TWAMP Light reflectors.
The 7705 SAR does not initiate TWAMP Light Test sessions, but only acts as a TWAMP Light Reflector. The 7705 SAR does not support PM sessions.
The 7705 SAR-Hm supports TWAMP Light Test sessions and PM sessions.

Perform the following steps:

1. Click the TWAMP Light Reflector tab and click Create or select an existing TWAMP Light Reflector and click Properties.The TWAMP Light Reflector (Create|Edit) form opens.
2. Configure the parameters on the General tab as required.
3. Click Create in the Reflector Prefixes panel. The Prefix TWAMP Light Reflector (New Instance) form opens.
4. Configure the parameters as required and click OK. The TWAMP Light Reflector (Create) form reappears.

   Prefixes are added to the reflector to determine which PM Sessions can target the reflector. You should specify the prefix address and length if you require masking. Only those TWAMP Light test sessions with a Source IP matching a prefix will be valid. A maximum of 50 prefixes can be added to a reflector. TWAMP Light sessions without a matching prefix will cause an alarm to be generated.
5. Close the form. The Routing Instance (Edit) form reappears.

c. Add a TWAMP Light Session to the routing instance:

Perform the following steps:

1. Click on the TWAMP Light Sessions tab and click Add.The TWAMP Light Session (Create) form opens.
2. Perform 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) to complete the configuration of the TWAMP-Light Session.

d. Add a PM Session to the routing instance:

Perform the following steps:

1. Click on the PM Sessions tab and click Add. The PM Session, New Create form opens.
2. Perform 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) to complete the configuration of the PM Session.

11

Click on the Routing tab and perform the following as required:

Perform the following steps:

1. Configure the required general parameters.

   For 7210 SAS NEs, the value for the Maximum Number of Equal Cost Routes parameter cannot be greater than the value of the Max ECMP Route Destinations (Active) parameter in the system resource profile. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

   For 7250 IXR NEs, you must configure the system resource profile appropriately. Create ECMP Profiles to provide both an LDP and an IP profile type; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380).

   See Chapter 28, “Routing protocol configuration” for more information about AS and confederation AS parameters for BGP using the GUI.
2. Configure the Enable Entropy Label parameter, as required.
3. Configure the required parameters in the LDP Shortcut Enabled panel.
4. Configure the required parameters in the Maximum Number of Multicast Routes panel.
5. Assign an ingress multicast information policy to the routing instance in the Multicast Path Management Info Policy panel.

   See 52.13 “To configure an ingress multicast information policy” (p. 1720) for information about how to configure an ingress multicast information policy. See Step 37 in this procedure to view the data from the operational channels after actual traffic passes through the virtual router/routing instance from a specific multicast source for a specific multicast group.
6. Configure the required parameters in the TTL Propagate panel.
7. Configure the required parameters in the LSP BFD panel.
8. Configure the Single SFM Overload Admin State parameter in the Single SFM Overload panel.
9. Configure the parameters in the Static Route Hold-Down Time panel.

The Initial Hold-Down Time, Hold-Down Time Multiplier, and Maximum Hold-Down Time parameters are configurable when you enable the Enable Hold-Down Time parameter.

10. Configure the IPv4 Max Size and IPv6 Max Size parameters in the Flowspec panel, as required.

12

Click on the WLAN GW tab as required to view or configure WLAN gateway functionality on 7750 SR devices routing instances.

Note: Prior to completing this step, verify that the WLAN GW parameter on the General tab on the Routing Instance (Create|Edit) form is enabled.

Perform the following steps:

1. Configure the parameters on the General tab as required.

   The Interim Update parameter must be enabled if the Include Counters and Hold Down Time parameters are to be configured.

2. Click on the Mobile GW Address Map tab and click Create or select an existing mobile gateway address map. The Mobile Gateway Address Map (Create|Edit) form opens.
3. Configure the parameters as required.
4. Click on the Select button to specify a mobile gateway/peer profile policy.
5. Click Create or choose an existing mobile gateway/peer profile and click Properties. The Mobile Gateway/Peer Profile (Create|Edit) form opens. See 64.20 “To configure a mobile gateway/peer profile” (p. 1862) for information about how to create a Mobile Gateway/Peer Profile. See Step 26 in this procedure to view WLAN GW Tunnel information from the operational channels after actual traffic passes through the routing instance.
6. If the WLAN GW is intended for a home LAN extension configuration, click on the Cross Connect tab.

   Configure the required parameters and select a WLAN GW group.

7. Save your changes and close the form. The Routing Instance (Edit) form reappears.

13

To configure home LAN extension functionality on the routing instance, click on the Home LAN Extension tab.

Select an ISA WLAN GW group and configure the other parameters.

14

To configure home ISA service chaining functionality on the routing instance, click on the ISA Service Chaining tab.

1. On the ISA Groups tab click Create or choose an existing ISA group and click Properties. The Service Chaining ISA Group (Create|Edit) form opens.
2. Select an ISA-NAT or ISA-WLAN GW group and click OK to close the form.
3. Click on the VTEP tab and configure the VXLAN VTEP Range, Start and End parameters.

15

Click on the Interfaces tab if required, to view or configure L3 interfaces assigned to the routing instance. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) to complete the configuration of L3 interface.

16

Click on the Address tab if required, to configure an IP address for the L3 interfaces.

Perform the following steps:

1. Choose an address in the list and click Click Properties. The IP Address form opens.
2. Configure the parameters as required.
3. Save your changes and close the form. The Routing Instance form reappears.

   Note:

   The IP Address, Prefix Length, and Broadcast Address Format cannot be changed for the System interface.

17

Click on the Static Routes tab if required, to view or configure static routes assigned to the routing instance. See 27.13 “To configure a static route on a routing instance” (p. 852) to complete the configuration of the static route.

18

Click on the Static Route Prefixes tab if required, to view or configure static route prefixes.

Perform the following steps:

1. Click Create. The Static Route Prefix form opens.
2. Configure the parameters.

   Routing instances that support BGP large communities can configure up to 12 communities in a static route prefix.
3. Save your changes and close the form. The Routing Instance form reappears.

19

Click on the NAT tab if required, to view or configure NAT configurations assigned to the routing instance. See 30.9 “To configure NAT on a routing instance” (p. 1086) to complete the NAT configuration.

20

Click on the RSVP Tunnel Interfaces tab if required, to view or configure tunnel interface entries assigned to the routing instance.

Perform the following steps:

1. Click Create to configure a new entry or select an existing RSVP Interface and click Properties. The RSVP Tunnel Interface (Create|Edit) form opens.

Note:

You can alternatively create a RSVP tunnel interface by right-clicking on a routing instance icon in the navigation tree routing view, and choose Create RSVP Tunnel Interface from the menu.

2. Click Select to choose a Point-to-Multipoint LSP. The Select - RSVP Tunnel Interface form opens. See 31.18 “To create a Point-to-Multipoint LSP” (p. 1137) for information about how to configure a Point-to-Multipoint LSP.
3. Choose an RSVP tunnel interface and click OK. The RSVP Tunnel Interface (Create) form reappears.
4. Configure all other parameters as required.
5. Save your changes and close the form. The Routing Instance form reappears.

21

Click on the LDP Tunnel Interfaces tab if required, to view or configure LDP Tunnel Interfaces assigned to the routing instance. Perform the following as required:

a. Configure a new interface.

Perform the following steps:

1. Click Create or select an existing LDP tunnel interface and click Properties. The LDP Tunnel Interface (Create|Edit) form opens.

   Note:

   You can alternatively create a LDP tunnel interface by right-clicking on a routing instance icon in the navigation tree routing view, and choose Create LDP Tunnel Interface from the menu.

   LDP tunnel interfaces can be created using the system address or any local interface address on the root node.

   An LDP tunnel can be created on intermediate nodes (that is, between direct and indirect nodes).

2. Configure the parameters as required.

   Note:

   When the Root Node parameter is enabled, the LDP tunnel interface you are creating is associated with a root node, and you must specify a unique P2MP ID and Sender Address.

   When the Root Node parameter is not enabled, the LDP tunnel interface you are creating is associated with a leaf node, and you must specify the same P2MP ID and Sender Address as the root node.

   The 7450 ESS only supports LDP tunnel interfaces in Mixed Mode.

   You must enable the OAM Test Root Node parameter to add an interface to an OAM test suite or to create a P2MP LSP Ping test.

3. Click Apply. The LDP Tunnel Interface (Create) form is refreshed.
4. Save and close the form or proceed to the next sub-step if required.

b. Add the interface to an OAM test suite if required.

Perform the following steps:

1. Click on the OAM tab and click Add on the Test Suite tab to assign the interface to an OAM test suite. The Select Test Suite form opens.
2. Click Search and choose the appropriate test suite. See 89.12 “To create an STM test suite” (p. 2951) for information about how to configure an OAM test suite.
3. Save the form or proceed to the next sub-step if required.

c. Create a P2MP LSP Ping test if required.

Perform the following steps:

1. Click on the P2MP LSP Ping tab and click Create or select an existing P2MP LSP Ping test and click Properties. The P2MP LSP Ping form is displayed. See 90.28 “To create and run a MPLS P2MP LSP ping OAM diagnostic test from the STM” (p. 3033) to complete the P2MP LSP Ping test configuration.
2. Save and close the form or proceed to the next sub-step if required.

d. Click on the Test Entity Results tab and click Search to view the results of the OAM test suite or P2MP LSP Ping test execution if required.

22

Click on the Multi-Homing Interface tab if required, to view or create multi-homing interface entries assigned to the routing instance. The Multi-Homing Interface is a loopback interface used in multi-homing resiliency for a pair of protected routers. When active, the Primary interface can be used to advertise reachability information of the alternate router to the rest of the network. The Secondary interface is used to resolve routes advertised by the alternate router in the event that router becomes unavailable. This mechanism applies to both IP and VPN traffic.

Note: You can alternatively create a multi-homing interface by right-clicking on a routing instance icon in the navigation tree routing view and choose Create Multi-homing Interface from the menu.
Multi-homing interfaces are not supported on 7705 SAR-Hm.
Only one Primary and one Secondary multi-homing interface can be created for a router.

You can configure a multi-homing interface on the 7950 XRS, 7750 SR, and 7450 ESS (in mixed mode) network elements. Chassis mode “D” must be enabled.

1. Click Create or select an existing Mult-Homing Interface and click Properties. The Mult-Homing Interface, (Create|Edit) form opens.
2. Configure the parameters as required.
3. Click on the Addresses tab and Click Create or select an existing IP address and click Properties. The IP Address, Routing Instance (Create|Edit) form opens.
4. Configure the IP Address parameter.

   Note:

   Only IPv4 addresses are supported when creating the multi-homing interface.
5. Save and close the forms. The Routing Instance form reappears.

23

Click on the Route Aggregation tab if required, to view or configure route aggregates that can be generated into the virtual router on the routing instance.

Perform the following steps:

1. Click Create or select an existing route aggregation and click Properties. The Aggregation (Create|Edit) form opens.
2. Configure the parameters as required.
3. Save your changes and close the form.

24

Click on the BGP Confederations tab, if required to view or create BGP confederations on the routing instance. See 28.30 “To configure a BGP confederation” (p. 916) to complete the BGP confederation configuration.

25

Click on the GTP tab to configure S11 and uplink interfaces on the routing instance; see 27.14 “To configure GTP on a routing instance” (p. 853).

26

Click on the WLAN GW Tunnels tab, if required to display WLAN GW Tunnel information from the operational channels after actual traffic passes through the routing instance. The tunnel is dynamically created by the NE to carry traffic from the access point to the WLAN-GW and terminates on the WLAN-GW ISA. The tunnel can be a GRE or VLAN type.

Perform the following steps:

1. Click Resync WLAN GW Tunnels to retrieve the up-to-date tunnel information from the NE. A warning message displays.
2. Acknowledge the warning message as required.
3. Click Search to refresh the data.

27

Click on the Source Address tab, if required to view or configure an override source IP address or L3 interface for use by a selected IP application on the routing instance.

Perform the following steps:

1. Click Create or select an existing Source Address and click Properties. The Source Address, (Create|Edit) form opens.
2. Configure the parameters as required.

   Note:

   If you choose the Interface Index option for the Source Address Termination parameter, the router must have a L3 network interface or an IES L3 access interface previously created on the routing instance to proceed. Click Select to choose an interface. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information about how to

configure a L3 network interface; see 78.28 “To configure an L3 access interface on an IES site” (p. 2472) for information about how to configure an IES L3 access interface.

If you choose the IP Address option for the Source Address Termination parameter, you must enter a IPv6 address for the Source IP Address.

3. Save your changes and close the form. The Routing Instance form reappears.

28

Click on the Advertisement tab to configure RDNSS advertisement options, if required.

Perform the following steps:

1. Click Create or select an existing entry and click Properties. The DNS Options form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

29

Click on the Local DHCP server tab if required, to view a list of DHCPv4 or DHCPv6 servers assigned to the routing instance or click Create to configure a new server. See 27.5 “To configure a local DHCPv4 server on a routing instance” (p. 839) or 27.6 “To configure a local DHCPv6 server on a routing instance” (p. 842) to complete the configuration.

30

Click on the RADIUS tab and the appropriate sub-tab as required, to view or create a new RADIUS server or a RADIUS proxy server assigned to the routing instance. See 27.8 “To configure a RADIUS server on a routing instance” (p. 846) or 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847) to complete the configuration.

31

Click on the PCP Servers tab if required, to view or create a PCP server assigned to the routing instance. See 67.4 “To configure a NAT PCP server on a base routing instance” (p. 1886) to complete the configuration.

32

Click on the PCEP PCC tab if required, to view or configure PCEP PCC attributes. See 27.10 “To configure a PCEP PCC” (p. 849) to complete the configuration.

33

Click on the SR Prefix SID tab if required, to configure a node SID.

Perform the following steps:

1. Click Create or select an prefix SID and click Properties. The Segment Routing Prefix SID (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

34

Click on the Seamless BFD tab if required, to view or configure Seamless BFD peers. See 28.26 “To configure Seamless BFD ” (p. 912) to complete the configuration.

35

Click on the IPSec Security Policies tab if required, to view or create VSR or 7705 SAR-Hm IPSec Security Policies.

Perform the following steps:

1. Click Create or select an existing IPSec Security policy and click Properties. The IPSec Security policy (Create|Edit) form opens.
2. Configure the required parameters.
3. Click on the IPsec Security Policy Entries tab, and add one or more security policy entries.

   Use the following steps:

   a. Click Create. The Security Policy Entry (Create) form opens.
   b. Configure the Entry ID parameter.
   c. Configure the parameters in the Local IP Address panel.
   d. Configure the parameters in the Local IPv6 Address panel.
   e. Configure the parameters in the Remote IP Address panel.
   f. Configure the parameters in the Remote IPv6 Address panel.
   g. Save your changes and close the form.

36

Click on the Origin Validation tab if required, to view or create a new RPKI session for BGP SIDR prefix origin validation assigned to the routing instance. See 28.34 “To configure BGP SIDR prefix origin validation” (p. 930) for additional information about configuring BGP SIDR prefix origin validation.

Perform the following steps:

1. Click Create or select an existing RPKI Session and click Properties. The RPKI Session, Site (Create|Edit) form opens.
2. Configure the parameters as required.
3. Click on the Static Routes tab to assign a cache server IP address and click Create. The Origin Validation Static Route Entry form opens

   Note:

   A static VRP entry is only configurable on a base router instance.
4. Configure the parameters as required.
5. Save your changes and close the form. The Routing Instance form reappears.

37

Click on the Mcast Path Mgmt Channel tab if required, to display the data from the operational channels after actual traffic passes through the virtual router/routing instance from a specific

multicast source for a specific multicast group. You must click Search to refresh the data. See 15.82 “To view the operational multicast channel properties of an MDA” (p. 541) for additional information about how to view the operational multicast channel properties of an MDA.

38

Click on the Self Generated Traffic tab if required, to configure QoS for SGT. See 27.15 “To configure QoS for self-generated traffic on a routing instance” (p. 854).

39

Click on the Statistics tab if required, to collect and view performance statistical information about the routing instance for example, routing, DHCPv6 drop, or virtual router ICMPv6 in/out statistics. See the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting statistics.

40

Click on the Deployment tab if required, to monitor the deployment status of the routing instance. See 11.5 “To monitor the deployment status of a network object” (p. 336) for information about monitoring the deployment status of a network object.

41

Click on the Faults tab if required, to view the routing instance alarm information such as the object alarms, affecting alarms, aggregated alarms, and related object alarms.

42

Click on the Delay Metric tab if required, to associate a link measurement template.

43

Save your changes and close the form.

END OF STEPS

27.3 To configure a CPM virtual routing instance

27.3.1 Purpose

This procedure describes the configuration of a virtual routing instance on a 7450 ESS, 7750 SR, 7850 VSA, or 7950 XRS NE.

27.3.2 Steps

1

In the navigation tree routing view, expand Network→*NE*.

2

Right-click on the NE and choose Create CPM Virtual Routing Instance. The Routing Instance (Create) form opens.

3

Configure the VRF/Routing Instance Name parameter and click Apply. The Routing Instance (Edit) form opens.

4

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

5

Configure the parameters as required.

6

Click on the Interfaces tab if required, to view or configure network interfaces assigned to the routing instance. See 27.19 “To create a network interface on a CPM virtual routing instance” (p. 865) to complete the configuration of the interface.

7

Click on the Address tab if required, to modify the IP address for the network interfaces.

8

Click on the Static Routes tab if required, to view or configure static routes assigned to the routing instance. See 27.13 “To configure a static route on a routing instance” (p. 852) to complete the configuration of the static route.

9

Click on the Deployment tab if required, to monitor the deployment status of the routing instance. See 11.5 “To monitor the deployment status of a network object” (p. 336) for information about monitoring the deployment status of a network object.

10

Click on the Faults tab if required, to view the routing instance alarm information such as the object alarms, affecting alarms, aggregated alarms, and related object alarms.

11

Save your changes and close the form.

END OF STEPS

27.4 To configure a cellular interface on a 7705 SAR-Hm

27.4.1 Steps

1

In the navigation tree routing view, expand Network→7705 SAR-Hm→Routing Instance.

2

Right-click on the routing instance and choose Create Interface. The Create Network Interface step form opens.

3

Configure the required general properties. Select PDN for the Domain parameter.

4

Configure the remaining parameters as required.

Choose the cellular port in the Select Port step.

5

Save your changes and close the form.

END OF STEPS

27.5 To configure a local DHCPv4 server on a routing instance

27.5.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Local DHCP Servers tab and click Create or select an existing DHCPv4 server and click Properties. The Local DHCP Server (Create|Edit) form opens.

4

Configure the parameters on the General tab as required.

5

Click Select in the Local User Database panel if required to choose a local database to provided local authentication. The Local User Db - Local DHCP Server, Routing Instance form opens.

Perform the following steps:

1. Click Search to list the available databases. See 74.9 “To configure a local user database for subscriber host authentication” (p. 2025) for information about configuring a local user database.
2. Select a database and click OK. The Local DHCP Server (Create|Edit) form reappears.

6

Click on the Failover tab to configure DHCP server failover support.

Perform the following steps:

1. Configure the parameters on the General tab as required.

   The Administrative State parameter must be configured to Down if you are configuring a failover DHCP server peer.
2. Click on the Peer tab to configure a failover DHCP server peer.
3. Click Create or select an existing failover peer and click Properties. The Failover Peer (Create|Edit) form opens.
4. Click Select to choose a MC peer group and click OK. See 40.4 “To configure an MC peer group” (p. 1330) for information on configuring an MC peer group.
5. Configure the Sync Tag parameter.The Sync Tag parameter setting must be the same on both failover servers.
6. Click Node Redundancy to configure node redundancy for the failover peer if required. The Manage Node Redundancy form opens, preconfigured for MC peer groups. The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add failover peer sync tags.
7. Click Create or select an existing MC peer group and click Properties. The MC Peer Group (Create|Edit) form opens. See 40.4 “To configure an MC peer group” (p. 1330) to complete the MC peer group configuration.
8. Save your changes and close the forms. The Local DHCP Server (Create|Edit) form reappears.

7

Click on the IP Address Pools tab to configure address pools for the local DHCPv4 server.

Perform the following steps:

1. Click Create or select an existing IP address pool and click Properties. The IP Address Pool (Create|Edit) form opens.
2. Configure the parameters on the General tab as required.

8

Click on the Failover tab and repeat Step 6 to configure DHCP server failover support.

9

Click on the Subnets tab and click Create to add a subnet to the IP address pool or select an existing subnet and click Properties. The Subnet (Create|Edit) form opens.

Perform the following steps:

1. Configure the parameters on the General tab as required.
2. Click on the Address Ranges tab and click Create or select an existing address range and click Properties. The Subnet Address Range (Create|Edit) form opens.
3. Configure the parameters as required. You must exclude static IP addresses from the subnet address range because static IP addresses are dedicated.
4. Save your changes and close form. The Subnet (Create|Edit) form reappears.

10

Click on the Options tab.

Perform the following steps:

1. Click Create or select an existing option entry and click Properties. The IP Address Pool Option (Create|Edit) form opens.
2. Configure the parameters as required.

   The Value parameter is configurable when the Type parameter is set to ASCII String or Hex String.

   The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
3. Save your changes and close the form. The IP Address Pool (Create|Edit) form reappears.

11

Click on the Free Addresses tab and click Search to observe the minimum number of free addresses in the pool. If the actual number of free addresses in the pool falls below this configured minimum, an alarm is generated. A value of 0 specifies No Minimum.

12

Click on the Declined Addresses tab and click Search to observe the number of declined addresses in the pool.

13

Save your changes and close the IP Address Pool form.

14

To add a sticky lease to an IP address pool, perform the following steps:

1. Click on an IP address pool entry and click Create Sticky Lease. The Sticky Lease Action (Create) form opens.
2. Configure the required parameters.

   The Host Name parameter value must be unique. The MAC Address can be duplicated across multiple sticky lease objects, provided that each has a different corresponding Circuit ID value. Similarly, the Circuit ID value can be duplicated across multiple sticky lease objects, provided that each has a different corresponding MAC Address.

15

Save your changes and close the forms.

END OF STEPS

27.6 To configure a local DHCPv6 server on a routing instance

Note: The tabs and parameters that are configurable vary depending on the NE.

27.6.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance form opens.

3

Click on the Local DHCP Servers tab, then the DHCPv6 tab.

4

Click Create or select an existing DHCPv6 server and click Properties. The Local DHCP Server (Create|Edit) form opens.

5

Configure the parameters on the General tab as required.

The ID Type and Value parameters are configurable only if the Server ID Type parameter is set to Enterprise.

6

Configure DHCP server failover support, if applicable.

Perform the following steps:

1. Click on the Failover tab and configure the parameters on the General tab as required.

   The Administrative State parameter must be configured to Down if you are configuring a failover DHCP server peer.
2. Click on the Peer tab to configure a failover DHCP server peer.
3. Click Create or select an existing failover peer and click Properties. The Failover Peer (Create|Edit) form opens.
4. Click Select to choose a MC peer group and click OK. See 40.4 “To configure an MC peer group” (p. 1330) for information on configuring an MC peer group.
5. Configure the Sync Tag parameter.The Sync Tag parameter setting must be the same on both failover servers.
6. Click Node Redundancy to configure node redundancy for the failover peer if required. The Manage Node Redundancy form opens, preconfigured for MC peer groups. The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add failover peer sync tags.
7. Click Create or select an existing MC peer group and click Properties. The MC Peer Group (Create|Edit) form opens. See 40.4 “To configure an MC peer group” (p. 1330) to complete the MC peer group configuration.
8. Save your changes and close the forms. The Local DHCP Server (Create|Edit) form reappears.

7

Click on the IPv6 Address Pools tab.

Perform the following steps:

1. Click Create or select an existing entry and click Properties. The IPv6 Address Pool (Create|Edit) form opens.
2. Configure the parameters on the General tab.

8

Click on the Excluded Prefixes tab, if applicable.

Perform the following steps:

1. Click Create or select an entry and click Properties. The Excluded Prefix Address Pool (Create|Edit) form opens.
2. Configure the parameters.

9

Click on the Failover tab and repeat Step 6 to configure DHCP server failover support, if applicable.

10

Click on the Prefixes tab to add prefixes to the IPv6 address pool, if applicable.

Perform the following steps:

1. Click Create or select an existing prefix entry and click Properties. The Prefix (Create|Edit) form opens.
2. Configure the parameters on the General tab as required.
3. Click on the Options tab.
4. Click Create or select an existing option entry and click Properties. The Prefix Option (Create|Edit) form opens.
5. Configure the required parameters.

   The Number parameter is configurable when the Option parameter is set to Custom Option.

   The Value parameter is configurable when the Type parameter is set to ASCII String, Hex String, or Domain.

   The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
6. Save your changes and close the form.
7. Click on the Minimum Free Thresholds tab, if applicable.
8. Click Create or select an existing minimum free threshold entry and click Properties. The Prefix Minimum Free Threshold (Create|Edit) form opens.
9. Configure the required parameters.

   The Minimum Threshold and Minimum Number parameters are mutually exclusive. They cannot both be configured at the same time.
10. Save your changes and close the forms.

11

On the IPv6 Address Pool form, click on the Options tab and click Create or select an existing option entry and click Properties. The IP Address Pool Option (Create|Edit) form opens.

Perform the following steps:

1. Configure the parameters as required.

   The Number parameter is configurable when the Option parameter is set to Custom Option.

   The Value parameter is configurable when the Type parameter is set to ASCII String, Hex String, or Domain.

   The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
2. Save your changes and close the form.

12

Click on the Minimum Free Thresholds tab, if applicable.

Perform the following steps:

1. Click Create or select an existing minimum free threshold entry and click Properties. The Address Pool Minimum Free Threshold (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the forms.

END OF STEPS

27.7 To perform a Force Partner Down action on a local DHCP server failover

27.7.1 Purpose

The Force Partner Down action attempts to force a transition of the local DHCP server failover facility from a noCommunication state to a partnerDown state.

27.7.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance form opens.

3

Click on the Local DHCP Servers tab, then Click on the DHCPv4 or DHCPv6 t

4

Click on the DHCPv4 or DHCPv6 tab.

5

Select the DHCP server on which you want to perform a Force Partner Down action and click Properties. The Local DHCP Server form opens.

6

Click on the Failover tab.

7

Click Force Partner Down in the Force Partner Down Action panel.

The fields in the Force Partner Down Action panel are updated to reflect the result of the Force Partner Down action.

8

Save your changes and close the forms.

END OF STEPS

27.8 To configure a RADIUS server on a routing instance

27.8.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the RADIUS tab.

4

Click Create or select an existing RADIUS server and click Properties. The RADIUS Server (Create|Edit) form opens.

5

Configure the parameters as required.

6

Click Select to assign a Python Policy if required. The Select Python Policy forms opens.

Perform the following steps:

1. Click Search to list the available policies or click Create. See 58.3 “To configure a Python policy” (p. 1808) to complete the Python policy configuration.
2. Select a policy and click OK. The RADIUS Server (Create|Edit) form reappears.

7

Click Select in the RADIUS Script Policy panel to assign a RADIUS script policy, if required. The Select Script Policy - RADIUS Server form opens.

Perform the following steps:

1. Click Search to list the available policies or Click. See 64.23 “To configure a RADIUS script policy” (p. 1865) to complete the RADIUS Script policy configuration.
2. Select a policy and click OK. The RADIUS Server (Create|Edit) form reappears.

8

Save your changes and close the forms.

END OF STEPS

27.9 To configure a RADIUS proxy server on a routing instance

27.9.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the RADIUS tab, then the RADIUS Proxy Servers tab.

4

Click Create or select an existing RADIUS proxy server and click Properties. The RADIUS Proxy Server (Create|Edit) form opens.

5

Configure the parameters as required.

6

Select a Python Policy, if required.

If you intend to configure a WLAN GW group with a distributed Python policy in Step 7, the Python policy you select in this step must be the same Python policy that is distributed to the WLAN GW group.

7

Click Select in the WLAN GW panel to assign a ISA-WLAN GW Group, if required.

8

Configure default RADIUS policies:

Perform the following steps:

1. Select a RADIUS authentication policy.
2. Select a RADIUS accounting policy.

   Note:

   If you configure a ISA-WLAN GW group in Step 7, the default RADIUS policies must be a ISA RADIUS policy.

9

Click on the Caching tab and configure the parameters as required.

10

Click on the Interfaces tab and click Create or select an existing network Interface and click Properties. The RADIUS Proxy Interface form opens.

Perform the following steps:

1. Click Select and choose an interface from the Select Interface form.
2. Select an interface and click OK. The RADIUS Proxy Server (Create|Edit) form reappears.

11

Click on the Attribute Matching/Users tab to configure RADIUS attribute matching type and matching entries for RADIUS policies.

Perform the following steps:

1. On the Type tab, configure the Type and Vendor ID parameters.
2. Click on the Entry tab.
3. Click Create to create a new RADIUS attribute matching entry, or select an existing entry in the list and click Properties. The Attribute Matching Entry form appears.
4. Configure the Entry ID, Prefix Match String, and Suffix Match String parameters as required.
5. Select a RADIUS authentication policy and a RADIUS accounting policy.
6. Save your changes and close the form.

12

Save your changes and close the forms.

END OF STEPS

27.10 To configure a PCEP PCC

27.10.1 Purpose

Use this procedure to create a PCC client and PCC Peer if needed. Each router acting as a PCC instantiates a PCEP session to the PCE in its domain.

PCEP session information and statistics are supported.

27.10.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the PCEP PCC tab and click Create.

4

Configure the parameters as needed.

5

To create a peer, click Create in the Peer panel and configure the parameters.

END OF STEPS

27.11 To configure a PCE Association

27.11.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the PCEP PCC tab and click Create. The PCEP-PCC (Edit) form opens.

4

Click on the PCE-Associations tab.

5

To configure diversity association:

1. Click on the Diversity sub-tab.
2. Click Create. The PCEP-PCC PCE-Association Diversity (Edit) form opens.
3. Configure the parameters and click OK.

6

To configure policy association:

1. Click on the Policy sub-tab.
2. Click Create. The PCEP-PCC PCE-Association Policy (Edit) form opens.
3. Configure the parameters and click OK.

7

Save your changes and close the forms.

END OF STEPS

27.12 To configure UDP relay, DHCP snooping, and DHCP Option 82 on OmniSwitch routing instances

Note: DHCPv6 snooping is supported only for standard VLAN on the OS 6450.

Note: The routing instance must be the default routing instance.
UDP Server Statistics/UDP Service Statistics will be lost as part of the parent change from Bridge Instance to Routing Instance.

27.12.1 Steps

1

In the navigation tree routing view, expand Network→OmniSwitch→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the UDP Relay tab to configure UDP relay if required.

4

Configure the required parameters in the UDP BOOTP/DHCP Relay General Configuration panel.

5

Configure the parameters in the Relay Information Configuration panel if you need to use the DHCP Option 82 functionality.

6

Configure the PXE Support parameter if you need support for PXE devices.

7

Click on the DHCP Snooping tab to configure DHCP snooping.

In the DHCPv6 Snooping Configuration panel, configure the DHCPv6 Snooping Mode parameter from the drop-down menu.

8

Configure the parameters as required.

9

Click on the Relay Services tab to configure UDP port relay services if required.

Perform the following steps:

1. Click Create.
2. Configure the UDP Relay Service parameter.

   If you selected the Other option, configure the Relay Service Port and Relay Service Description parameters.
3. Save your changes and close the form. The Routing Instance (Edit) form reappears.

10

Click on the Relay Destinations tab and click Create. The Create UDP Service Destinations form opens.

11

Click Select. The Select UDP Service - UDP Relay Service Destination form opens.

Perform the following steps:

1. Choose a UDP relay service and click OK.
2. Configure the Forwarding Address parameter, if required.
3. Click Next. The Select VLANs step is displayed.
4. Configure the search filter and click Search. The Select VLANs form displays the VLANs that match the search criteria.

5. Choose one or more VLANs to be used to forward the UDP traffic.
6. Click Finish. The selected VLANs appear in the relay destinations list and the Routing Instance (Edit) form reappears.

12

Save your and close the form.

END OF STEPS

27.13 To configure a static route on a routing instance

27.13.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Static Routes tab and click Create or select an existing static route and click Properties. The Static Route (Create|Edit) form opens.

Note: You can alternatively create a static route by right-clicking on a static route icon in the navigation tree routing view, and choose Create Static Route from the menu.

4

Configure the required general parameters.

5

Configure the required parameters in the Destination panel.

To create an indirect static route, see 27.16 “To configure LSP entries with indirect static routes” (p. 855) .

6

To assign a network interface for the static route if IPv6 is enabled on the routing instance, select an interface in the Destination panel. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for information about how to configure a network Interface.

7

Configure the required parameters in the Other panel.

8

Configure the required parameters in the CPE Check panel.

9

To assign a prefix list policy to the static route, select a prefix list policy in the Prefix panel. See 54.7 “To configure a prefix list policy” (p. 1752) for information about how to configure a Prefix List policy.

You cannot specify a Prefix List policy if the BFD Enabled or Enable CPE Check parameters are enabled for the static route.

10

Configure the parameters in the Source/Destination Class Index panel.

11

Save your changes and close the form.

END OF STEPS

27.14 To configure GTP on a routing instance

27.14.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

On the General tab, set the Enable GTP Uplink parameter to Enabled.

4

Click on the GTP tab to configure S11 and uplink interfaces.

5

Click on the On the S11 tab to configure S11 interfaces and peer profile mappings:

1. On the Interfaces tab, click Create. The S11 Interface (Create) form opens.
2. Select an interface and an APN policy.
3. Save the changes and close the form.
4. Click on the Peer Profile Mapping tab and click Create. The S11 Peer Profile Mapping (Create) form opens.

5. Configure the Address Prefix and Prefix Length parameters.
6. Select a Peer Profile.
7. Save the changes and close the form.

6

Click on the On the Uplink tab to configure an uplink interface and peer profile mappings:

1. On the General tab, configure the APN Network Identifier and PDN Type parameters.
2. Click on the Peer Profile Mapping tab and click Create. The Uplink Peer Profile Mapping (Create) form opens.
3. Configure the Address Prefix and Prefix Length parameters.
4. Select a Peer Profile.
5. Save the changes and close the form.

7

Save your changes and close the form.

END OF STEPS

27.15 To configure QoS for self-generated traffic on a routing instance

27.15.1 Before you begin

NEs produce SGT for various applications, for example Telnet, SNMP, SSH. For each application, you can configure the DSCP or dot1p value for the traffic generated by that application. You can also map DSCP values to forwarding classes.

27.15.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Self Generated Traffic tab and perform the following as required:

a. Configure DSCP values for the required SGT applications.
   1. Click on the DSCP Marking tab and select an application in the list, then click Properties. The Application DSCP Marking form opens.
   2. Configure the DSCP parameter.
   3. Save your changes and close the form.

b. Map DSCP names to forwarding classes.
   1. Click on the DSCP Mapping tab and select a DSCP name in the list, then click Properties. The Application DCSP Mapping form opens.
   2. Configure the Forwarding Class parameter.
   3. Save your changes and close the form.

c. Configure dot1p values for the required SGT applications.
   1. Click on the Dot1p Marking tab and select an application in the list, then click Properties. The Application Dot1p Marking form opens.
   2. Configure the Dot1p parameter.
   3. Save your changes and close the form.

4

Save your changes and close the Routing Instance (Edit) form.

END OF STEPS

27.16 To configure LSP entries with indirect static routes

27.16.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Static Routes tab and select an existing static route and click Properties. The Static Route (Edit) form opens.

4

In the Destination panel, set the Type parameter to Indirect.

5

In the Destination panel, configure the IP Address parameter.

6

In the IGP Shortcut panel, set one of the following parameters to true.

- RSVP-TE
- SR-ISIS

- SR-TE
- SR-OSPF

The LSP Entries panel appears.

7

Click Create. The Indirect LSP Entry (Create) form opens.

8

Configure the LSP Name parameter or select an LSP Name that exists on the NE in the LSP panel.

9

Save your changes and close the forms.

END OF STEPS

27.17 To create an L3 network interface on a routing instance

27.17.1 Purpose

Perform this procedure to create L3 network interfaces on a routing instance. This procedure provides an example for configuring a network interface on a 7950 XRS with IPv6 enabled.

Note: The steps in this procedure depend on the device type, device version, and the parameter options that are selected during configuration. If a step in this procedure does not apply to the interface type that you are configuring on your device, proceed to the next applicable step.

An L3 interface on an OmniSwitch enables IP routing on a VLAN. Without an L3 interface, traffic is bridged within the VLAN or across connections to the same VLAN on other switches.

See 27.18 “To configure L3 network interfaces” (p. 863) to re-configure a L3 network interface after it is created. See 27.1.17 “Workflow to configure NE routing and forwarding” (p. 823) for information about how to configure other L3 network interfaces on routing instances.

You can create the following L3 interface types:

- ICMP interface
- system interface
- unnumbered interface
- management interface
- control tunnel
- IGMP interface
- PIM interface
- MLD interface

27.17.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Create Interface. The Create Network Interface step form opens.

3

Configure the parameters on the Define General Properties form as required.

When you create an unnumbered MPLS-TP instance, you must configure a Static ARP in Step 11 before you can turn the Administrative State parameter to Up.

4

Click Next. Perform one of the following as required:

a. If you chose the Numbered option for the Class parameter in Step 3 , the Configure IP Address form opens.

   On supporting NEs, you can create multiple IP addresses. For each address you need to create, perform the following:

   1. Click Create to associate an IP address for the interface. The IP Address, Routing Instance (Create) form opens,
   2. Configure the parameters as required.

      The Broadcast Address Format parameter only appears if the IP Address parameter is set to an IPv4 address.

      The parameters in the IPv6 panel only appear if the IP Address parameter is set to an IPv6 address.
   3. Select a Track SRRP Instance, if required.
   4. Save the changes and close the form.

b. If you chose the Unnumbered option for the Class parameter in Step 3 , the Configure Unnumbered Interface form opens. Configure the parameters as required.

5

Click Next to associate the L3 interface with a physical port. The Select Port form open.

Perform the following steps:

1. Select a port.
2. Select a LAG link mapping profile, if required to assign a specific LAG egress port to a SAP or interface, and control how egress traffic is handled if the specified port fails. See 13.21 “To create a LAG link mapping profile” (p. 443) for more information about creating a LAG link mapping profile.

6

Click Next to associate a network policy for the interface. The Select Network Policy form opens. Network policies are used to determine QoS settings based on the packet CoS bits on the ingress and egress of the network.

Note: The form name and policy selection options vary depending on the NE. See Chapter 50, “QoS policies” for information about network policies.

Perform the following steps:

1. Select a network policy in the Network Policy panel. See 50.41 “To configure a QoS network policy” (p. 1568) for information about configuring a network policy.

   Note:

   If the network policy you selected has Port Redirect Groups on its Ingress or Egress Forwarding Classes, you must specify both an Ingress or egress queue group template policy and Instance ID.

   If you select a network policy with a forwarding class mapped to a queue group queue ID, you must ensure that the mapping queue group queue ID is in the selected queue group template policy.

2. Select a queue group template policy in the Ingress Forwarding Plane Redirect and Egress Port Redirect panels. Queue group template policies allow SAP or IP interface forwarding classes to be redirected from the typical queue mapping to a shared queue. See 50.74 “To configure a queue group ingress template policy” (p. 1617) or 50.75 “To configure a queue group egress template policy” (p. 1619) for information about configuring an ingress/egress queue group template policies.

   Note:

   You must ensure that the port you selected has a network egress queue group with the same name as the queue group template policy created on it.

   Queue group template policies are not applicable to L3 interfaces associated with HSMDA ports.

3. Configure the Queue Group Instance ID parameter or click Select to associate a queue group instance for the interface.

If you are configuring an L3 access interface for 7210 SAS or 7250 IXR nodes, go to Step 7 . Otherwise, go to Step 8.

7

Click Next to configure accounting statistics collection for the interface. The Select Accounting Policy form opens. Select an accounting policy and configure the required parameters.

8

Click Next to assign ACL filters to the interface. The Select ACL Filters form opens. ACL filters are used to filter out IP traffic that matches user-defined criteria defined in ACL IP filter or ACL IPv6 policies.

Perform the following steps:

1. Select an ingress/egress ACL IP filter policy. See 51.5 “To configure an ACL IP filter policy” (p. 1671) for information about configuring an ACL IP filter policy.
2. Select an ingress/egress ACL IPv6 filter policy. See 51.6 “To configure an ACL IPv6 filter policy” (p. 1677) for information about configuring an ACL IPv6 filter policy

9

Click Next. The Configure ICMP form opens. Configure the ICMP parameters.

10

Click Next. If the IPv6 Allowed parameter is enabled on the Define General Properties form, the Configure ICMPv6 form opens. Configure the IPv6 parameters as required.

11

Click Next. The Configure ARP form opens.

Perform the following steps:

1. Configure the parameters as required.
2. Click Create to statically associate an IP or MAC address to the interface. The Static ARP (Create) form opens.
3. Configure the parameters as required.
4. Save your changes and close the form.

12

Click Next. The Configure Proxy ARP form opens. Configure the parameters as required.

Proxy ARP allows a device, such as a router, to answer ARP requests that are intended for another device. This allows a device to reach a remote subnet without configuring routes to the subnet or a default gateway device.

13

Click Next. The Configure Proxy ARP Policies form opens. Configure the parameters as required.

Note: You must not leave an empty policy parameter between two configured policy parameters. For example, do not configure the Policy 1 and Policy 3 parameters and leave the Policy 2 parameter unconfigured. The NFM-P reorders the policies and moves the policy entered for the Policy 3 parameter to the Policy 2 parameter.

14

Click Next. If the IPv6 Allowed parameter is enabled on the Define General Properties form, the Configure Neighbor Discovery form opens.

Perform the following steps:

1. Click Create to configure a neighbor. The Neighbor Discovery (Create) form opens.
2. Configure the parameters as required.
3. Save your changes and close the form.

15

Click Next. The Configure Proxy Neighbor Discovery form opens.

Perform the following steps:

1. Configure the parameters as required.
2. Click Select to associate a routing policy statement to the interface. You can associate up to five routing policy statements for each interface. See 54.5 “To configure a routing policy statement” (p. 1745) to create a routing policy statement.

16

Click Next. The Configure Secure Neighbor Discovery form opens. Configure the parameters as required.

17

Click Next. If NTP or SNTP is enabled on the NE, the Configure NTP form opens. Configure the Broadcast parameter.

Note: When the Broadcast parameter is enabled, a time protocol such as NTP or SNTP must be enabled and configured on the device, or the device ignores broadcast time packets received on the interface. See Chapter 8, “Device commissioning and management” for information about configuring NTP or SNTP on a device.

18

Click Next. The Configure DHCP - General form opens. To configure IPv4 DHCP for the interface:

Perform the following steps:

1. Configure the required parameters.

   The Remote ID String parameter is configurable when the Remote ID parameter is set to Remote IDString.
2. Select a Python policy, if required.
3. Click Next. The Configure DHCP - Server form opens.
4. Configure the Server 1 through Server 8 parameters.

19

Click Next. The Configure VRRP form opens. Click Create to add a VPPR instance. See 37.4 “To create and configure a VRRP instance” (p. 1283) to create a VRRP instance.

VRRP instances are members of a virtual router that provide backup if a router fails in a statically configured LAN. VRRP instances are created from network interfaces on a router, either as a master owner through which IP packets are routed before a failover, or as a backup non-owner that assumes the packet-forwarding (master) role after a failover. Before you create or add a VRRP instance to a virtual router, ensure that the master and backup interfaces have the same VRID and occupy the same subnet. See 37.4 “To create and configure a VRRP instance” (p. 1283) to create a VRRP instance.

20

Click Next. If the IPv6 Allowed parameter is enabled on the Define General Properties form, the Configure IPv6 VRRP form opens. Click Create to add a VPPR IPv6 instance. See 37.4 “To create and configure a VRRP instance” (p. 1283) to create a VRRP instance.

21

Click Next. The Configure Router Advertisement form opens.

Perform the following steps:

1. Click Create to add a router advertisement entry. The Router Advertisement (Create) form opens.
2. Configure the parameters on the General tab as required.

   If you are configuring the L3 interface for an IPv6 VRRP instance, then the Send Advertisement and Use Virtual MAC Address parameters must both be enabled.
3. Click on the DNS Options tab and click Create. The DNS Options form opens.
4. Configure the parameters as required.
5. Click on the Prefix tab and click Create. The Router Advertisement Prefix (Create) form opens.
6. Configure the parameters as required.

   Note:

   Each Lifetime (seconds) parameter is configurable when the associated No Expiry parameter is disabled.
7. Save the changes and close the forms.

22

Click Next. The Unicast RPF form opens. Configure the parameters as required.

The URPF Check State IPv6 and URPF Check Mode IPv6 parameters are only configurable when the IPv6 Allowed parameter is enabled is enabled on the Define General Properties form.

23

Click Next. The Configure DHCP - General form opens. Configure the parameters as required to configure IPv4 DHCP for the interface.

Note: The Remote ID String parameter is configurable when the Remote ID parameter is set to Remote IDString.

24

Click Next. The Configure DHCP - Server form opens. Configure the Server 1 through Server 8 parameters.

25

Click Next. The Configure Network Domain form opens. Click Add Network Domain to add a network domain to a network interface. See 27.21 “To create a network domain” (p. 867) to create a network domain.

26

Click Next. The Configure BFD form opens. To configure Bidirectional Forwarding Detection for the interface:

Perform the following steps:

1. Set the Administration Status parameter to Up.
2. Configure the parameters as required.
3. Click OK.

   Note:

   You cannot enable BFD on an interface, if BFD is not configured on the interface. You cannot set the administrative status of an interface to disabled, when protocols using the interface have BFD enabled. See Chapter 28, “Routing protocol configuration” for information about enabling and disabling BFD for routing protocols.

27

Click Next. If the IPv6 Allowed parameter is enabled on the Define General Properties form, the Configure IPv6 BFD form opens. Configure the parameters as required.

28

Click Next. The Select Security form opens. Click Select to assign a DoS protection policy or DDoS protection policy to a network interface. See 27.1.9 “DoS protection on L3 network interfaces” (p. 820) or 27.1.10 “DDoS protection on network and access interfaces” (p. 821) in this chapter for more information.

29

Click Next. The Configure Cflowd Sampling form opens.

Note: You must enable cflowd globally on an NE before you can configure cflowd collectors; see 12.10 “To enable and configure global Cflowd sampling on an NE” (p. 347)
.
You can configure cflowd sampling for unicast and multicast traffic on a routing network interface; up to two cflowd sampling objects can be configured under the parent interface, one for unicast and one for multicast.

Perform the following steps:

1. Click Create. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click OK.

30

Click Finish. The Summary form opens.

31

Enable the View the newly created Network Interface parameter to view the interface configuration after closing the form.

32

Save your changes and close the form.

END OF STEPS

27.18 To configure L3 network interfaces

27.18.1 Purpose

Perform this procedure to re-configure L3 network interfaces created in 27.17 “To create an L3 network interface on a routing instance” (p. 856) or to configure interfaces that are associated with a routing instance but not previously configured.

27.18.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→*Interface type*.

2

Right-click on the interface that you want to configure and choose Properties. The Network Interface (Edit) form opens.

3

Click on the appropriate tab and configure the parameters as required. Table 27-1, “Network interface form tabs” (p. 863) describes the tabs you can use to modify an existing routing interface.

Note: The network interface form tabs listed in Table 27-1, “Network interface form tabs” (p. 863) depend on the device type, device version, and the parameter options that were initially configured for the network interface. If a tab does not apply to the interface type that you are configuring on your device, proceed to the next applicable tab that applies.

Table 27-1 Network interface form tabs

Tab	Description
General	To configure the general parameters associated with the network interface.

Table 27-1 Network interface form tabs (continued)

Tab	Description
Port	To create an association between the network interface and a physical port or channel. The system interface is not associated with a port.
Policies	To assign QoS policies, ingress or egress IP ACL filter policies, and accounting template policies to the network interface.
Protocols	To view a list of protocols that are enabled on the interface, and to apply protocols.
Multicast	To view a list of multicast protocols that are enabled on the interface, and to apply protocols.
Security	To assign a DoS protection policy or DDoS protection policy to a network interface. See 27.1.9 “DoS protection on L3 network interfaces” (p. 820) or 27.1.10 “DDoS protection on network and access interfaces” (p. 821) in this chapter for more information.
BFD	To configure BFD parameters to a network interface.
Addresses	To assign IP addresses, and subnet and broadcast address formats to a network interface. To assign Track SRRP Instance to IP addresses, if required.
ARP	To configure or view ARP and proxy ARP parameters on the network interface.
ICMP	To configure ICMP parameters on the network interface.
ICMPv6	To configure ICMPv6 parameters on the network interface if IPv6 is enabled.
Neighbor Discovery	To configure neighbor discovery and proxy neighbor discovery for static routes on the network interface.
DHCP	To configure DHCP Relay parameters or select a Python policy.
Local DHCP	To assign local DHCP servers to the network interface. To associate a local DHCPv6 server, the IPv6 Allowed parameter must be enabled on the General tab.
DHCP Client	To enable DHCP and configure a DHCP client.
Administrative Groups	To assign administrative group policies to a network interface. See 54.13 “To configure an administrative group policy” (p. 1758) to create an administrative group policy.
Shared Risk Link Groups	To assign shared risk link group policies to a network interface. See 54.14 “To configure a Shared Risk Link Group policy” (p. 1759) to create a shared risk link group policy.
Network Domain	To add a network domain to a network interface. See 27.21 “To create a network domain” (p. 867) to create a network domain.
VRRP	To add a VRRP instance to an existing virtual router that assumes routing responsibilities from a failed router in a statically configured LAN. See 37.4 “To create and configure a VRRP instance” (p. 1283) to create a VRRP instance.
MEPs	To add network interface facility MEPs to the network interface. Each network interface can support one facility MEP.
Advertisement	To configure router advertisement on the network interface.
NTP	To enable SNTP broadcasts on the network interface.
RCA Result	To view the results of RCA activities on the network interface.
Zone	To assign a 7705 SAR security zone policy to a network interface. See 68.9 “To configure a security zone policy for a 7705 SAR” (p. 1894) to create a security zone policy.

Table 27-1 Network interface form tabs (continued)

Tab	Description
QoS	To assign an egress network queue policy to a network interface. See 50.46 “To configure a network queue policy” (p. 1581) or 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583) to create a network queue policy.
Deployment	To monitor the deployment status of the interface such as any warnings when a configuration change is not fully deployed to an NE.
Clear Statistics Status	To view the status of all clear requests on the network interface.
RP Statement Actions items	To view the routing policy statement associated with the network interface. See 54.5 “To configure a routing policy statement” (p. 1745) to create a routing policy statement.
Statistics	To view the statistics associated with the network interface.
TCA	To assign a TCA policy to a network interface. See 19.4 “To configure a TCA policy” (p. 709) to create a TCA policy.
Faults	To view alarms that are raised against the network interface, or related alarms that affect the network interface.

4

Save your changes and close the form.

END OF STEPS

27.19 To create a network interface on a CPM virtual routing instance

27.19.1 Purpose

Perform this procedure to create network interfaces on a CPM virtual routing instance. Up to two network interfaces can be created on the virtual routing instance with up to two ports each. A limited number of attributes are supported on the interface.

See 27.20 “To configure network interfaces on a CPM virtual routing instance” (p. 866) to re-configure a network interface after it is created. See 27.1.17 “Workflow to configure NE routing and forwarding” (p. 823) for information about how to configure other L3 network interfaces on routing instances.

27.19.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Create Interface. The Create Network Interface step form opens.

3

Configure the parameters on the Define General Properties form as required.

4

Click Next. The Configure IP Address form opens.

Perform the following steps:

1. Click Create to associate a IP address for the interface. The IP Address, Routing Instance (Create) form opens.
2. Configure the parameters as required.
3. Save the changes and close the form.

5

Click Next to associate the interface with a physical port. The Select Port form opens.

Select a port. Only A/4 or B/4 ports can be associated with the interface.

6

Save your changes and close the form.

END OF STEPS

27.20 To configure network interfaces on a CPM virtual routing instance

27.20.1 Purpose

Perform this procedure to re-configure network interfaces created in 27.19 “To create a network interface on a CPM virtual routing instance” (p. 865) or to configure interfaces that are associated with a CPM virtual routing instance but not previously configured.

27.20.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→*Interface*.

2

Right-click on the interface that you want to configure and choose Properties. The Network Interface (Edit) form opens.

3

Click on the appropriate tab and configure the parameters as required. Table 27-2, “CPM virtual routing instance network interface form tabs” (p. 867) describes the tabs you can use to modify an existing routing interface.

Table 27-2 CPM virtual routing instance network interface form tabs

Tab	Description
General	To configure the general parameters associated with the network interface.
Port	To create an association between the network interface and a physical port or channel. The system interface is not associated with a port.
Addresses	To assign an IP address and subnet to a network interface. Only one primary IP address can be associated with a network interface.
Statistics	To view the statistics associated with the network interface.
TCA	To assign a TCA policy to a network interface. See 19.4 “To configure a TCA policy” (p. 709) to create a TCA policy.
Deployment	To monitor the deployment status of the interface such as any warnings when a configuration change is not fully deployed to an NE.
Faults	To view alarms that are raised against the network interface, or related alarms that affect the network interface.

4

Save your changes and close the form.

END OF STEPS

27.21 To create a network domain

27.21.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→Network Domains.

2

Right-click on the network domains icon and choose Create Network Domain. The Network Domain (Create) form opens.

3

Configure the parameters as required.

4

Click OK. The Domain Network (Create) form closes and a network domain displays in the navigation tree routing view.

END OF STEPS

27.22 To associate a network interface or service tunnel with a network domain

27.22.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→Network Domains→domain-*user defined domain*.

2

Right-click on the domain-*user defined domain* icon and choose Properties. The Network Domain - domain-*user defined domain* (Edit) form opens.

3

Associate a network interface with a network domain.

Perform the following steps:

1. Click on the Network Interfaces tab and click on Add Network Interface. The Add Network Interface form opens
2. Click on Search to display all the Network Interfaces.
3. Choose a network interface and click OK. The selected network interface is displayed on the Network Domain form.

4

Associate a service tunnel with a network domain.

Perform the following steps:

1. Click on the Service Tunnels tab and click on Add Tunnel. The Add Tunnel form opens.
2. Click on Search to display all service tunnels.
3. Choose a service tunnel and click OK. The selected service tunnel is displayed on the Network Domain form.

5

Save your changes and close the form.

END OF STEPS

27.23 To remove a network interface or service tunnel from a network domain

27.23.1 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→Network Domains→domain-*user defined domain*.

2

Right-click on the domain-*user defined domain* icon and choose Properties. The Network Domain - domain-*user defined domain* (Edit) form opens.

3

Remove a network interface from a network domain.

Perform the following steps:

1. Click on the Network Interfaces tab and click Remove Network Interface. The Remove Network Interface form opens.
2. Click on Search to display all network interfaces associated with the network domain.
3. Choose a network interface and click OK. The selected network interface is remove from the network domain.

4

Remove a service tunnel from a network domain.

Perform the following steps:

1. Click on the Service Tunnel tab and click Remove Tunnel. The Remove Tunnel form opens.
2. Click on Search to display all service tunnels associated with the network domain.
3. Choose a service tunnel and click OK. The selected service tunnel is remove from the network domain.

5

Save the network domain and close all open forms.

END OF STEPS

27.24 To list and view routing instances and child objects

27.24.1 Purpose

Perform this procedure to list and view the properties associated with routing instances and child objects such as IP and source addresses, assigned interfaces, protocol and site assignments, statistics, and faults.

27.24.2 Steps

1

Choose Manage→Networking→Routing Instances from the NFM-P main menu. The Manage Routing Instances form opens.

2

Choose an object from the Select Object Type drop-down menu.

3

Configure the search fields as required, and click Search. The search results display showing the selected object general properties associated with the routing instances and child objects.

4

As required, choose an object and click Properties to get additional information on the selected object. The selected object properties form opens.

5

Select the appropriate tab on the form to view the current parameter settings, deployment information, or fault information.

6

Save your changes and close the form.

END OF STEPS

27.25 To view and clear DHCP leases or prefixes assigned to a routing instance

27.25.1 Purpose

Depending on the device being queried, you can view DHCPv4 or DHCPv6 leases or prefixes, or both types may be supported by the device.

27.25.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Local DHCP Servers tab, then the appropriate DHCPv4 or DHCPv6 sub-tab. The list of associate DHCP servers is displayed.

4

Choose a DHCP server and click Properties. The Local DHCP Server (Edit) form opens.

5

Perform one of the following:

a. View DHCPv4 leases or prefixes.

   Perform the following steps:

   1. Click on the Leases tabs.
   2. Click Search on the General tab to refresh the data.
   3. Choose a lease and click Properties. The Local DHCP Server Leases form opens.
   4. View the lease information on the form.

b. View DHCPv6 leases or prefixes.

   Perform the following steps:

   1. Click Show Leases. The Local DHCPv6 Server Show Leases form opens.
   2. To view all active leases on the server, click OK.
   3. To view the lease information for a specific prefix, enter an IP address in the Prefix parameter and click OK.
   4. To view the detailed lease information for a specific prefix, enter the IP address in the Prefix parameter, select the Detail check box, and click OK. The form displays the CLI details for the specific prefix.

c. Clear DHCPv4 leases.

   Perform the following steps:

   1. Click on the Leases tab.
   2. Click on the Regular tab to clear DHCP leases.
   3. Click Search to refresh the data.
   4. Select the DHCP lease you want to clear and click Clear DHCP Lease.

      If you want to clear all of the leases on the local DHCPv4 server, click Clear All DHCP Leases.
   5. Click on the Sticky tab to clear sticky leases.
   6. Select the sticky lease you want to clear and click Clear Sticky Lease.
   7. If you want to clear all sticky leases with a specific host name prefix, click Clear Sticky Lease With Host Name Prefix.

      Specify a host name prefix and click Ok.
   8. Click on the Clear Status tab to check the progress/success of the lease clear command.

6

Save your changes and close the form.

END OF STEPS

27.26 To view DHCPv6 log events

27.26.1 Purpose

This procedure describes how to view the following log events on a DHCPv6 server:

- host conflict log
- lease not owner log
- pool unknown log

27.26.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Local DHCP Servers tab, then the Local DHCPv6 tab. The list of associate DHCPv6 servers displays.

4

Choose a DHCPv6 server and click Properties. The Local DHCPv6 Server (Edit) form opens.

5

Click on the Logs tab, then the appropriate sub-tab for the type of log that you wish to display.

6

Configure the filter criteria and click Search. A list of event logs appears.

7

Choose a log entry and click Properties. The Log form opens.

8

View the log entry and close the form.

END OF STEPS

27.27 To configure DHCP clients on SAR devices

27.27.1 Overview

The SAR-Hm can be configured to support IP/MPLS services over WLAN. This requires that the SAR-Hm act as a WLAN station, and alternatively act as both an AP and station concurrently to create a WLAN topology as necessary.

You can configure DHCP clients on a network interface with WLAN or an Ethernet port.

The 7705 SAR-Hmc can be configured to support IP/MPLS services on Network Interface with Ethernet Port.

27.27.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance→Interface.

2

Right-click on the interface that you want to configure and choose Properties. The Network Interface (Edit) form opens.

3

Click on the DHCP Client tab.

4

Click Create, or select a DHCP client and click Properties. The DHCP Client form opens.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

27.28 To list MVPN Extranet objects for a NE

27.28.1 Purpose

Perform this procedure to display a list of MVPN extranet objects that exchange route information through a given NE. The Show MVPN-Extranet menu command issues a show router mvpn-extranet CLI command to the selected NE, and returns a list of subscriber VRF IP addresses. The command can be issued from a base routing instance, or from a VPRN site.

27.28.2 Steps

1

Perform one of the following:

a. In the navigation tree routing view, expand Network→*NE*→Routing Instance.

b. Run the Show MVPN-Extranet command from a VPRN site.

Perform the following steps:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Configure the filter criteria. A list of services appears at the bottom of the Manage Services form.
3. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
4. In the navigation tree, expand Network→*NE*→Routing Instance.

2

Right-click on the Routing Instance icon and choose Show MVPN-Extranet from the menu. The Show MVPN-Extranet form opens.

3

In the Parameters field, you can type filter arguments for the show router mvpn-extranet CLI command. For example:

- mvpn-extranet [detail] [srcmvpn <*source_service_id*>] [recvmvpn <*reciever_service_id*>]

where

[detail] provides a detailed route table in the command output.

[srcmvpn <*source_service_id*>] filters the output route table to a source VRF, based on the service ID of the source VRF. The variable range is 0 to 4294967295.

[recvmvpn <*reciever_service_id*>] filters the output route table to a destination VRF, based on the service ID of the destination VRF. The variable range is 0 to 4294967295.

4

Click OK. The Show MVPN-Extranet form displays table of MVPN extranet objects.

END OF STEPS

27.29 To display show router fp-tunnel information for a routing instance

27.29.1 Purpose

Perform this procedure to display show router fp-tunnel information for a routing instance including the IOM/IMM label, next hop, and outgoing interface information for BGP, LDP and RSVP tunnels.

27.29.2 Steps

1

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the Routing Instance icon and choose Show FP Tunnel Table from the menu. The FP RTR Show FP Tunnel Table form appears.

3

Configure the parameters as required.

4

Click OK. The fp rtr Show FP Tunnel table opens and a CLI session is initiated.The FP tunnel table information for a routing instance is displayed.

END OF STEPS

27.30 To configure a Multi-Chassis shunt interface on a base routing instance or VPRN routing instance

27.30.1 Steps

1

Perform one of the following:

1. Configure a Multi-Chassis shunt interface on a base routing instance.

   Perform the following steps:

   a. In the navigation tree routing view, expand Network→*NE*→Routing Instance.
   b. Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

2. Configure a Multi-Chassis shunt interface on a VPRN routing instance.

   Perform the following steps:

   a. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
   b. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
   c. On the VPRN Service tree, click on the site to which you need to configure a Multi-Chassis shunt interface.
   d. Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the IPSec tab, then on the MC Shunt Interface tab.

3

Create or select an entry and click Properties. The Multi-Chassis Shunt Interface (Create/Edit) form opens.

4

Configure the required parameters.

5

Save your changes and close the form.

END OF STEPS

27.31 To configure a Multi-Chassis shunting profile on a base routing instance or VPRN routing instance

27.31.1 Steps

1

Perform one of the following:

1. Configure a Multi-Chassis shunting profile on a base routing instance.

   Perform the following steps:

   a. In the navigation tree routing view, expand Network→*NE*→Routing Instance.
   b. Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

2. Configure a Multi-Chassis shunting profile on a VPRN routing instance.

   Perform the following steps:

   a. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
   b. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
   c. On the VPRN Service tree, click on the site to which you need to configure a Multi-Chassis shunt interface.
   d. Right-click on the routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the IPSec tab, then on the MC Shunting Profile tab.

3

Create or select an entry and click Properties. The Multi-Chassis Shunting Profile (Create/Edit) form opens.

4

Configure the required parameters.

5

Click on the MC Peer tab.

6

Create or select an entry and click Properties. The Multi-Chassis Shunting Profile Peer (Create/Edit) form opens.

7

Configure the required parameters.

8

Configure the Shunt Interface parameter.

9

Save your changes and close the forms.

END OF STEPS

28 Routing protocol configuration

Routing protocol configuration overview

28.1 Routing protocol configuration overview

28.1.1 Overview

A routing protocol specifies how routers communicate with each other and enables them to select routes between any two nodes in a network. Routing protocols typically determine the specific choice of a route based on the following metrics:

- number of network-layer devices along a path (hop count)
- bandwidth
- load
- delay
- MTU
- reachable addresses
- cost

Routing protocols store the results of these metrics in a routing table. Routing protocols use the information and path parameters to compile a network topology. The information about routes that is used to update routing tables can be modified using routing policies. Some devices also support a RIB API with read and write access to their routing tables.

Each router has a prior knowledge only of networks attached to it directly. A routing protocol shares this information first among immediate neighbors, and then throughout the network. This way, routers gain knowledge of the topology of the network.

You can use the network, routing, IS-IS, and OSPF views in the NFM-P GUI navigation tree to view and configure the device routing protocol support. The supported protocols are displayed in the navigation tree as child objects to the routing instance.

A device can support multiple routing protocols. The supported protocols are device dependent. The NFM-P only displays the protocols that can be configured on each device type. Supported routing protocols on the NFM-P include:

- 28.3 “IPv6” (p. 881)
- 28.5 “BGP” (p. 883) and 28.5.2 “MP-BGP” (p. 886)
- 28.7 “RIP and RIPng” (p. 888)
- 28.8 “LDP” (p. 888)
- 28.9 “IS-IS” (p. 891)
- 28.10 “OSPFv2 and OSPv3” (p. 893)
- 28.12 “RSVP” (p. 896) (enabled by default)
- 28.13 “L2TP” (p. 898) (enabled by default)
- 28.19 “WPP” (p. 907)
- MPLS
- MPLS-TP
- Multicast routing
  - 28.14 “PIM” (p. 901)
  - 28.15 “IGMP” (p. 905)
  - 28.16 “MSDP” (p. 905)
  - 28.17 “MLD” (p. 906)
- 28.18 “Bridging” (p. 907)
- 28.19 “WPP” (p. 907)
- 28.20 “BIER” (p. 908)
- 28.21 “IPSec” (p. 908)

Note: For an OmniSwitch, the NFM-P supports viewing the status of routing protocols that are enabled using the CLI.

L2TP is enabled by default on a base routing instance, but not on a VPRN routing instance.

Routing protocols can be configured to import or export routes from other routing protocols using routing policies. See 27.1.12 “NE routing policies” (p. 821) in 27.1 “NE routing and forwarding” (p. 817) .

28.2 Area-based routing considerations with protocol usage

28.2.1 Overview

When you configure a device, you configure the routing protocols used to determine how routers communicate with each other and the topology of packet handling between different devices in the network.

Networks can be grouped into areas for a number of reasons including scaling, administrative convenience, equipment compatibility, or business structure, For example, an area can be a collection of network segments within an AS that have been administratively assigned to the same group.

Area topology is concealed from the rest of the AS. This means a significant reduction in routing traffic. Routing in the AS takes place on three levels, depending on whether the source and destination of a packet reside in:

- the entire path is limited to a single area (known as intra-area routing)
- the path spans multiple areas all owned by the same carrier (known as inter-area routing)
- the path spans multiple carriers (known as inter-carrier routing)

In intra-area routing, the packet is routed based on information found within the area; no routing information from outside the area is used. This protects intra-area routing from the injection of bad routing information. Two devices, which are not area border routers, and belonging to the same area, have identical topological databases.

Devices that are aware of more than one area are called area border routers. In this case, all devices in an AS do not have an identical topological database. An area border router has a

separate topological database for each area it is connected to. ASs share routing information, such as routes to each destination and information about the route or AS path, with other ASs using BGP.

28.3 IPv6

28.3.1 Overview

The NFM-P supports IPv6 for control-plane addressing on IPv6-enabled devices.

Although the implementation of IPv6 is driven by the diminishing IPv4 address space, it is increasingly necessary to use a protocol that is designed to handle more complex network applications, such as broadband voice and video transmission, IP transit, Internet exchange peering, and other large enterprise applications.

28.3.2 Transition from IPv4 to IPv6

The transition from IPv4 to IPv6 is occurring in stages as network providers, service providers, and end users migrate existing applications and equipment to the new version. Control-plane forwarding of IPv6 packets is an important part of this transition; it allows isolated IPv6 hosts and smaller IPv6 networks to peer across an IPv4 network, using a mechanism such as 6over4 tunneling.

With 6over4 tunneling, a host encapsulates IPv6 packets in IPv4 packets for transport across an IPv4 network. Routers that identify 6over4 encapsulation remove the IPv4 encapsulation before they forward the packets to other native IPv6 hosts. The 6over4 mechanism uses the IPv4 multicast infrastructure for neighbor discovery.

28.3.3 IPv6 benefits

The general benefits of IPv6 include:

- **simplified header format and fixed header length**
  An IPv6 header contains fewer fields than an IPv4 header. IPv6 excludes obsolete IPv4 header fields and processes option fields only when they contain values. IPv6 also standardizes the size of the packet header to 40 bytes to streamline packet processing. These features reduce packet-processing overhead and make routing more efficient.
- **addressing enhancements**
  IPv6 increases the IP address size from 32 bits to 128 bits to support a greater number of NEs and to provide a more versatile addressing hierarchy. IPv6 also supports address auto-configuration.
  The scalability of multicast routing is improved by the presence of a Scope field. IPv6 introduces anycast addressing, which designates a group of disparate NEs as the recipient of a specific data stream.
- **improved scalability and extensibility**
  Built-in traffic optimization makes IPv6 highly scalable, and IPv6 supports future routing technology enhancements using protocol extensions.
- **flow-labeling capability**
  IPv6 features packet prioritization which allows the labeling of packets that require special handling, such as real-time service for VoIP conferences.

- **improved privacy and security**
  Authentication, encryption, and data integrity features are mandatory components for which IPv6 supports standardized extensions.

28.3.4 IPv6 support on the NFM-P

The NFM-P entities that support IPv6 configuration include the following:

- BGP for base routing instances and VPRN services
- MVPN IPv6 for BGP routing instances
- MVPN IPv6 for BGP policy statements
- IS-IS adjacencies
- multicast routing
- static routes
- ICMP
- routing policies
- access ingress and egress policies
- CPM filter policies
- PPP
- IES and VPRN SAPs
- IES and VPRN bidirectional forwarding detection
- VLL Ipipe
- PIM-SSM
- PIMv6 on VPRN routing instances and interfaces
- MLD under IES and VPRN

An NFM-P operator enables IPv6 on an interface during interface creation.

The benefits of the NFM-P IPv6 implementation include:

- integration with existing IPv4 configurations on many property and configuration forms
- automatic validation of IP addresses on the client GUI, regardless of the IP version
- support of compressed IPv6 addresses when repeated address octets are present
- support for IPv6 statistics
- simultaneous IPv4 and IPv6 support by routing protocols
- separate IPv4 and IPv6 administrative and operational states on an interface

28.3.5 Accepted IPv6 address formats

The NFM-P accepts IPv6 addresses in the following formats:

- colon-hexadecimal, or *x:x:x:x:x:x:x:x*
  where *x* is a 16-bit hexadecimal number from 0 to FFFF
- a combination of colon-hexadecimal and dotted-decimal, or *x:x:x:x:x:x:d.d.d.d*
  where
  x is a 16-bit hexadecimal number from 0 to FFFF
  d is an 8-bit decimal number from 0 to 255

Using a combination of colon-hexadecimal and dotted-decimal formats may be convenient in an environment that supports the use of IPv4 and IPv6 addresses.

The NFM-P allows IPv6 address compression for an address that contains repeated zero values. You can use two adjacent colons to represent any group of repeated zero values in an IPv6 address. For example:

2001:DB8::

expands to

2001:0DB8:0000:0000:0000:0000:0000:0000

It is not necessary to supply the leading zeros for a number in an IPv6 address. For example, 2001:DB8::3C:5 is a valid IPv6 address.

28.4 BFD

28.4.1 Overview

Bidirectional Forwarding Detection (BFD) is a communication error detection protocol. BFD can provide failure detection on any kind of bidirectional path or link between systems, including physical links, virtual circuits, LAGs, and MPLS LSPs. When BFD is in use, switches periodically exchange BFD Hello packets over each path between them. If a switch stops receiving BFD Hello packets after a configured interval, the neighboring switch, or the path hop, is assumed to have failed.

BFD global configuration is performed in the NE properties.

BFD sessions are enabled at the path or interface level, for example, in the properties of an LSP.

Seamless BFD

Seamless BFD (SBFD) offers a streamlined version of BFD. An SBFD module allocates SBFD session discriminators to local entities and distributes them to remote entities. The advertised discriminator information allows remote entities to more quickly initiate and perform BFD connectivity testing.

28.5 BGP

28.5.1 Overview

BGP is an inter-AS routing protocol. An AS is a network or a group of devices logically organized and controlled by a common network administration. BGP enables devices to exchange network reachability information. AS paths are the routes to each destination. There are two types of BGP: IBGP and EBGP.

- IBGP is used to communicate with peer devices in the same AS. Routes received from a device in the same AS are not advertised to other devices in the same AS but can be advertised to an EBGP peer.
- EBGP is used to communicate with peers in a different AS. Routes received from a device in a different AS can be advertised to both EBGP and IBGP peers.

See 79.1.6 “Inter-AS connections” (p. 2517) in “VPRN service management” (p. 2513) for more information about connecting ASs in a VPRN.

You can use the NFM-P to enable BGP on the device and perform the following functions:

- set the AS values for the routing instance
- create confederations of group-managed devices

- create BGP peer groups
- create neighbors (peers) within the BGP peer groups

The NFM-P supports the configuration of IPv6 addresses for BGP peering on the base routing instance of supporting devices.

The base BGP routing instance and the policy statement support MVPN IPv6 on supporting devices.

The NFM-P supports IPv4 and IPv6 routes with and without implementation labels. The base router maintains separate BGP RIBs for each of the following types of routes: IPv4, label-IPv4, IPv6 and label-IPv6. Each VPRN maintains separate BGP RIBs for each of the following types of routes: IPv4, label-IPv4 (if CSC is enabled), IPv6 and label-IPv6.

A device can only belong to one AS. After the neighbor relationship is established between devices, they exchange BGP open messages, which contain information such as AS numbers, BGP versions, router IDs, hold-time values, and keepalive messages. This information determines the status of the BGP session. Peer relationships are defined by configuring the IP addresses of the devices that are peers of the local BGP system.

The following figure shows a simple BGP example using groups, subconfederations, and confederations.

Figure 28-1 BGP example

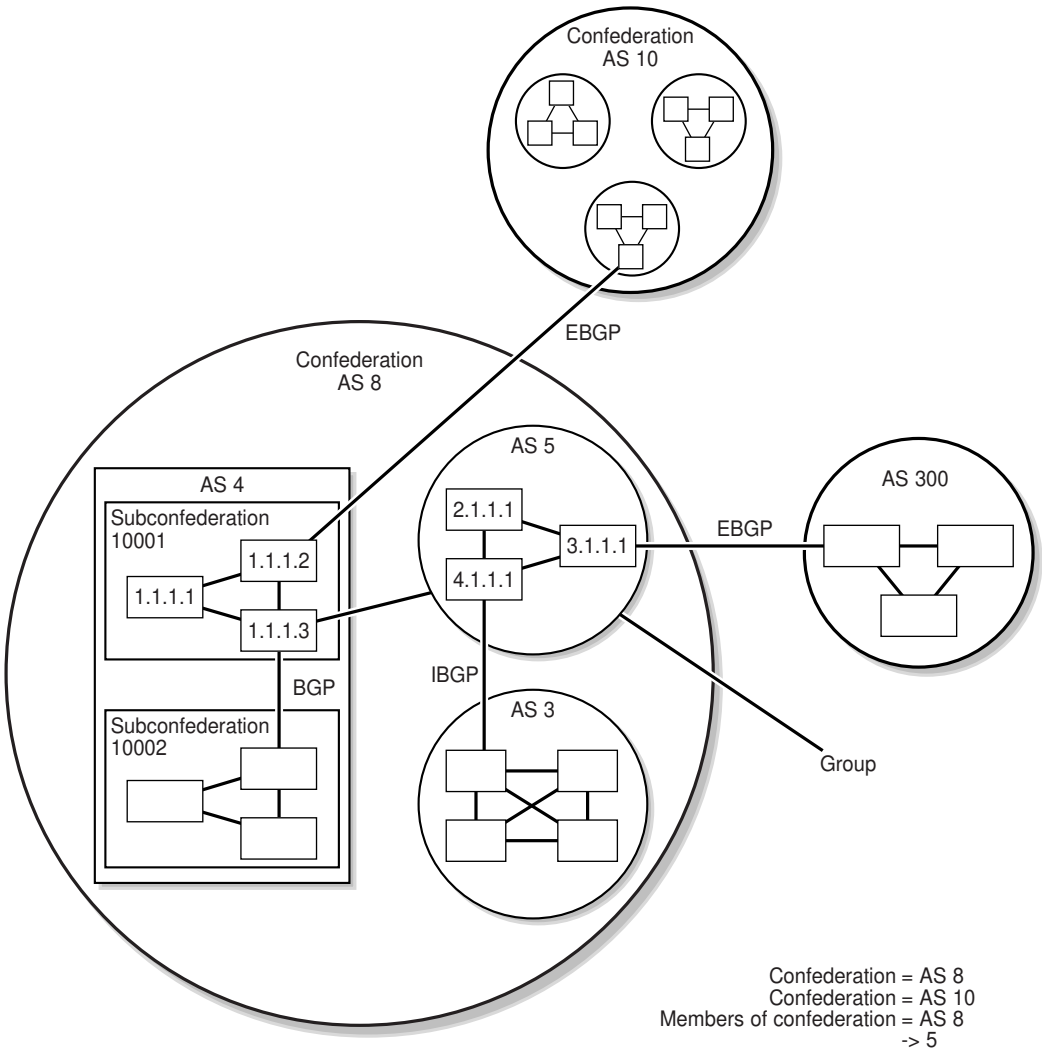

In a standard BGP configuration, all BGP-enabled devices within an AS have a full mesh of BGP peerings to ensure all externally learned routes are redistributed through the entire AS. This is needed because IBGP does not re-advertise routes learned from one IBGP peer to another IBGP peer. However, as more devices are added, scaling the IBGP mesh can become an issue. You can use the following subdividing methods to solve this scaling issue:

- BGP confederations: subdivides a large AS into smaller ASs yet still advertise as a single AS to external peers. The intent is to reduce IBGP mesh size.
- Subconfederations: further subdivide ASs; within each smaller AS, or confederation, IBGP is still used, however EBGP is used between subconfederations. This means less meshing between peers is required.
- Route reflector: ASs are divided into groups called clusters. Each cluster contains at least one route reflector. The route reflector redistributes routing updates to all devices in its cluster. Because the route reflector provides all the routing updates, the other devices in the cluster do not maintain a BGP mesh.
- BGP Optimal Route Reflection (BGP-ORR): BGP-ORR uses IGP on a route reflector to advertise the best path to the BGP-ORR client groups.

  Client groups sharing the same or similar IGP topology can be grouped as one BGP peer group, with BGP-ORR enabled. You can configure one of the client NEs as the primary NE for BGP-ORR so that the IGP metric from that NE is used to select the best path and advertise it to the clients in the same BGP peer group. Optionally, you can also select other NEs as secondary or tertiary backup NEs, to be used when the primary NE goes down or is unreachable.

By default, BGP routers will advertise only the best path to other BGP peers. When the best path fails or a better path is discovered, the old routing information is overwritten. If BGP PIC Add-Path is enabled, peers will advertise multiple paths for the same destination, all of which will be retained in their routing tables. In the event of a failure or path withdrawal message, the router is able to update its forwarding entry for that destination without having to receive new path advertisements. This results in faster convergence after path failure.

28.5.2 MP-BGP

MP-BGP is an extension to BGP that allows different types of addresses (known as address families) to be distributed in parallel. Standard BGP only supports IPv4 Unicast addresses; MP-BGP supports both IPv4 and IPv6 addresses and supports Unicast and Multicast variants of each. MP-BGP allows information about the topology of IP multicast-capable routers to be exchanged separately from the topology of normal IPv4 Unicast routers. The routes associated with multicast routing are used by PIM to build multicast data distribution trees.

BGP distributes IP routing information between networks. The distribution of IP routing information between VPRNs requires MP-BGP. MP-BGP addressing uses an 8-byte RD with a 4-byte or 16-byte IP address, depending on whether IPv4 or IPv6 is used. The requirements for MP-BGP configuration are the following.

- MP-BGP must be enabled on the participating PE devices.
- All PE devices must be configured as BGP peers.

When PE devices learn routing information from the CE devices in a VPRN configuration, the routing information is shared using MP-BGP. The RD is used to associate the new routing information with a VPRN instance.

The PE devices distribute the route information to the other CE devices in the VPRN. Each learned route is assigned an MPLS label that is distributed to the devices.

When customer packets arrive at a PE device, the packets are encapsulated with the MPLS label that corresponds to the learned route for the packet destination.

The MP-BGP multicast extension can be applied to build separate routing tables for multicast paths. IGPs such as OSPF and IS-IS can import multicast and Unicast routing information to the multicast and Unicast routing tables. The multicast routing information can subsequently be used by PIM to perform RPF lookups.

28.5.3 BGP SIDR prefix origin validation

BGP SIDR prefix origin validation can help prevent BGP prefix spoofing. BGP speakers access cryptographically secured information about the AS numbers that are authorized to originate routes for different prefixes.

If a BGP speaker that supports SIDR prefix origin validation receives a route from an EBGP peer and the origin AS of the route is incorrect for the advertised prefix, the route is considered invalid and may be given a lower preference or not propagated further. When a device is configured with RPKI router sessions, the device stores the information it receives from cache servers in a local database called the origin validation database. The entries in this database are called VRP (Validated ROA Payload) entries.

See 28.34 “To configure BGP SIDR prefix origin validation” (p. 930) for information about configuring BGP SIDR prefix origin validation.

28.6 BMP

28.6.1 Overview

BMP is a protocol used for monitoring BGP sessions. BMP allows a BGP router (BMP client) to advertise unprocessed routing information it has received from its peers to a monitoring station. The peer can be configured to send pre policy route information, post policy information, or both. The monitoring station can then monitor routing table size, identify issues and monitor trends in table size and update and withdraw frequency.

A BMP client is configured to create monitoring sessions using TCP with one or more monitoring stations.

The following are some of the messages provided by BMP during a monitoring session:

- Route Monitoring: the initial dump of all routes received from a peer and ongoing updates to the list of paths advertised by a peer
- Statistics Reports: output of statistics sent periodically to the monitoring station
- Peer Up Notification: sent whenever a peering session is established
- Peer Down Notification: sent whenever a peering session is terminated, including reasons for termination

28.7 RIP and RIPng

28.7.1 Overview

RIP is an IGP that uses a distance-vector algorithm to determine the best route to a destination, using hop count as the deciding factor. In order for the protocol to provide complete information about routing, every device in the domain must participate in the protocol. RIP, a UDP-based protocol, updates its neighbors, and the neighbors update their neighbors. RIP advertises IPv4 routing information, and RIPng advertises IPv6 routing information.

Unlike OSPF and other link-state protocols, RIP directly advertises reachability information to its neighbors. RIP advertises reachability information by sending prefix, mask, and either hop count or cost metric data. Each device running the RIP protocol advertises all RIP devices periodically by sending RIP update PDUs. The route with the lowest metric is advertised as the best route.

The NFM-P supports the configuration of RIPv1 and RIPv2 on network and access interfaces, as well as on IES and VPRN subscriber interfaces and group interfaces.

The NFM-P supports the configuration of IPv6 addresses for RIP on the base routing instance of the 7750 SR (chassis modes C and D, and chassis modes A and B with mixed mode enabled) and the 7705 SAR-Hm.

28.8 LDP

28.8.1 Overview

LDP is used to distribute labels in non-traffic-engineered MPLS applications. Routers can establish LSPs across a network by mapping network-layer routing information directly to the data link layer switched paths. After the LDP distributes the labels to the LSR, the LSR assigns the label to an FEC, and then informs all other LSRs in the path about the label and how the label switches data accordingly.

An FEC is a collection of common actions associated with a class of packets. LDP helps establish the LSP by providing a set of procedures that LSRs can use to distribute labels.

When a service tunnel is configured on managed routers using LDP signaling in an MPLS environment, LDP sessions are set up based on the configured hello and other PDU values. If another service tunnel is created to the same destination, the LDP session is reused.

The LDP sessions between LSRs:

- find and establish LDP peers in the managed network
- exchange label mappings for each LSR
- exchange label bindings

After all the LSRs are LDP-aware and the LSP is created, forwarding can occur as follows:

1. An FEC is associated with the LSP.
2. The FEC maps the packets to the LSP.

3. The next LSR that is part of the LSP splices incoming FEC labels to the outgoing FEC label of the next hop.

There are three types of LDP:

- T-LDP
- LDP DU
- LDP DoD

T-LDP is used to distribute labels for VLL and VPLS. T-LDP allows the targeting of remote devices that are not directly connected as targeted peers.

LDP DU can be used to create tunnels between PEs for IP-VPN services.

LDP DoD allows LDP peers to request label bindings for specific FECs.

The following figure shows an example of LDPs that are used in a simple Layer 2 and Layer 3 service provider network.

Figure 28-2 LDP sample network

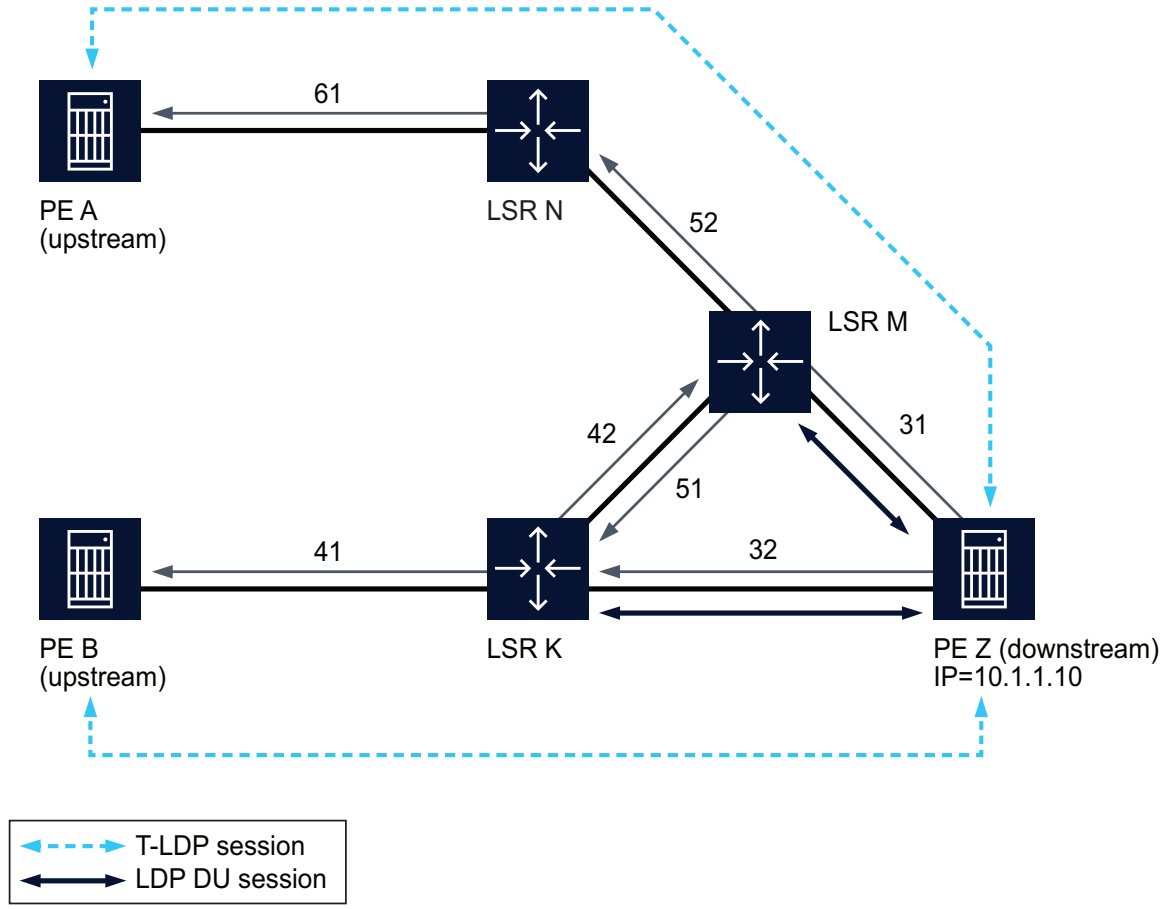

- The solid straight lines between the devices indicate IP connectivity. These are directly connected peers.
- The dotted bidirectional curved lines indicate T-LDP sessions. These are targeted peers that are not directly connected.
- The solid bidirectional curved lines indicate LDP DU sessions. This example only shows two instances, between M and Z, and Z and K. If there is IP connectivity between all the devices, then all of the devices have LDP DU sessions.

Provider edge router Z advertises the labels for its address 1.1.1.10 to adjacent link state devices M and K. Routers M and K distribute the labels for that address to the rest of the network.

If provider edge router A wants to send a VPN-labeled packet to router Z, it uses label 61 as the outer label. When the packet reaches router N, outer label 61 is swapped for outer label 52 and the packet continues downstream to router M. Router M then swaps out label 52 for either outer label 31 or 42, depending on the router M label selection algorithm. If outer label 31 is selected, the packet reaches router Z directly, which then continues to route the packet to the customer site based on the VPN label.

If a router receives more than one request for an IP address, and the router does not support ECMP, it selects the first label it receives. If the router supports ECMP, it selects the label with the lowest cost path, or selects both labels if the cost paths are identical.

When a router supports ECMP, LSRs can have multiple equal cost paths to an IP address. ECMP LDP retains all labels it receives from multiple next hop peers. The forwarding plane contains multiple next hops for an FEC and as a result, provides load balancing for LDP-based LSPs.

When a router supports ECMP LDP and a device configured as the next hop is no longer valid, for example a session between peers is lost or the peer withdraws its label, a new valid LDP next hop peer is selected and the forwarding plane is updated.

Load balancing across LSPs for LDP over RSVP is also supported. When ECMP is enabled, all equal cost LSP endpoints are installed in the routing table by the IGP for consideration by LDP. LDP selects the LSP with the lowest LSP metric to determine the next hop. If multiple LSPs are available with equal cost, then ECMP is utilized until the ECMP count is exhausted. An LDP tree is created by sending LSP Trace messages along an ECMP path to downstream LSR NEs.

Weighted ECMP is supported for fine control of load balancing, especially where capacity of local links is unequal or variable.

28.8.2 LDP for P2MP

LDP support for P2MP can be enabled on an LDP interface. You can configure an LDP interface for multicast traffic forwarding towards a downstream NE. LDP configuration allows an exchange of P2MP FEC via an established session to a peer, and the use of next hops over an interface.

LDP configuration is supported on a tunnel interface under a base routing instance and under protocols such as PIM and IGMP. A P2MP ID is used as an index for LDP-based tunneling instead of an LSP ID.

Note: BFD and source redundancy are only supported on an RSVP tunnel interface.

28.9 IS-IS

28.9.1 Overview

IS-IS is a link-state interior gateway protocol that uses the shortest path first algorithm to determine a route. Routing decisions are made using the link-state information. IS-IS entities include:

- networks, which are autonomous system routing domains
- intermediate systems, which are routers
- end systems, which are network devices that send and receive PDUs

End systems and intermediate system protocols allow devices and NEs to identify each other. The IS-IS protocol sends link state updates periodically through the network, so each device can maintain current network topology information.

Large networks, or autonomous systems, are supported by IS-IS using a two-level hierarchy. This divides a large area into more manageable, smaller areas. The first level (level 1) of routing is performed within an area. The second level (level 2) of routing is performed between areas, as shown in the following figure.

Figure 28-3 IS-IS routing areas example

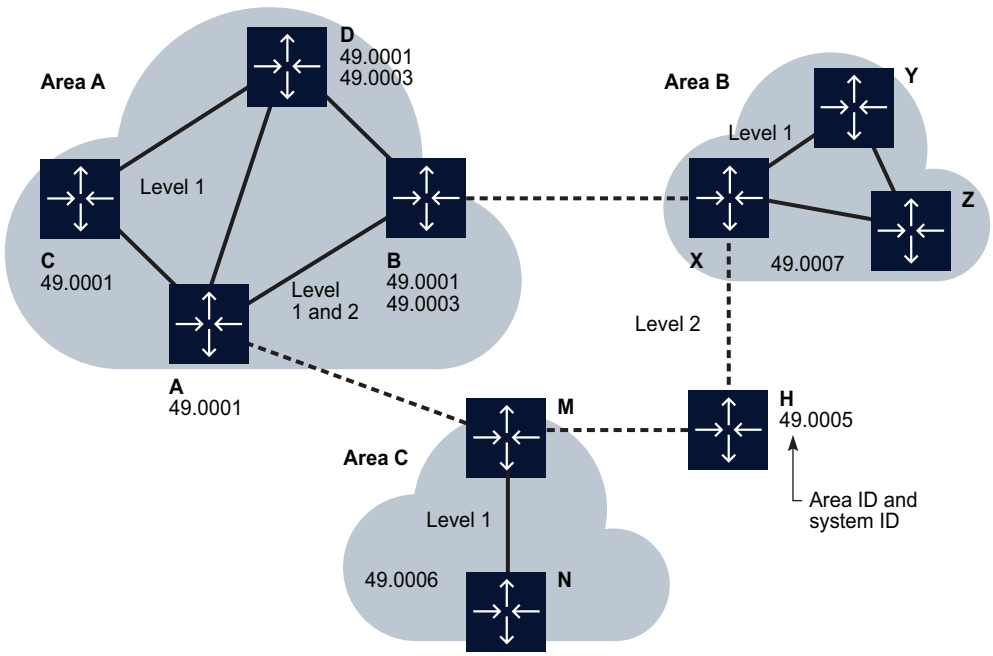

Level 2 areas are also called backbones, similar to an OSPF backbone area. All traffic traversing different areas must traverse the backbone. A device can be configured as level 1, level 2, or both level 1 and 2. In this example, routers A, B, M, H, and X form the level 2 IS-IS backbone. The connection between routers A and B carries both level 1 and level 2 link-state PDUs (LSPs).

However, level 1 devices are only aware of their own area's topology, and must forward traffic to a level 1/2 device to forward the data to another area.

Two devices are in the same level 1 area when they have level 1 adjacency. Level 1 adjacency occurs when the area IDs are common and there is a level 1 connection between the devices. Level 2 adjacency occurs when it has at least one level 1 and 2, or one level 2 interface configured.

The NFM-P supports the configuration of IPv6 addresses for IS-IS adjacencies.

Note: If two neighboring devices in the same level 1 area run both level 1 and 2, they establish both a level 1 and a level 2 adjacency.

When LDP over RSVP is enabled for IS-IS, LSPs can be used by the IGP to calculate its SPF tree. The IGP then provides LDP with all of the ECMP IDP next-hops and tunnel endpoints that the IGP identifies as the lowest cost path to the destination. If an IGP calculation and an LDP over RSVP have the same cost, LDP chooses an LDP over RSVP tunnel over an IGP route and ECMP between the two types is not considered. The type and number of tunnels that are to be considered by LDP depend on the IGP costs, where the lowest cost between the tunnel endpoint and the target is selected.

After IS-IS is configured, routing occurs as follows:

1. Hello PDUs are sent to IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
2. IS-IS neighbor relationships are formed.
3. Link-state PDUs (LSPs) are created based on local interfaces and prefixes that are learned from adjacent devices.
4. The devices flood LSPs to adjacent neighbors, and build a link-state database.
5. A shortest path tree is calculated by the IS-IS and the routing table is built.

28.9.2 Flexible Algorithm support

Flexible IGP algorithms allow IGPs to compute constraint-based paths over the network. A flexible algorithm definition can be configured for each NE in the global NE properties. The flexible algorithm definition provides topology constraints on a per-NE basis. A flexible algorithm definition requires an administrative group; see 54.13 "To configure an administrative group policy" (p. 1758).

A set of flexible algorithm definitions creates a flexible algorithm topology. If flexible participation has been configured on an ISIS instance, the IGP uses constraint-based shortest path first and the flexible algorithm definitions to find the best paths through the flexible algorithm topology.

28.9.3 IPv6 Traffic Engineering support

The NFM-P supports IPv6 Traffic Engineering (TE) extensions for IS-IS.

In addition, ISIS supports advertising which protocol is enabled on a given TE-link (SR-TE, RSVP-TE, or both) via the Application Specific Link Attributes sub-TLV. This allows the advertising router to send potentially different Link TE attributes for RSVP-TE and SR-TE applications.

To enable this feature, in summary, you must:

- configure the Interface Name parameter in the routing instance. The IP address of the specified interface is used as the IPv6 Traffic Engineering Router ID.
- configure traffic engineering on the ISIS Site.

28.10 OSPFv2 and OSPv3

28.10.1 Overview

OSPF is a hierarchical link-state interior gateway protocol that operates within ASs and is used in IP networks. OSPFv2 applies to IPv4 addressing whereas OSPFv3 applies to IPv6 addressing. OSPFv2 is defined in RFC 2328; OSPFv3 is defined in RFC 5340.

OSPF packets are routed based on the destination IP address of the IP packets. Each OSPF router collects link-state information to build a network topology based on OSPF areas. This topology is used to apply the Dijkstra algorithm to calculate the shortest path to each destination in the network.

OSPFv3 can run in Sparse or Only mode. In Sparse mode OSPFv3 uses Legacy LSAs to compute the shortest path and extended LSA is used by segment routing. In Only mode OSPFv3 uses extended LSA for SPF computation. See the appropriate device documentation for more information about OSPF.

When LDP over RSVP is enabled for OSPF, LSP can be used by the IGP to calculate its SPF tree. The IGP then provides LDP with all of the ECMP IDP next-hops and tunnel endpoints that the IGP identifies as the lowest cost path to the destination. If an IGP calculation and an LDP over RSVP have the same cost, LDP chooses an LDP over RSVP tunnel over an IGP route and ECMP between the two types is not considered. The type and number of tunnels that are to be considered by LDP depend on the IGP costs, where the lowest cost between the tunnel endpoint and the target is selected.

28.10.2 Flexible Algorithm support

Flexible IGP algorithms allow IGPs to compute constraint-based paths over the network. A flexible algorithm definition can be configured for each NE in the global NE properties. The flexible algorithm definition provides topology constraints on a per-NE basis. A flexible algorithm definition requires an administrative group; see 54.13 “To configure an administrative group policy” (p. 1758).

A set of flexible algorithm definitions creates a flexible algorithm topology. If flexible participation has been configured on an OSPF instance, the IGP uses constraint-based shortest path first and the flexible algorithm definitions to find the best paths through the flexible algorithm topology.

28.10.3 OSPF areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical OSPF area. The topology of the area is hidden from the rest of the AS, which significantly reduces OSPF protocol traffic. Routing in the AS takes place on two levels, depending on whether the source and

destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; routing information obtained from outside the area is not used. Routers that belong to more than one area called an ABR. An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

28.10.4 OSPF neighbors

Routers in the same broadcast domain or at each end of a point-to-point telecommunications link form adjacencies when they detect each other. The detection occurs when a router identifies itself in a hello OSPF protocol packet. This is called a two-way state and is the most basic relationship. The routers in an Ethernet or frame relay network select a designated router and a backup designated router which act as a hub to reduce traffic between routers. OSPF uses both Unicast and multicast to send hello packets and link state updates.

The OSPF protocol establishes and maintains neighbor relationships in order to exchange routing updates with other routers. The neighbor relationship table is called an adjacency database in OSPF. Provided that OSPF is configured correctly, OSPF forms neighbor relationships only with the routers directly connected to it. In order to form a neighbor relationship between two routers, the interfaces used to form the relationship must be in the same area. An interface can only belong to a single area.

28.10.5 OSPF super-backbone

The OSPF super-backbone provides an additional layer of hierarchy in OSPF. The OSPF super-backbone functions include:

- loop prevention
- handling LSAs received from the CE
- managing VPN IPv4 routes received by BGP

The PE routers that connect OSPF areas to the super-backbone function as ABRs in the OSPF areas to which they are attached. To achieve full compatibility, the PE routers can also serve as ASBRs in non-stub areas.

PE routers insert inter-area routes from other areas into the area in which the CE router is present. The CE routers are not involved at any level and are not aware of the super-backbone or other OSPF areas that exist outside of the super-backbone.

When you configure the super-backbone, all destinations that are learned by PEs with matching domain IDs become inter-area routes.

See the appropriate device documentation for more information about the OSPF super-backbone.

28.11 Segment routing

28.11.1 Overview

Segment routing enables an IGP to perform shortest path routing and source routing using segments. A segment can represent a local prefix of an NE, a specific adjacency of the NE, a

service context, or a specific explicit path over the network. For each segment, the IGP advertises a segment identifier, or SID. The SID can be used to resolve forwarding in several contexts, such as BGP labelled routes, VPRN and EVPN auto-binding of tunnels, and remote loop-free alternate.

When segment routing is used together with MPLS, the segment identifier is a standard MPLS label. An NE forwarding a packet using segment routing pushes one or more MPLS labels. See 31.6 “To configure an MPLS instance” (p. 1116) for more information about how to configure the segment routing label range.

Each prefix SID represents a network global IP address. Therefore the SID index for a prefix must be unique in the network. All routers in the network are expected to configure and advertise the same prefix SID index range for an IGP instance. However, the label value used by each router to represent this prefix can be local to that router with the use of a start label.

Two mutually exclusive prefix SID types can be configured: global and local.

With the global prefix SID type, the IGP instance assumes the start label value is the lowest label value in the SRGB and the prefix SID index range size is equal to the range size of the SRGB. If the global option has been selected for one IGP instance, all IGP instances on the system will be restricted to do the same.

With the local prefix SID type, the user configures a subset of the SRGB by specifying the start label value and the prefix SID index range size. Note that all resulting net label values (start-label + index) must be within the SRGB or the configuration will fail. Also, the NFM-P checks for overlaps of the resulting net label value range across IGP instances and will strictly enforce that the ranges do not overlap.

You can use the NFM-P to configure segment routing and to assign an SID index or label to ISIS and OSPF instances; see 28.79 “To configure IS-IS segment routing” (p. 978) and 28.80 “To configure OSPF segment routing” (p. 980).

28.11.2 Segment routing policy

A segment routing policy specifies a source-routed path from a head-end router to an end-point, and it specifies the traffic flows that should be steered into that source-routed path. An SR policy can be explicitly configured or learned from a protocol such as BGP or PCEP. For information about configuring a BGP site to learn segment routing policies, see 28.78 “To enable SR policy support on a BGP site, peer, or peer group” (p. 977). For information about creating a segment routing policy, see 28.77 “To create a segment routing policy” (p. 976).

28.11.3 Segment routing tree

A segment routing tree is an SID based multicast tree. It is very similar to a P2MP LSP.

An SR tree policy can be used in the same manner as a P2MP RSVP-TE or mLDP, for example, it can be used as a provider tunnel under an MVPN.

An SR tree policy represents a multicast tree from a root to a set of leaves. The SR tree policy can contain redundant trees, each with its own preference. This redundancy is presented via candidate paths. Each candidate path represents a P2MP tree with its own traffic engineering constraints. These candidate paths can be optimized based on link failures or IGP optimizations. As such each candidate path can contain multiple P2MP LSPs represented by path Instance IDs. The candidate path can perform make before break between these path instances. The candidate path with the highest preference becomes the active candidate path.

An SR tree policy is relevant only on the root node where the segment routing tree is instantiated. An SR tree policy does not have any forwarding information for the P2MP LSP. It only contains root and leaf information and the traffic engineering information for the tree to be set up from the root to the leaves. The forwarding information is part of the replication policy. Therefore the transit and leaf nodes only contain the replication policy.

For information about creating a segment routing tree, see 28.81 “To create a segment routing tree” (p. 981).

Replication segments

A replication segment is a forwarding instruction for a P2MP LSP. It contains the incoming replication SID and a set of outgoing interfaces and their corresponding set of outgoing labels. A replication policy consists of a treeID, rootID and path-instanceID (LSP-ID). If a replication segment is shared between multiple roots or P2MP LSPs, it does not have a rootID. A shared replication segment has a replication segment identifier denoted in the treeID. This type of replication segment can be used for FRR.

28.11.4 Segment Routing with IPv6

SR-IPv6 allows Segment Routing to be deployed over non-MPLS networks and/or in areas of the network where MPLS is not present. An IPv6 packet consists of an IPv6 header, extension headers, and payload. The Segment Routing header (SRH) is an extension header added to IPv6 packets to implement Segment Routing IPv6 (SRv6) based on the IPv6 forwarding plane. It specifies an IPv6 explicit path and stores IPv6 segment lists that function in the same way as segment lists in SR-MPLS.

The segment lists in an SRv6 SRH are processed from the bottom up, which is different from SR-MPLS processing. Another difference between SRv6 and SR-MPLS is that segments in SRv6 SRHs are not removed after being processed by nodes.

SRv6 segments are identified using segment identifiers (SIDs) encoded as IPv6 addresses. An SRv6 SID consists of two parts: Locator and Function. The Locator part provides the location function. The Function part identifies an instruction bound to the node that generates the SRv6 SID.

Each SRv6-capable node keeps a local SID table containing all SRv6 SIDs generated on the node. The destination SID identifies a destination node. It is similar to a node SID in SR-MPLS. In shortest path routing, the destination SID is encoded in the Destination Address (DA) field of the outer IPv6 header. Unlike usual segment routing, SID value is derived from the locator prefix configured and the end function.

SRv6 is supported from NFM-P Release, 22.2 R1 onwards.

28.12 RSVP

28.12.1 Overview

RSVP is a network control protocol that hosts use to request specific qualities of service from the network for specific data streams. RSVP is also used to deliver QoS requests to all devices in a data path and to establish and maintain the state information required to provide the requested service quality.

RSVP is not a routing protocol. RSVP operates using Unicast and multicast routing protocols.

RSVP consults local routing tables to relay RSVP messages. By default, RSVP is enabled on all devices that support it.

RSVP requests typically result in the reservation of resources on each device in the data path. MPLS use this RSVP mechanism to set up traffic-engineered LSPs. RSVP requests resources for unidirectional flows only. RSVP treats a sender as logically distinct from a receiver, although the same application process may act as both a sender and a receiver.

The RSVP hello protocol detects the loss of a neighbor NE or the reset of a neighbor RSVP state information. In standard RSVP, neighbor monitoring occurs as part of the RSVP soft-state model. The reservation state is maintained as cached information that is first installed and then periodically refreshed by the ingress and egress LSRs.

If the state is not refreshed within a specified time interval, the LSR discards the state because it assumes that either the neighbor NE has been lost or its RSVP state information has been reset.

28.12.2 Diff-Serv Traffic Engineering support

The NFM-P supports Diff-Serv Traffic Engineering, or TE, extensions for RSVP. Diff-Serv TE extensions provide the ability to manage bandwidth in an MPLS network on a per-TE-class basis. With Diff-Serv TE, an LER can perform this on a per-class basis. Therefore, you can set different limits for admission control of LSPs in each TE class over each link in the network.

You do this by setting a bandwidth constraint, which configures the percentage of the RSVP interface bandwidth that each CT shares. The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the Maximum Reservable link bandwidth TE parameter, that is, the link bandwidth multiplied by the RSVP interface subscription parameter.

This configuration also exists at the RSVP interface level and the value specifically configured for the interface overrides the globally-configured value. The bandwidth constraint value can be changed on the fly. You are also allowed to specify the bandwidth constraint for a CT which is not used in any of the TE class definitions and which is not used by any LSP originating or transiting this NE.

To enable this feature, in summary, you must:

- configure Diff-Serv Classes on the router RSVP Routing Instances
- configure protocol properties to allow Diff-Serv Classes on the required router RSVP Interfaces
- configure traffic engineering and specify the Class Type that the required dynamic LSPs belongs to
- configure the Diff-Serv Class Type that the required LSP Paths belongs to. This overrides the value set at the LSP level.
- configure Diff-Serv to LSP FC mappings on the required service tunnels (SDPs). User-entered mappings of FC to LSP name are validated automatically (to avoid configuration conflicts) by checking with the RSVP module.

28.12.3 Dark bandwidth accounting support

The bandwidth availability advertised by IGP in a TE network accounts for RSVP-TE LSP bandwidth requirements. However, other types of labeled traffic may exists on the same links. The bandwidth used by this traffic is called dark bandwidth. You can configure accounting for MPLS-SR related dark bandwidth.

To enable this feature, you must:

- enable collection of Aux Stats, that is, statistics for dark bandwidth, on the MPLS routing instance
- enable and configure dark bandwidth accounting on the required router RSVP Interfaces

28.13 L2TP

28.13.1 Overview

L2TP is a session-layer protocol that extends the PPP model by allowing L2 and PPP endpoints to reside on different devices that are interconnected by a PSN. L2TP extends the PPP sessions between the CPE and PPP/L2TP termination point on the L2TP network server (LNS), via an intermediate L2TP access concentrator (LAC). The LAC is the initiator of session-generated L2TP tunnels; the LNS is the server that waits for new tunnels. Manually configured and initiated L2TP tunnels can be initiated and stopped from either the LNS or LAC.

After an L2TP tunnel is established, the network traffic between the peers is bidirectional. If a tunnel carrying a session fails, another tunnel from the same tunnel group re-establishes the session. Within each L2TP tunnel, one or more L2TP sessions can exist. Each L2TP session transports PPP packets.

The NFM-P supports the configuration and management of the following:

- ISA-LNS groups
- L2TP sites
- L2TP tunnel group profiles
- L2TP tunnel profiles

28.13.2 ISA-LNS groups

The NFM-P supports the creation and configuration of ISA-LNS groups on the 7750 SR. ISA-LNS groups provide LNS PPP session termination. An ISA-LNS group is associated with specific L2TP inbound peers and groups. Session traffic is automatically balanced across the available active ISA broadband application MDAs in the group.

The following operations can be performed on an ISA-LNS group member:

- drain MDA—prevents new sessions from being accepted
- stop drain—allows new sessions to be established
- stop MDA sessions—terminates the active sessions

See 13.12 “To configure an ISA-LNS group” (p. 426) for information about creating and configuring an ISA-LNS group.

28.13.3 L2TP sites

By default, L2TP is enabled on a base routing instance, and an L2TP site is created. An L2TP icon appears in the Routing view of the navigation tree. An L2TP site does not exist by default on the base routing instance of a 7450 ESS in non-mixed mode.

Certain LAC parameter values can be configured on an L2TP site which can, in turn, be inherited by any child L2TP tunnel group profiles of the site. The Tunnel Session Limit parameter value can only be inherited by an L2TP tunnel group profile.

A tunnel selection blacklist function provides a means to temporarily suspend (or blacklist) tunnels that return one or more specified return codes, indicating that the tunnel is inoperative. The user can specify the manner in which blacklisted tunnels are handled:

- how many tunnels can be blacklisted at any given time
- how long a tunnel can remain blacklisted
- action taken on a blacklisted tunnel when the maximum blacklist period expires

See 28.89 “To configure L2TP on a routing instance” (p. 989) for information about managing L2TP sites.

L2TP is not enabled by default on a VPRN site. To enable L2TP on a VPRN site, see 79.11 “To configure a VPRN site” (p. 2545) .

28.13.4 L2TP tunnel group profiles

An L2TP tunnel group profile represents the configuration for a group of L2TP tunnels. L2TP tunnel group profiles must be configured on the LNS NE. If a local user database is used on the LAC for session authentication, an L2TP tunnel group profile must be configured on the LAC. When you create an L2TP tunnel group profile, you can configure certain LAC parameter values to be inherited from the parent L2TP site. In turn, the profile can be used by its child L2TP tunnel profile to inherit certain parameter values. The inherited parameter values are used as the default values during L2TP tunnel profile configuration.

The following operations can be performed on an L2TP tunnel group profile that has a configured L2TP tunnel profile:

- drain—prevents new sessions from being accepted
- stop drain—allows new sessions to be established
- stop—closes tunnel instances and terminates the active sessions

See 28.89 “To configure L2TP on a routing instance” (p. 989) for information about managing L2TP tunnel group profiles.

28.13.5 L2TP tunnel profiles

You can create and manage L2TP tunnel profiles from the Tunnel Group Profile properties form. When you create an L2TP tunnel profile, you can configure certain LAC parameter values to be inherited from the parent L2TP tunnel group profile and L2TP site.

The following operations can be performed on an L2TP tunnel profile:

- start instance—attempts to create new tunnels

- stop instance—attempts to close the tunnels
- drain—prevents new sessions from being accepted
- stop drain—accepts new sessions

See 28.89 “To configure L2TP on a routing instance” (p. 989) for information about managing L2TP tunnel profiles.

28.13.6 L2TP tunnel instance endpoints

The endpoints of an operational L2TP tunnel are represented by tunnel instance endpoints. You can view tunnel instance endpoints from the Tunnel Instance Endpoints tab on the L2TP Site (Edit) form, and from the Tunnel Group Profile or Tunnel Profile properties form if the endpoint is created using the profile configuration.

An L2TP tunnel instance endpoint is automatically created when:

- a start operation is performed on an L2TP tunnel profile
- an incoming L2TP session is established using group and tunnel profiles
- RADIUS authentication returns a configuration for the tunnel when an incoming L2TP session is authenticated and PPP session authentication determines that L2TP is used, at which point an L2TP session is established.

The following operations can be performed on an L2TP tunnel instance endpoint:

- drain—prevents new sessions from being accepted
- stop drain—allows new sessions to be established
- stop—closes the tunnel instances and terminates the active sessions
- resync tunnel instance endpoints—updates all L2TP tunnel status objects on a routing instance to the current state

See 28.92 “To view L2TP tunnel instance endpoints on a subscriber instance” (p. 994) for information about viewing L2TP tunnel instance endpoints.

28.13.7 L2TP peers

An L2TP site can have none or multiple L2TP peers. L2TP peer information is available on the Peers tab of the L2TP Site form. Information about L2TP tunnels for a specific L2TP peer is available on the Tunnels tab of the L2TP Peer properties form.

The following operations can be performed on an L2TP peer:

- drain—prevents new sessions from being accepted
- stop drain—accepts new sessions

28.13.8 L2TP tunnels and tunnel endpoints

You can view information about L2TP tunnels and tunnel endpoints from the L2TP Tunnel - Endpoint A - Endpoint B form.

28.13.9 L2TP sessions

You can view information about L2TP sessions from the L2TP Session (View) form. You can access the form from the following places in the GUI:

- Manage Residential Subscribers form
- Sessions tab on the L2TP Site (Edit) form
- Sessions tab on the L2TP Tunnel Instance Endpoint (View) form
- L2TP panel on the PPPoE Session form (LAC only)
- L2TP Session button on the Subscriber Host form (LAC only)

See 28.93 “To view L2TP sessions” (p. 994) for information about viewing L2TP sessions from the Manage Residential Subscribers form.

28.13.10 PPP sessions

You can view information about PPP sessions from the PPP Session, PPPoE Session, and PPPoA Session forms. You can access the forms from the following places in the GUI:

- Manage Residential Subscribers form
- LNS PPP panel in the L2TP Session (View) form (LNS only)
- PPP Session button in the Subscriber Host form

See 28.94 “To view PPP sessions” (p. 995) for information about viewing PPP sessions.

28.14 PIM

28.14.1 Overview

PIM is a component of multicast routing that defines the one-to-many or many-to-many transmission of information. You can use the following variations for PIM configurations:

- sparse mode
- dense mode
- source-specific multicast
- bidirectional

Sparse mode is the most common PIM configuration. Sparse mode is used for data transmission to NEs in multiple Internet domains that contain a small ratio of NEs that subscribe to the multicast traffic. Dense mode is used when a large ratio of the potential NEs subscribe to the multicast traffic. In source-specific multicast, paths originate at a single, defined source. Bidirectional PIM is not source-specific.

28.14.2 Anycast RP

Anycast RP for PIM-SM enables fast convergence when a PIM RP router fails. The receivers and sources rendezvous at the closest RP after the router failure. Anycast RP allows an arbitrary number of RPs for each group in a single, shared-tree PIM-SM domain. Triple play configurations

that distribute multicast traffic using PIM-SM realize the benefits of fast RP convergence, which helps to avoid the loss of multicast data streams or IPTV delivery to the end user.

Anycast RP for PIM-SM environments is supported in the base routing PIM-SM instance of the service router, and in VPRN instances that are configured with PIM on supporting NEs.

Anycast RP for PIM requires the completion of the following configuration information:

- An IP address specified as the RP address. This address is statically configured, or distributed using a dynamic protocol, to all PIM routers in the domain.
- A set of routers in the domain designated as RPs for the RP address. These routers form an anycast RP.
- Each of these routers is configured with a loopback interface that uses the RP address. The RP address, or a prefix that includes the RP address, is injected into the unicast routing system inside the domain.
- Each router in the anycast RP also needs a separate IP address, to be used for communication between the RPs.
- Each router in the anycast RP is configured with the addresses of all other routers in the anycast RP. The addresses must be consistently configured in all RPs in the set.

The following figure shows a sample implementation of anycast RP for PIM-SM.

Figure 28-4 Sample implementation of anycast RP for PIM-SM

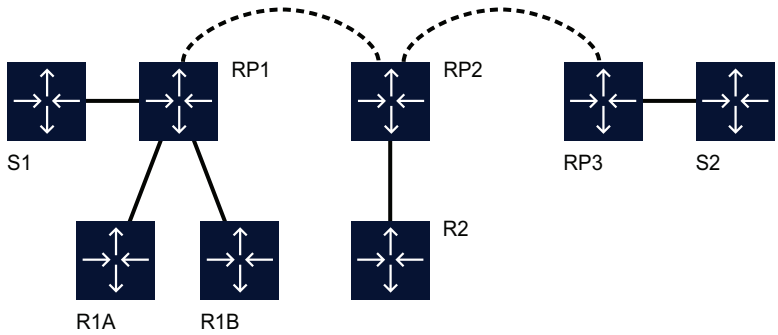

The following figure summarizes the sequence of events for the Anycast RP implementation shown in Figure 28-4, “Sample implementation of anycast RP for PIM-SM” (p. 902).

Table 28-1 Sequence of events for Anycast RP fast convergence

Sequence	Event description
1	S1 sends a multicast packet.

Table 28-1 Sequence of events for Anycast RP fast convergence (continued)

Sequence	Event description
2	The router connected to S1 forms a PIM registration message to send to the Anycast RP address. The Unicast routing system delivers the PIM registration message to the nearest RP, in this case RP1.
3	RP1 receives the PIM registration message, decapsulates the message, and sends the packet down the shared tree to the R1A and R1B receivers.
4	RP1 is configured with the IP address for RP2 and RP3. Since the registration message did not come from one of the RPs in the anycast RP set, RP1 assumes that the packet came from a designated router.
5	RP1 sends a copy of the registration message from the S1 designated router to RP2 and RP3. RP1 uses its IP address as the source address for the PIM registration message.
6	RP1 can join the source tree by sending a join message to S1. However, RP1 must create a source-specific state.
7	RP2 receives the registration message from RP1, decapsulates the message, and send the packet down the share tree to the R2 receiver.
8	RP2 sends a registration-stop message back to RP1. RP2 can wait to send the registration-stop message if it decides to join the source tree. RP2 should wait until it receives data from the source tree before it sends the registration-stop message. If RP2 does wait for the data, the registration-stop message is sent to RP1 when it receives the next registration message. If RP2 does not wait for the data, the registration-stop message is immediately sent to RP1.
9	RP2 can join the source tree by sending a join message to S1. However, RP2 must create a source-specific state.
10	RP3 receives the registration message from RP1 and decapsulates the message. No receivers joined for the group, so RP3 discards the packet.
11	RP3 sends a registration-stop message back to RP1.
12	RP3 creates a source-specific state, so when a receiver joins after S1 starts sending traffic, RP3 can quickly join the source tree for S1.
13	RP1 processes the registration-stop message from RP2 and RP3. RP1 can cache the receipt of registration-stop messages from the RPs in the anycast RP set. (The cache of messages is completed on a per-RP or per-source-specific basis.) The cache of messages increases the reliability of the delivery of registration messages to each RP. Subsequent registration messages received by RP1 are sent only to the RPs in the anycast RP set that have not previously sent registration-stop messages from the source-specific entry.
14	RP1 sends a registration-stop message to the DR under the following conditions: • after receiving a registration message from the DR • if all RPs in the anycast RP set have returned registration-stop messages for a specific source-specific route

28.14.3 SPT switchover thresholds

SPT switchover thresholds allow you to configure the switchover threshold, in Kbps, for the group prefixes. The threshold value determines when the router switches from the shared tree to the source-specific tree. The switchover is attempted only if the traffic rate on the shared tree for the group exceeds the configured threshold.

28.14.4 L3 Multicast Load Balancing for ECMP

The NFM-P distributes multicast traffic by balancing the load based on the total available multicast bandwidth on all ECMP paths. You can configure ECMP rebalancing on the VPRN and base PIM Routing Instance configuration forms.

NFM-P ECMP rebalancing has the following important characteristics:

- Multicast load balancing over ECMP links is enabled, by default.
- The rebalancing timer is set to 30 minutes, by default.
- Distribution of the multicast groups over the available links is processed based on the bandwidth configured for the specified group address. If the bandwidth is not configured for the multicast stream, then the configured default value is used.
- If a link failure occurs, the load on the failed channel is distributed to the remaining channels. The bandwidth required to accommodate the load from the failed link is evenly distributed over the remaining links.
- If an additional link becomes available for a specific multicast channel, it is then treated in an equivalent manner to the other links of the interface.
- A manual (operator-initiated) re-balance command is typically used to re-evaluate the current balance, with regard to bandwidth utilization. If necessary, multicast streams can subsequently be moved to different links to achieve a balance.
- In an automatic timed re-balance, the system re-balances multicast streams over the available links, based on the configured bandwidth and interval. If no links have been added or removed, or it is determined that no multicast streams will benefit from a re-balance, then it is not be implemented.
- When multicast load re-balancing is not enabled, any ECMP changes are not optimized. However, whenever a link is added, an attempt is made to balance the number of multicast streams on all the available ECMP links. This may however, not result in balanced bandwidth utilization of all the ECMP links.
- Only a single re-balance command can be executed at any specific time. If a re-balance is in progress and a manual re-balance command is entered, it is rejected and a message is displayed informing the user who a re-balance is already in progress.

28.14.5 VRRP-aware PIM

VRRP aware PIM enables PIM to track the state of a VRRP instance and know if the associated VRRP interface is the master. This feature is supported on base router, IES and VPRN interfaces.

PIM monitors the state of the VRRP interface using an operational group. The operational group is up when the VRRP interface is the master, and down for all other VRRP states. A VRRP instance can only be associated with one operational group, and an operational group can only have one associated VRRP instance.

If the monitored interface is the VRRP master, PIM becomes the DR by setting its priority to the configured Operational DR Priority value. Priorities must be configured so that the Operational DR Priority is the highest priority on the IP interface.

If a PIM router is the DR and then receives an indication from VRRP that the interface is no longer the VRRP master, PIM will relinquish the DR role by setting its priority back to the default or configured priority value.

If the VRRP instance or operational group is not configured, PIM will operate as normal with the default or configured priority value.

Two operational groups are supported per PIM interface, one for IPv4 and one for IPv6. A change in operational group status is address family independent; IPv4 and IPv6 priorities are configured independently of each other.

28.15 IGMP

28.15.1 Overview

IGMP is a multicast protocol which service providers can use to establish multicast group memberships on a LAN. Within the LAN, end users use IGMP to communicate with a local multicast router, which then uses PIM to distribute the IGMP messages to other local and remote multicast routers. Multicast routers send regular membership queries to IGMP hosts which respond with membership reports. Multicast routers can use these reports to determine which hosts are interested in receiving particular multicast messages.

Configuring a group interface query source IP address on the IGMP site allows a VPLS IGMP snooping instance to register the port facing an ESM port as a multicast router port.

In a multi-chassis implementation on an NE configured with active/active MC-LAG, multicast IGMP traffic from an access network is forwarded to both multicast redirection interfaces. This would cause the downstream device to receive duplicate IGMP messages. The Redundant Multicast option allows you to identify a single IGMP interface as the designated forwarder to the redirection interface, ensuring that the downstream device only receives the IGMP traffic once.

IGMP operates above the network layer on IPv4 networks.

28.16 MSDP

28.16.1 Overview

MSDP is a protocol that enables multiple PIM-SM domains to communicate with each other using their own RPs. MSDP also enables multiple RPs in a single PIM-SM domain to establish MSDP mesh-groups and to synchronize information between anycast RPs about the active sources being served by each anycast RP peer. The 7950 XRS, 7750 SR, and the 7450 ESS in mixed mode support MSDP.

Each PIM-SM domain has its own RPs and MSDP enables these RPs to inform each other about active sources. When an active source is detected, the RPs send PIM-SM explicit join messages to the active source. When RPs in remote domains know about active sources, they can pass on this information to their local receivers and multicast data can be forwarded between the domains.

The RP learns about a new multicast source within its domain through the PIM register mechanism and encapsulates the first data packet in a source active message. After an RPF check, the MSDP source active message is flooded by each peer to its MSDP peers until it reaches every MSDP router in the interconnected networks. If the receiving MSDP peer is an RP, and the RP has a (*,G)

entry for the group in the source active message, the message is accepted. The RP creates an (S,G) state for the source and join to the shortest path tree for the state of the source.

The encapsulated data is forwarded down the RP shared tree. When the packet is received by a receiver's last hop router, the last-hop may also join the shortest path tree to the source.

The following figure shows a sample of how data flows from a source in Domain A to a receiver in Domain E in an MSDP implementation for PIM-SM.

Figure 28-5 Sample implementation of MSDP for PIM-SM

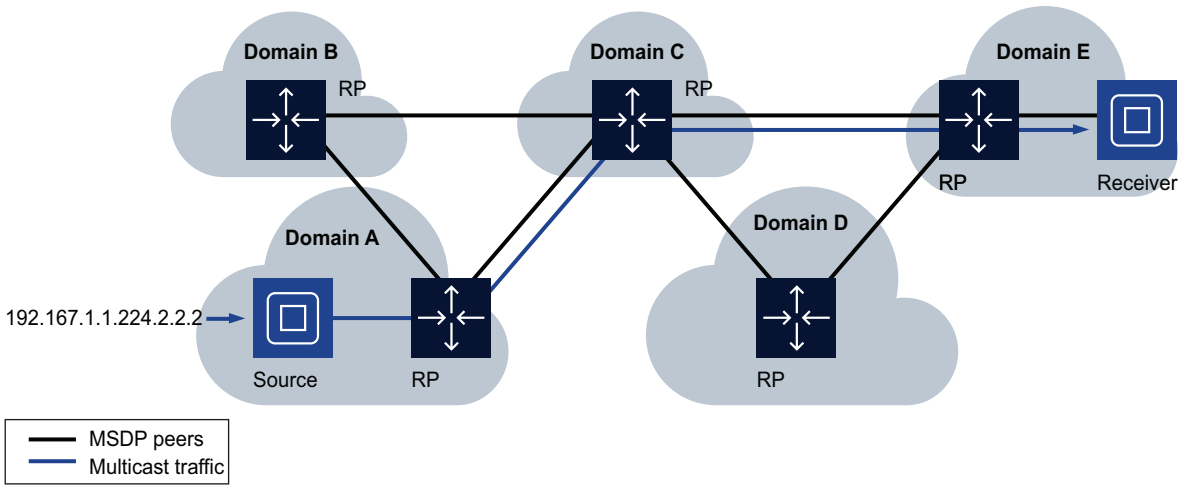

28.17 MLD

28.17.1 Overview

MLD is an asymmetric protocol used by IPv6 routers to discover the presence of multicast listeners, that is, NEs that wish to receive multicast packets. MLD specifies separate behaviors for multicast address listeners and multicast routers.

The purpose of MLD is to enable each multicast router to discover, for each of its directly attached links, which multicast addresses and which sources have interested listeners on that link. The information gathered by MLD is provided to whichever multicast routing protocol is used by the router, such as PIM, to ensure that multicast packets are delivered to all links where there are listeners interested in such packets.

An MLD group interface is configurable on a base routing instance, VPRN routing instance, and on an IES site. To allow the delivery of IPv6 multicast content to subscribers, group interfaces are configurable under MLD. MLD report/done/query messages use the link-local address as their source address. An MLD policy is applied to the subscriber profile to allow unique behavior per subscriber.

To support MCS for subscribers using MLD, MLD information is synced across chassis. During a failover scenario, standby MLD becomes active and PIM sends a join message to all necessary multicast sources.

MLD is supported on the 7950 XRS, in chassis mode C or D, and on the 7750 SR in chassis mode A or B with mixed mode enabled. MLD is supported on all variants of the 7705 SAR.

MLD under IES and VPRN is supported on the 7950 XRS, the 7750 SR in chassis mode D, and the 7450 ESS in chassis mode D with mixed mode enabled.

28.18 Bridging

28.18.1 Overview

Bridging is used between an OmniSwitch device and a 7450 ESS to create a fast ring. The fast ring bridges the end-user devices, such as BTV set-top boxes, and Layer 2 services, such as VPLS, that distribute multicast traffic. If required, you can configure the devices with MSTP or another STP.

STP is active by default on an OmniSwitch. The default mode of operation is 1x1 mode using RSTP. A loop-free network topology is automatically calculated based on default STP switch, bridge, and port parameter values.

Additional OmniSwitch configuration such as enabling and configuring OmniSwitch learned port security parameters for VLAN ports, DHCP relay, and DHCP snooping can be done while configuring other bridge parameters.

You can configure MVRP as part of the bridging configuration on the OmniSwitch. MVRP provides a mechanism to maintain the contents of dynamic VLAN registration entries for each VLAN, and to propagate that information to other bridges.

In the NFM-P, a bridge is represented in the network view below a device, similar to the routing instance of a device.

28.19 WPP

28.19.1 Overview

WPP (Web Portal Protocol) runs between a BNG and a Web portal server. It is used for Web portal authentication of WLAN users (DHCP hosts). WPP can be enabled on base routing instances or VPRN routing instances. WPP version 2.0 has an updated packet format and additional message types.

Subscriber creation triggered by WPP messages (log in via HTTP) supports a topology where the subscriber host IP address is provided via the CMTS NE. The ESM host is created after the client is authenticated through the Web portal.

In the case of WPP traffic on an SRRP dual homing configuration, the Web portal is not aware of SRRP mastership. A network interface on a base routing instance, or an L3 access interface on an IES or VPRN service can be configured with the Track SRRP Instance option, which specifies the master NE to which WPP traffic is directed.

28.20 BIER

28.20.1 Overview

Bit Indexed Explicit Replication (BIER) offers an architecture for forwarding multicast data packets without the need for multicast trees or for intermediate routers to maintain a per-multicast flow state. BIER information is distributed in the network using the underlay IGP.

A BIER domain is a connected set of bit forwarding routers, each with a unique ID. A BIER domain can be divided into sub-domains for scalability without a linear increase in size of the BIER header. For example, in IS-IS, a BIER sub-domain is an IS-IS multi topology, where ipv4-unicast is a single sub-domain and ipv4-multicast is another sub-domain.

BIER is supported on FP4 network interfaces. BIER is not supported on FP3 or earlier cards, or on access interfaces.

If a chassis has a mix of FP3 and FP4 network ports, BIER is signaled on all FP3 and FP4 interfaces which are part of the IS-IS.

28.20.2 BIER templates

BIER is configured by creating a BIER template. The BIER template contains the sub-domain to multi-topology mapping and other BIER configurations, such as the bit forwarding router ID and BIER prefix.

28.21 IPSec

28.21.1 Overview

IPSec tunnel parameters for VSR or 7705 SAR-Hm can be configured from a routing interface on the NE. See Chapter 34, “IPsec” for general information about IPSec.

Routing protocol configuration workflow and procedures

28.22 Routing protocol configuration workflow

28.22.1 Stages

1

Configure a base routing instance or VRF instance on a device; see 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

2

Enable the routing protocols to be supported on the device or routing instance. The supported protocols are device dependent. The options are:

- IPv6; enabled during interface creation
- BGP and MP-BGP; see 28.29 “To enable BGP on a routing instance” (p. 916) and 28.31 “To configure global-level BGP” (p. 918) , 28.32 “To configure peer-group-level BGP” (p. 922) , and 28.33 “To configure peer-level BGP” (p. 926)
- RIP and RIPng; see 28.44 “To enable RIP or RIPng on a routing instance” (p. 937) , 28.45 “To configure global-level RIP or RIPng” (p. 937) , 28.46 “To configure group-level RIP or RIPng” (p. 938) , and 28.47 “To configure interface-level RIP or RIPng” (p. 939)
- LDP; see 28.50 “To enable LDP on a routing instance” (p. 940) and 28.51 “To configure global-level LDP” (p. 941)
- IS-IS; see 28.59 “To enable IS-IS on a routing instance” (p. 952) and 28.60 “To configure IS-IS on a routing instance” (p. 953)
- OSPFv2 and OSPFv3; see 28.66 “To enable OSPF on a routing instance” (p. 963)
- RSVP; see 28.85 “To configure RSVP on a routing instance” (p. 985)
- L2TP; see 28.89 “To configure L2TP on a routing instance” (p. 989)
- WPP; 28.134 “To create a web portal routing instance” (p. 1051)
- Bridging; see 28.128 “To configure bridging on an OmniSwitch” (p. 1043)
- MPLS; see 31.5 “To enable MPLS on a routing instance” (p. 1115)
- MPLS-TP; see 32.3 “To enable MPLS-TP on a routing instance” (p. 1168)
- Multicast
  - PIM; see 28.97 “To enable PIM on a routing instance” (p. 998)
  - IGMP; see 28.103 “To enable IGMP on a routing instance” (p. 1016)
  - MSDP; see 28.112 “To enable MSDP on a routing instance” (p. 1026)
  - MLD; see 28.121 “To enable MLD on a base routing instance” (p. 1034)

Note: L2TP is enabled by default on a base routing instance, but not on a VPRN routing instance.

3

Configure L3 network interfaces as required if you need to enable multicast routing protocols. See the 27.1.17 “Workflow to configure NE routing and forwarding” (p. 823) in Chapter 27, “NE routing and forwarding” .

4

Configure the appropriate routing policies as required; see Chapter 54, “Routing policies” .

5

Review the following protocol workflows as required to configure additional functionality:

a. BGP and MP-BGP; see 28.28 “Workflow to configure BGP and MP-BGP” (p. 915)

b. RIP; see 28.43 “RIP and RIPng configuration workflow” (p. 936)

c. LDP; see 28.49 “Workflow to configure LDP” (p. 940)

d. IS-IS; see 28.58 “Workflow to configure IS-IS” (p. 952)

e. OSPFv2 or OSPFv3; see 28.65 “Workflow to configure OSPFv2 and OSPFv3” (p. 962)

f. RSVP; see 28.84 “Workflow to configure RSVP” (p. 985)

g. L2TP; see 28.88 “Workflow to configure L2TP” (p. 988)

h. PIM; see 28.95 “Workflow to configure PIM” (p. 997)

i. IGMP; see 28.102 “Workflow to configure IGMP” (p. 1016)

j. MSDP; see 28.111 “Workflow to configure MSDP” (p. 1025)

k. MLD; see 28.120 “Workflow to configure MLD” (p. 1033)

l. bridging; see 28.127 “Workflow to configure bridging on an OmniSwitch” (p. 1043)

m. WPP; see 28.133 “Workflow to configure WPP” (p. 1051)

n. BIER; see 28.135 “Workflow to configure BIER” (p. 1053)

o. IPSec; see 28.137 “Workflow to configure IPSec ” (p. 1055)

6

Configure routing protocol monitoring as required; see 28.39 “Workflow to configure BMP” (p. 933).

7

Configure the protocol for the remote device as required.

BFD and SBFD configuration workflow and procedures

28.23 BFD and SBFD configuration overview

28.23.1 General information

BFD configuration consists of the following:

- creation of a BFD template policy

  A BFD template policy defines the fast forwarding path failure detection requirements for BFD packets/sessions such as the transmit/receive timer interval, the minimum echo receive interval, and the number of consecutive BFD messages that can be missed before the BFD session status changes to down.
- configuration of BFD or SBFD on a path or interface

Configuration of Seamless BFD requires configuration of an SBFD reflector and peer. The reflector is configured in the NE properties and peers are added to the routing instance.

28.24 Workflow to configure BFD

28.24.1 Stages

1

Configure a BFD template policy; see 28.25 “To configure a BFD template policy” (p. 911).

2

Configure BFD or SBFD on the path or interface as needed; see the configuration procedures for the path or interface type.

3

Configure SBFD; see 28.26 “To configure Seamless BFD ” (p. 912).

28.25 To configure a BFD template policy

28.25.1 Purpose

You can apply a BFD template policy to the following:

- LSP-path bindings; see 31.13 “To configure a Dynamic or segment routing TE LSP” (p. 1132) .
- Proactive OAM template on an enabled MPLS-TP routing instance; see 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) .
- Bidirectional MPLS-TP LSP; see 32.7 “To create a bidirectional MPLS-TP LSP” (p. 1173) .
- PW Template Binding; see 77.110 “To configure a site for BGP AD or BGP VPLS” (p. 2408) .
- Spoke SDP binding

- For VLL sites; see 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) , or 76.35 “To configure a spoke SDP binding with an L2TPv3 tunnel on a VLL Epipe site” (p. 2165).
- For a VPLS; see 77.101 “To configure BFD on a VPLS SDP binding” (p. 2400) .
- For an IES; see 78.12 “To configure BFD on an IES spoke SDP binding” (p. 2441) .
- For a VPRN; see 79.71 “To configure BFD on a VPRN spoke SDP binding” (p. 2642) .

28.25.2 Steps

1

Choose Policies→BFD Templates from the NFM-P main menu. The Manage BFD Templates form opens.

2

Click Create. The BFD Template, Global Policy (Create) form opens.

3

Configure the required parameters.

4

Click Apply. The BFD Template (Edit) form refreshes with additional tabs.

5

As required, click on the appropriate tabs and click Search to determine to which network objects the BFD template policy is applied.

6

Click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

28.26 To configure Seamless BFD

28.26.1 Purpose

Use this procedure to create a Seamless BFD reflector and peer. The reflector is configured in the NE properties and peers are added to the routing instance.

28.26.2 Steps

To configure a Seamless BFD reflector

1

In the navigation tree equipment view, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then the BFD sub-tab.

3

Click Create. The Seamless BFD Reflector (Create) form opens.

4

Configure the parameters and click OK.

The reflector is added to the list in the NE properties and the routing instance.

To configure a Seamless BFD peer

5

In the navigation tree routing view, expand Network→*NE*→Routing Instance.

6

Click on the Seamless BFD tab, then the Peers sub-tab.

7

Click Create. The Seamless BFD Peer (Create) form opens.

8

Configure the parameters and click OK.

9

Save your changes and close the forms.

END OF STEPS

BGP configuration workflow and procedures

28.27 BGP configuration workflow and procedures

28.27.1 Overview

The BGP command hierarchy consists of three levels:

- global
- peer group
- peer (also known as neighbor)

The following figure shows the network view open to show Confederations and BGP settings.

Figure 28-6 BGP in the navigation tree network view

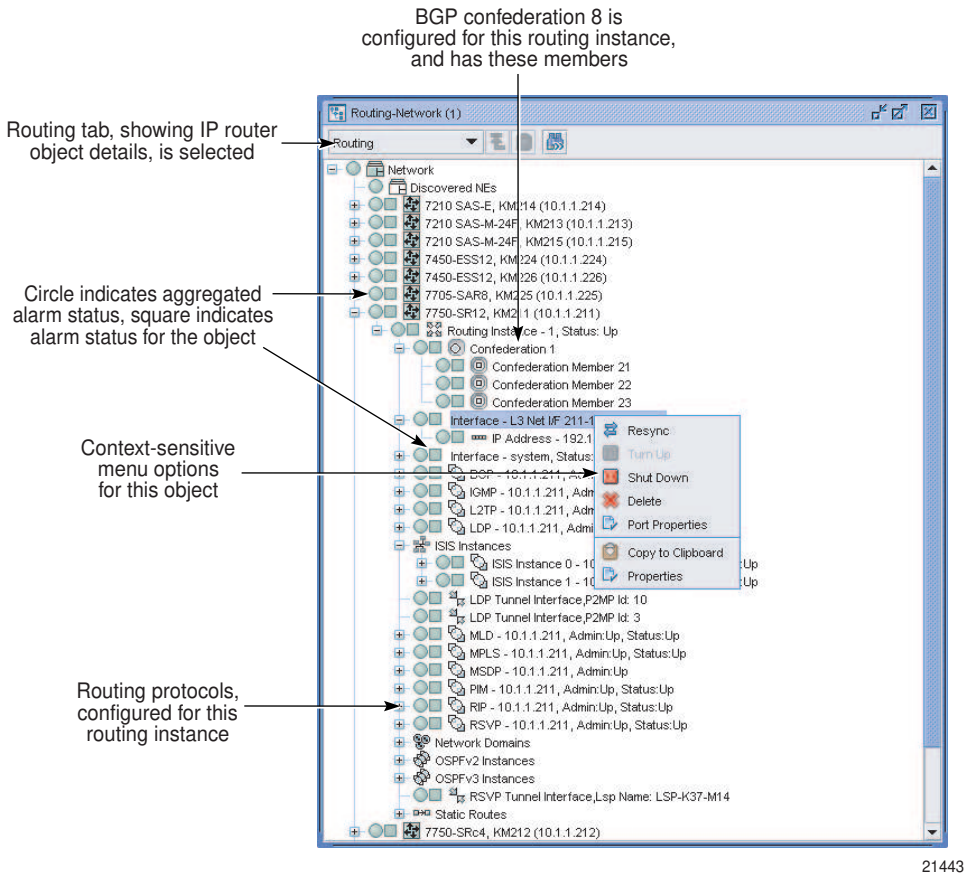

BGP parameters are initially applied at the global level. The parameters are inherited by the group and peer levels. Parameters can be modified and overridden on a level-specific basis.

Many of the hierarchical BGP commands can be modified at different levels. BGP group-level parameters take precedence over BGP global-level parameters. BGP peer-level parameters take precedence over group- and global-level parameters.

The following procedures describe how to configure BGP.

28.28 Workflow to configure BGP and MP-BGP

28.28.1 Stages

1

Enable BGP on a routing instance; see 28.29 “To enable BGP on a routing instance” (p. 916) .

2

If BGP mesh scaling is an issue due to the number of NEs in an AS, configure a BGP confederation; see 28.30 “To configure a BGP confederation” (p. 916) .

3

If you are configuring MP-BGP, enable multicast IPv4 on the base routing instance of each NE that is to participate in a VPRN; see 28.31 “To configure global-level BGP” (p. 918) and Chapter 79, “VPRN service management” .

4

Create at least one BGP peer group; see 28.32 “To configure peer-group-level BGP” (p. 922) .

5

Configure peer-level BGP as required; see 28.33 “To configure peer-level BGP” (p. 926) .

6

Configure BGP SIDR prefix origin validation as required; see 28.34 “To configure BGP SIDR prefix origin validation” (p. 930) .

7

Configure long-lived graceful restart as required; see 28.37 “To configure long-lived graceful restart on a BGP site” (p. 932).

8

Enable or disable BGP peering as required; see 28.35 “To enable or disable BGP peering” (p. 931) .

9

Configure a site for BGP AD or BGP VPLS as required, see 77.110 “To configure a site for BGP AD or BGP VPLS” (p. 2408) .

10

Verify BGP routes against a routing policy statement as required; see 54.23 “To verify BGP routes against a routing policy statement” (p. 1769)

11

Verify a BGP prefix list against a routing policy statement as required; see 54.24 “To verify a BGP prefix list against a routing policy statement” (p. 1770) .

28.29 To enable BGP on a routing instance

28.29.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties.

3

Click on the Protocols tab and configure the BGP Enabled parameter.

Note: You must configure an AS number before enabling BGP. Configure the Autonomous System parameter on the Routing tab of the Routing Instance (Edit) form.
If a confederation is required, configure the Confederation Autonomous System parameter on the Routing tab on the Routing Instance (Edit) form.

4

Save your changes and close the form.

END OF STEPS

28.30 To configure a BGP confederation

28.30.1 Purpose

For BGP confederations, the following rules apply:

- A device can only belong to one confederation.
- Multiple devices can belong to one BGP confederation.

You must configure BGP on the device and configure global-level BGP parameters before you configure BGP confederations, as described in 28.29 “To enable BGP on a routing instance” (p. 916) and 28.31 “To configure global-level BGP” (p. 918) .

28.30.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the Routing Instance icon and choose Properties.

3

Click on the Routing tab.

4

Enter the confederation number for the Confederation Autonomous System parameter and click Apply.

5

Click on the BGP Confederations tab.

6

Click Create to add a BGP confederation, or choose a BGP confederation and click Properties to configure an existing BGP confederation.

Note: You can only have one BGP confederation per device.

7

Add a new member to the confederation:

Perform the following steps:

1. Click on the Members tab.
2. Click Create.
3. Configure the Member AS parameter as the number of ASs for the confederation. The member AS number represents the BGP instance of the device.
4. Click OK.

   You can verify the confederation membership by opening the Confederation icon to view icons that represent the members of the confederation.

8

Close the form.

END OF STEPS

28.31 To configure global-level BGP

28.31.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→BGP.

2

Right-click on the BGP icon and choose Properties.

3

Click on the General tab and configure the required parameters.

4

Select a TCP key chain in the KeyChain panel, if required.

Note: You can assign a TCP key chain to a BGP site, group, or peer on a 7450 ESS in mixed mode, 7705 SAR, 7750 SR, or 7950 XRS.

5

Configure the required parameters in the Optimal Route Reflection Shortest Path First Timers panel.

6

Configure the required parameters in the Graceful Restart panel.

The graceful restart parameters are configurable when you enable the Graceful Restart parameter.

7

Configure the required parameters in the MED panel.

The parameters are used to find a way to leave the AS when there are multiple methods of leaving the AS.

8

Configure the required parameters in the MultiPath panel.

Set the MultiPath parameter to 1 to disable. When set from 2 to 16, multipath is enabled and BGP load shares traffic across the number of links specified. If the equal cost routes available are greater than the configured value, then routes with the lowest next hop IP address are chosen.

9

Configure the Backup Path parameter in the Backup Path panel.

10

Configure the required parameters in the Best Path Selection panel.

When you set the MED Compare parameter to off, the Strict parameter is set to true and is not configurable. When you set the MED Compare parameter to on, the Strict parameter is set to false and is not configurable.

To include routes from BGP peers for MVPN IPv4, MVPN IPv6, or both, in RR convergence, set the Ignore Router ID parameter to true and configure the Ignore Router ID Internal parameter.

11

Configure the required parameters in the OLC panel.

12

Click on the Behavior tab and configure the required parameters.

Note: You must set the Enable Add-Path Send parameter to True in order to configure the additional Add-Path parameters.
When you configure the Peer Tracking Policy, you can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression as described in 28.36 “To create a BGP policy expression” (p. 931).

13

Click on the AS Properties tab and configure the required parameters.

The Local AS parameters are used to configure a virtual AS. A virtual AS is used when a router (RTA) is moved from one AS (AS1) to another AS (AS2). However, the customer router (CR1) is configured to belong to the AS1. To avoid reconfiguring CR1 to belong to AS2, CR1 can continue to belong to AS1, but RTA has its local AS value set to AS1. RTA can advertise AS1 for routes advertised to CR1.

14

Click on the VPN tab and configure the required parameters.

Note the following:

- If you are configuring global-level BGP for BGP AD in VPLS or for BGP VPLS, you must enable the L2 VPN option for the Family parameter.
- You must set the AS Path Ignore parameter to True in order to configure the AS Path Ignore Family parameters.
- The Apply Import Route Policies and Apply Export Route Policies parameters are not configurable on a VPRN routing instance.
- The Apply Import Route Policies and Apply Export Route Policies parameters specify whether to apply the existing import and export route policies configured on the Import Policies and Export Policies tabs.
- Selecting an address family in the Rapid Update panel reduces the minimum route advertisement interval (MRAI) to zero. Families that are not selected will retain the default MRAI.

15

Configure next hop resolution:

1. Click on the Next-Hop Resolution tab.
2. Configure the Use BGP Routes, Resolve Label Routes, Allow Static Routes, and Weighted ECMP parameters as applicable.
3. To defer hardware fast-reroute protection to the SR tunnel, enable the Prefer Transport Frr parameter.
4. Select a routing policy statement.
5. Click Create Expression and configure the parameters to add a policy expression.
6. Configure the parameters in the Use Leaked Routes panel.
7. Choose a Shortcut Tunnel and click Properties.
8. Configure the required parameters for the Shortcut Tunnel, save the changes, and close the form.
9. In the Label Route Transport Tunnel panel, choose a family from the available entries and click Properties.
10. Configure the required parameters.

    You can choose segment routing as a tunnel option. Segment routing must be configured if you select the SR-ISIS or SR-OSPF parameters; see 28.11 “Segment routing” (p. 894).
11. Save your changes and close the form.

16

Configure ORF Extended Community:

1. Configure the Accept ORF parameter.
2. Configure the Send ORF parameter.

   If Send ORF is set to True, provide the RTL details:
   a. Click on the Create button.
   b. Enter site and route target details.
   c. Click OK.

17

Configure Optimal Route-Reflection:

1. Click on the Optimal Route-Reflection tab.
2. Click Create. The Optimal Route Reflection (Create) form opens.
3. Configure the parameters.
4. Click OK to create an optimal route-reflection location and close the form.

   You can create up to 16 locations.
5. Configure the Cluster ID and Optimal Route-Reflection location parameters on the General tab.

18

Click on the Segment Routing V6 tab and configure the required parameters.

19

Add a peer group.

Perform the following steps:

1. Click on the Groups tab and click Create.
2. Perform Step 3 to Step 22 of 28.32 “To configure peer-group-level BGP” (p. 922) .
3. Save the changes.

20

Add a peer.

Perform the following steps:

1. Click on the Peers tab.
2. Click Create.
3. Perform Step 3 to Step 18 of 28.33 “To configure peer-level BGP” (p. 926) .
4. Save the changes.

21

Configure dynamic peer prefixes as required:

1. Click on the Dynamic Peer Prefixes tab.
2. Click Create.
3. Configure the required parameters.
4. Save the changes.

22

Create a route target list as required:

1. Click on the Route Target List tab.
2. Click Create to create a route target.
3. Configure the required parameters.
4. Save the changes.

Routes with one or more of the route targets in the list are accepted or advertised to route-reflector clients. This is only applicable if the router is a route-reflector server.

23

Click on the Import Policies and Export Policies tabs and configure the required parameters.

Configure the import route policies to determine which routes are accepted from peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

Configure the export route policies to determine which routes are advertised to peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . A router performs no validation to ensure the policies match.

You can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression, as described in 28.36 “To create a BGP policy expression” (p. 931).

24

Click on the RIB Management tab to assign leak import and export policies:

1. On the IPv4 Leak Import Policies, IPv6 Leak Import Policies, Label IPv4 Leak Import Policies, and Label IPv6 Leak Import Policies tabs, assign up to 15 leak import policies to the BGP site as required.
2. Click on the IPv4 Leak Export Policies tab and assign up to 15 leak export policies to the BGP site.
3. Click on the Route Table Import tab and configure the required parameters.

25

Click on the Authentication tab and configure the required parameters.

26

Save your changes and close the forms.

END OF STEPS

28.32 To configure peer-group-level BGP

Note: For most parameters in this procedure, you can specify that the parameter value is inherited from the parent BGP configuration using the Inherit Value parameter.

If you disable value inheritance for a parameter, the available options are restricted, based on the parent parameter value and the protocol functionality. For example, if a parameter in the global-level BGP configuration is set to True, the only available option for the same parameter in the peer-group-level BGP configuration is False, unless a value of False violates a protocol rule, in which case the only available option is True.

The parameters that you configure for a BGP peer group take precedence over the parameters that are configured for global-level BGP.

28.32.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→BGP.

2

Right-click on the BGP icon and choose Create Group.

3

Configure the required parameters.

4

Select a TCP key chain in the KeyChain panel, if required.

Note: You can assign a TCP key chain to a BGP site, group, or peer on a 7450 ESS in mixed mode, 7705 SAR, 7750 SR, or 7950 XRS.

5

Configure the required parameters in the Graceful Restart panel.

The graceful restart parameters are configurable when the Graceful Restart parameter is enabled.

6

Configure the required parameters in the MED panel.

The parameters are used to find a way to leave the AS when there are multiple methods of leaving the AS.

7

Click Apply.

8

Configure the required parameters in the OLC panel.

9

Configure dynamic peer prefixes as required:

1. Click on the Dynamic Peer Prefixes tab.
2. Click Create.
3. Configure the required parameters.
4. Save the changes.

10

Click on the Behavior tab and configure the required parameters.

The Fault Tolerance parameter allows BGP routers to be more tolerant of certain UPDATE message errors, use less disruptive error recovery mechanisms, and provide better operational and diagnostics information.

The Idle Timeout parameter is configurable when the Never Timeout parameter is disabled.

11

Click on the AS Properties tab and configure the required parameters.

The Peer AS parameter specifies the peer AS for this specific group, and the behavior, internal or external. Multipath configurations are not supported at the BGP peer level.

12

Configure the required parameters in the Local AS panel.

The No Prepend Global AS is configurable only if you set the Local AS parameter to a value other than 0.

The Local AS parameters are used to configure a virtual AS. A virtual AS is used when a router (RTA) is moved from one AS (AS1) to another AS (AS2). However, the customer router (CR1) is configured to belong to the AS1. To avoid reconfiguring CR1 to belong to AS2, CR1 can continue to belong to AS1, but RTA has its local AS value set to AS1. RTA can advertise AS1 for routes advertised to CR1.

13

Configure the required parameters in the Remove Private AS panel.

14

Click on the VPN tab and configure the required parameters in the VPN panel.

Note the following:

- The Advertise Label, Apply Import Route Policies and Apply Export Route Policies parameters are not configurable on a VPRN routing instance.
- The Apply Import Route Policies and Apply Export Route Policies parameters specify whether to apply the existing import and export route policies configured on the Import Policies and Export Policies tabs.
- Selecting an address family in the Rapid Update panel reduces the minimum route advertisement interval (MRAI) to zero. Families that are not selected will retain the default MRAI.

15

Configure ORF Extended Community:

1. Configure the Accept ORF parameter.

2. Configure the Send ORF parameter.

   If Send ORF is set to True, provide the RTL details:

   a. Click on the Create button.

   b. Enter site and route target details.

   c. Click OK.

16

Click on the Add Path tab and configure the required parameters.

You must select the Enable Add-Path parameter in order to configure the additional Add-Path parameters.

IMPORTANT INFORMATION

For release 22.2 and later, the modification of parameters under Add Path will still work with pre 22.2 OSSI scripts. However, to enable or disable Add-Path parameter using pre 22.2 OSSI scripts, the following workaround is mandatory;

- To disable AddPath, specify **<sendAddPath>false</sendAddPath>** along with the **<paramsInheritanceMask>** tag with all AddPath bits removed.
- To enable AddPath, specify **<sendAddPath>true</sendAddPath>** along with the **<paramsInheritanceMask>** tag with AddPath bits specified.

Note: It is recommended to use either pre-22.2 and 22.2 OSSI script consistently for Add-Path Configuration and do not alternate.

Note: The inheritance from parent BGP configuration applies to all the parameters. The inheritance values are displayed as read-only fields for operational setup. For non-operational setup, inheritance values are not displayed.

17

Click on the Segment Routing V6 tab and configure the required parameters.

18

To add a peer, perform the following steps:

1. Click on the Peers tab.
2. Click Create.
3. Perform Step 3 to Step 18 of 28.33 “To configure peer-level BGP” (p. 926) .
4. Save your changes and close the form.

19

To configure prefix limits, perform the following steps:

1. Click on the Prefix Limit tab.

   Note:

For node releases earlier than 13.0R1, prefix limit parameters appear on the Behavior tab and the Prefix Limit tab is not available.

2. Click Create. The Prefix Limit (Create) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

20

Click on the Import Policies and Export Policies tabs and configure the required parameters.

Configure the import route policies to determine which routes are accepted from peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

Configure the export route policies to determine which routes are advertised to peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . A router performs no validation to ensure the policies match.

You can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression, as described in 28.36 “To create a BGP policy expression” (p. 931).

21

To configure interface instance, perform the following steps:

1. Click on the Dynamic Peer Interfaces tab.
2. Create or select an entry and click Properties. The Dynamic Peer Interface (Create/Edit) form opens.
3. Configure the required parameters in General tab.
4. Click on the AS Allowed Ranges tab and configure the Minimum and Maximum values.

   Note:

   Upto 32 AS Allowed Range can be configured.
5. Save the changes and close the form.

22

Click on the Authentication tab and configure the required parameters.

23

Save your changes and close the form.

END OF STEPS

28.33 To configure peer-level BGP

Note: For most parameters in this procedure, you can specify that the parameter value is inherited from the parent BGP configuration using the Inherit Value parameter.

If you disable value inheritance for a parameter, the available options are restricted, based on the parent parameter value and the protocol functionality. For example, if a parameter in the peer-group-level BGP configuration is set to True, the only available option for the same parameter in the peer-level BGP configuration is False, unless a value of False violates a protocol rule, in which case the only available option is True.

The parameters that you configure for a BGP peer take precedence over the parameters that are configured for group-level BGP.

28.33.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→BGP→Peer Group.

2

Right-click on the Peer Group icon and choose Create Peer.

3

Configure the general parameters.

4

Select a TCP key chain in the KeyChain panel, if required.

Note: You can assign a TCP key chain to a BGP site, group, or peer on a 7450 ESS in mixed mode, 7705 SAR, 7750 SR, or 7950 XRS.

5

Configure the required parameters in the Graceful Restart panel.

The graceful restart parameters are configurable when the Graceful Restart parameter is enabled.

6

Configure the required parameters in the MED panel.

The MED parameters are used to find a way to leave the AS when there are multiple methods of leaving the AS.

7

Click Apply.

8

Configure the required parameters in the OLC panel.

9

Click on the Behavior tab and configure the required parameters.

Note: The Fault Tolerance parameter allows BGP routers to be more tolerant of certain UPDATE message errors, use less disruptive error recovery mechanisms, and provide better operational and diagnostics information.

10

Click on the AS Properties tab and configure the required parameters.

The Peer AS parameter specifies the peer AS for the group to which the peer belongs, and the behavior, internal or external. Multipath configurations are not supported at the BGP peer level.

11

Configure the required parameters in the Local AS and Remove Private AS panels.

The No Prepend Global AS parameter is configurable when you set the Local AS parameter to a value other than 0.

12

Click on the VPN tab and configure the required parameters.

Note the following:

- To enable VPN IPv6, select VPN IPv6 and IPv6 for the Family parameter.
- To enable VPN IPv6, select IPv6 for the Advertise Label parameter.
- The Advertise LDP Prefix and Use Service Routes parameters are configurable when the Advertise Label parameter is enabled for IPv4 routes.
- The Apply Import Route Policies and Apply Export Route Policies parameters specify whether to apply the existing import and export route policies configured on the Import Policies and Export Policies tabs.
- Selecting an address family in the Rapid Update panel reduces the minimum route advertisement interval (MRAI) to zero. Families that are not selected will retain the default MRAI.

13

Click on the ORF Extended Community tab and perform the following:

1. Configure the Accept ORF parameter.
2. Configure the Send ORF parameter.

   If Send ORF is set to True, provide the RTL details:

   a. Click on the Create button.
   b. Enter site and route target details.

c. Click OK.

14

Click on the Add Path tab and configure the required parameters.

You must select the Enable Add-Path parameter in order to configure the additional Add-Path parameters.

IMPORTANT INFORMATION

For release 22.2 and later, the modification of parameters under Add Path will still work with pre 22.2 OSSI scripts. However, to enable or disable Add-Path parameter using pre 22.2 OSSI scripts, the following workaround is mandatory;

- To disable AddPath, specify **<sendAddPath>false</sendAddPath>** along with the **<paramsInheritanceMask>** tag with all AddPath bits removed.
- To enable AddPath, specify **<sendAddPath>true</sendAddPath>** along with the **<paramsInheritanceMask>** tag with AddPath bits specified.

Note: It is recommended to use either pre-22.2 and 22.2 OSSI script consistently for Add-Path Configuration and do not alternate.

Note: The inheritance from parent BGP configuration applies to all the parameters. The inheritance values are displayed as read-only fields for operational setup. For non-operational setup, inheritance values are not displayed.

15

To configure prefix limits, perform the following steps:

1. Click on the Prefix Limit tab.

   Note:

   For node releases earlier than 13.0R1, prefix limit parameters appear on the Behavior tab and the Prefix Limit tab is not available.
2. Click Create. The Prefix Limit (Create) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

16

Click on the Import Policies and Export Policies tabs and configure the required parameters.

Configure the import route policies to determine which routes are accepted from peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

Configure the export route policies to determine which routes are advertised to peers. The policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . A router performs no validation to ensure the policies match.

You can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression, as described in 28.36 “To create a BGP policy expression” (p. 931).

17

Click on the Segment Routing V6 tab and configure the required parameters.

18

Click on the Authentication tab and configure the required parameters.

19

Save the changes and close the form.

20

Configure the protocol for the far-end device, if applicable. Use CLI for devices that are managed outside the scope of the NFM-P.

END OF STEPS

28.34 To configure BGP SIDR prefix origin validation

28.34.1 Purpose

To help prevent BGP prefix spoofing, you can configure validation of the AS values that are received from EBGP peers. When a BGP speaker that supports the validation function receives a route AS value, the speaker can check the associated prefix for validity. If the origin AS is not correct for the advertised prefix, the route is considered invalid and treated according to the routing policy configuration.

28.34.2 Steps

1

Configure an RPKI session on a routing instance; see 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

2

Configure origin validation on an EBGP peer session.

Perform the following steps:

1. Configure the Origin Validation parameter on the Behavior tab on the Peer Group (Edit) form; see 28.32 “To configure peer-group-level BGP” (p. 922) .
2. Configure the Origin Validation parameter on the Behavior tab on the Peer, Peer Group (Create) form; see 28.33 “To configure peer-level BGP” (p. 927) .

3

Create a policy statement entry to match IPv4 or IPv6 routes in the routing information database. Configure the Origin Validation State parameter on the Action, From Criteria, and Default Action tabs; see 54.5 “To configure a routing policy statement” (p. 1745) .

4

On a BGP site properties form, on the General tab in the Best Path Selection panel, set the Compare Origin Validation State and Origin Invalid Unusable parameters to control how the origin validation states associated with routing information database entries are used in the BGP decision process. See 28.31 “To configure global-level BGP” (p. 918) .

5

On a Community Routing Policy, configure the Community Member parameter with the new origin validation state. See 54.8 “To configure a community policy” (p. 1752) .

END OF STEPS

28.35 To enable or disable BGP peering

28.35.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→BGP→Peer Group.

2

Click on the Peer Group icon to display the peers in the peer group.

3

Right-click on a peer and choose one of the following menu items:

a. Turn Up to activate a peer

b. Shut Down to deactivate a peer

4

Click Yes.

END OF STEPS

28.36 To create a BGP policy expression

28.36.1 Overview

Use this procedure to create logical expressions when configuring policy selections on BGP configuration forms.

28.36.2 Steps

1

Click on the Create Expression button in a BGP site, peer, or peer group configuration form.

2

Click on the Select button and choose a policy. Click on the Add button to add the policy to the expression. You can add up to 15 policies to a single expression.

3

Use the Boolean operator buttons (AND, OR, and NOT) to form expressions, or click on the text box to manually enter an expression. You can use parentheses to create more complicated expressions.

4

Save your changes and close the form.

END OF STEPS

28.37 To configure long-lived graceful restart on a BGP site

28.37.1 Steps

1

Click on the Graceful Restart tab on a BGP site or BGP peer configuration form.

2

Configure the parameters in the Long-lived Graceful Restart panel.

3

Create stale route timer configurations for each applicable family.

1. In the Family panel, click Create. The Long-lived Graceful Restart (New Instance) form opens.
2. Configure the stale route timer parameters and click OK. A new entry appears in the Family panel. You can create one entry for each family type.

4

Click OK and confirm to close the form.

END OF STEPS

BMP configuration workflow and procedures

28.38 BMP configuration overview

28.38.1 General information

BMP configuration consists of the following:

- configuring a BMP client, including creation of monitoring stations
- enabling monitoring on BGP sites, groups, or peers

BMP monitoring status is not inherited from global BGP settings. Stations that have been configured globally can be selected from a local BGP entity.

28.39 Workflow to configure BMP

28.39.1 Stages

1

Configure BGP on the NE or VPRN routing instance; see 28.28 “Workflow to configure BGP and MP-BGP” (p. 915) or 79.23 “To configure BGP on a VPRN routing instance” (p. 2557).

2

Configure a BMP client; see 28.40 “To configure an NE as a BMP client” (p. 933)

3

Enable BMP; see 28.41 “To enable BMP” (p. 934).

28.40 To configure an NE as a BMP client

28.40.1 Steps

1

In the navigation tree Equipment or Routing view, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then the BMP sub-tab.

3

Configure the Administrative State parameter.

4

Click on the Stations sub-tab.

5

From the Stations sub-tab, click Create. The Monitoring Station (Create) form opens.

6

Configure the parameters to create a monitoring station.

7

Repeat Step 5 to create additional monitoring stations as needed. Up to eight monitoring stations can be created.

8

Save your changes and close the forms.

END OF STEPS

28.41 To enable BMP

28.41.1 Steps

1

Perform one of the following:

a. Navigate to a BGP site, peer group, or peer on an NE.

In the navigation tree Routing view, expand to one of the following:

- Network→NE→Routing Instance→BGP
- Network→NE→Routing Instance→BGP→Peer Group.
- Network→NE→Routing Instance→BGP→Peer Group→Peer.

b. Navigate to a BGP site, peer group, or peer on a VPRN routing instance.

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
3. On the service navigation tree, expand to one of the following:
   - Sites→Routing Instance→Protocols→BGP
   - Sites→Routing Instance→Protocols→BGP→Peer Group.
   - Sites→Routing Instance→Protocols→BGP→Peer Group→Peer.

Right-click on the BGP entity and choose Properties. The *BGP entity* (Edit) form opens.

2

Click on the BMP tab.

3

Configure the Enable Monitoring parameter and click Apply. The Monitor and Stations panels appear.

4

Configure the All Stations parameter to enable all stations to monitor BGP, or select individual stations in the Stations panel.

5

Configure the parameters.

6

Save your changes and close the forms.

END OF STEPS

RIP and RIPng configuration workflow and procedures

28.42 RIP and RIPng configuration overview

28.42.1 Overview

The RIP command hierarchy consists of three levels:

- global
- group
- interface (also known as neighbor)

For RIP configuration, you must define at least one group and one interface. The parameters that are configured on the global level are inherited by the group and interface levels. Parameters can be modified and overridden on a level-specific basis.

Many of the hierarchical RIP commands can be modified on different levels. RIP group-level parameters take precedence over BGP global-level parameters. RIP interface-level parameters take precedence over peer-group and global-level parameters.

When configuring RIP on the 7705 SAR, 7750 SR, or 7705 SAR-Hm, a RIP routing site is created under the routing instance which can be seen in the routing network tree. Using this RIP routing site, RIP groups and RIP interfaces can be configured.

The following procedures describe how to configure RIP and RIPng.

28.43 RIP and RIPng configuration workflow

28.43.1 Stages

1

Enable RIP or RIPng on a routing instance; see 28.44 “To enable RIP or RIPng on a routing instance” (p. 937) .

2

Configure global-level RIP or RIPng as required; see 28.45 “To configure global-level RIP or RIPng” (p. 937) .

3

Configure group-level RIP or RIPng as required; see 28.46 “To configure group-level RIP or RIPng” (p. 938) .

4

Configure neighbor-level RIP or RIPng as required; see 28.47 “To configure interface-level RIP or RIPng” (p. 939) .

28.44 To enable RIP or RIPng on a routing instance

28.44.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties.

3

Click on the Protocols tab and select the RIP Enabled or RIPNG Enabled parameters.

4

Save your changes and close the form.

END OF STEPS

28.45 To configure global-level RIP or RIPng

28.45.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→RIP or RIPng.

2

Right-click on the RIP or RIPng icon and choose Properties.

3

Configure the required general parameters.

4

Click on the Behavior, Authentication, Import Policies and Export Policies tabs, and configure the required parameters.

5

Create one or more RIP or RIPng groups, as required.

Perform the following steps:

1. Click on the Group tab.
2. Click Create. The RIP Group (Create) form opens.

3. Configure the required parameters. See 28.46 “To configure group-level RIP or RIPng” (p. 937) .

6

Create one or more RIP or RIPng interfaces, as required.

Perform the following steps:

1. Click on the Interface tab.
2. Click Create.
3. Configure the required parameters. See 28.47 “To configure interface-level RIP or RIPng” (p. 939) .

7

Save and close the forms.

END OF STEPS

28.46 To configure group-level RIP or RIPng

Note: You can choose to inherit values from the global-level RIP configuration by selecting the Inherit Value parameter. If you choose not to inherit a parameter value from the global-level RIP configuration, only the parameter options that are not set in the parent configuration are available. For example, if the Check Zero parameter is set to false in the global-level RIP configuration, you can only set the parameter to true in the group-level RIP configuration.

The parameters that you configure for a RIP group take precedence over the parameters that are configured for the global-level RIP configuration.

28.46.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→RIP or RIPng.

2

Right-click on the RIP or RIPng icon and choose Create Group.

3

Configure the Name parameter.

4

Configure the required general parameters.

5

Click on the Behavior, Authentication, Import Policies and Export Policies tabs, and configure the required parameters.

6

Save your changes and close the form.

END OF STEPS

28.47 To configure interface-level RIP or RIPng

Note: You can choose to inherit values from the global-level RIP configuration by selecting the Inherit Value parameter. If you choose not to inherit a parameter value from the global-level RIP configuration, only the parameter options that are not set in the parent configuration are available. For example, if the Check Zero parameter is set to false in the global-level RIP configuration, you can only set the parameter to true in the interface-level RIP configuration.

The parameters that you configure for a RIP interface, also known as a RIP neighbor, take precedence over the parameters that are configured for the group- and global-level RIP configuration.

28.47.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→RIP or RIPng.

2

Right-click on the RIP or RIPng group icon and choose Create Interface.

3

Configure the required general parameters.

4

Select an RIP group.

5

Select an interface.

6

Click on the Behavior, Authentication, Import Policies and Export Policies tabs, and configure the required parameters.

7

Save your changes and close the form.

END OF STEPS

LDP configuration workflow and procedures

28.48 LDP configuration overview

28.48.1 Overview

An LDP instance in the NFM-P network navigation tree has the following child objects:

- Interfaces–The interfaces object contains the configured LDP interfaces for directly connected peers.
- Targeted Peers–The targeted peers object contains the indirectly connected peers.

28.49 Workflow to configure LDP

28.49.1 Stages

1

Enable LDP on a routing instance; see 28.50 “To enable LDP on a routing instance” (p. 940) .

2

Configure global-level LDP; see 28.51 “To configure global-level LDP” (p. 941) .

3

Configure one of the following options:

a. If two LDP peers are adjacent devices, create an LDP interface; see 28.52 “To configure an LDP interface” (p. 944) .

b. If two LDP peers are non-adjacent devices, create an LDP targeted peer; see 28.53 “To configure an LDP targeted peer” (p. 945) .

4

Create an LDP keychain; see 28.54 “To configure an LDP peer” (p. 946) .

5

Configure ECMP for LDP routing as required; see 28.55 “To configure ECMP for LDP routing” (p. 948) .

6

View LDP session information as required; see 28.56 “To view the LDP session information” (p. 949) .

28.50 To enable LDP on a routing instance

28.50.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Protocols tab and enable the LDP enabled parameter.

4

Save the changes.

END OF STEPS

28.51 To configure global-level LDP

28.51.1 Steps

1

Enable LDP on a router, as described in 28.50 “To enable LDP on a routing instance” (p. 941) .

2

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→LDP.

3

Right-click on the LDP icon and choose Properties. The LDP, Routing Instance (Edit) form opens.

4

Configure the Administrative State parameter.

5

Click on the Common tab and configure the required parameters.

Configure the Targeted Sessions Allowed parameter to true to configure the router for T-LDP. Targeted sessions are LDP sessions that distribute labels between indirectly connected peers.

The Neighbor Liveness Time (seconds) and Maximum Recovery Time (seconds) parameters are configurable when the Enforce Graceful Restart parameter is enabled.

6

Click on the Interface Properties tab and configure the required parameters. By default, the parameter values are inherited by all LDP interfaces.

7

Click on the Targeted Peer Properties tab and configure the required parameters. By default, these parameter values are inherited by all LDP targeted peers.

Note: If the Enable Hello Reduction parameter is disabled while the targeted LDP session to the peer is operationally up, the change in the hold time will take effect only when the targeted LDP session is re-established.

8

Click on the Policies tab and configure the required parameters.

9

Click on the Import Tunnel Table and Export Tunnel Table sub-tabs and configure the required parameters.

The tunnel table policies are used to allow LDP-capable PE devices to offer services to PE routers in areas or domains where BGP-labeled routes are not supported. They are used to determine which routes are advertised to LDP. Only policy entries supporting the LDP protocol are considered; all other policies are ignored. See also 28.33 “To configure peer-level BGP” (p. 927) for information on enabling the advertisement of LDP prefixes to BGP peers.

10

Click on the OAM tab, then the Ecmp sub-tab. See 28.55 “To configure ECMP for LDP routing” (p. 948) for more information about configuring ECMP for LDP interfaces.

11

Click on the Interfaces tab. See 28.52 “To configure an LDP interface” (p. 944) for more information about configuring LDP interfaces.

12

Click on the Peers tab. See 28.54 “To configure an LDP peer” (p. 946) for more information about creating peers.

13

Click on the Targeted Peers tab. See 28.53 “To configure an LDP targeted peer” (p. 945) for more information about creating and configuring targeted peers.

14

Click on the Authentication tab, configure the Type and Key parameters. And, assign an authentication keychain to the LDP.

15

Click on the Import Policies tab and configure the required parameters.

16

Click on the Export Policies tab and configure the required parameters.

17

Click on the Static Prefix FECs tab.

Perform the following steps:

1. Click Create. In the StaticFec (Create) form, configure the required parameters.
2. Click OK.
3. Click Yes.

18

Click on the LSP BFD tab.

Perform the following steps:

1. Click Create. In the LSP BFD(Create) form, configure the required parameters.

   The Prefix List and BFD Template policies attached to the LSP BFD must already exist.
2. Click OK.

19

Click on the Aggregate Prefix Match tab and configure the required parameters.

When the Aggregate Prefix Match Enabled parameter is enabled, the LDP installs a prefix binding in the LDP FIB by performing a longest match against an aggregate prefix in the routing table, as opposed to requiring an exact match of the prefix. The LDP prefix binding continues to be advertised on a per individual /32 prefix basis.

The Administrative State parameter can only be configured if the Aggregate Prefix Match Enabled parameter is enabled.

20

Configure the required parameters in the Prefix Exclude Policies panel.

The Prefix Exclude Policies can only be configured if the Aggregate Prefix Match Enabled parameter is enabled.

When routing policy statements are specified in the Prefix Exclude Policies panel, any prefixes defined in the routing policies are excluded from the aggregate prefix matching procedure. In this case, when LDP receives an FEC-label binding for this prefix, it performs an exact match of a specific FEC element prefix, as opposed to a longest match of one or more LDP FEC element prefixes.

21

Click on the Accounting tab. Click Create. In the AccountingFecPrefix (Create) form, configure the FEC Prefix parameter.

22

In the Egress Accounting Statistics panel, click Select and choose the required accounting policy. Only accounting policies with the CombinedLdpLspEgressStats statistics type are available for selection. Click OK and configure the required parameters.

See the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting statistics.

23

Save your changes and close the form.

END OF STEPS

28.52 To configure an LDP interface

Note: You can choose to inherit values from the global-level LDP configuration by selecting the Inherit Value parameter.

28.52.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→LDP.

2

Right-click on the LDP icon and choose Create Interface.

3

Select an interface in the Interface panel.

4

Configure the required general parameters.

5

Click on the Protocol Properties tab and configure the required parameters on the available tabs.

By default, the parameter values are inherited by all LDP interfaces.

For the Address Type parameter:

- Choose System to have the system IP address set up LDP sessions.

- Choose Interface to have the IP interface address set up LDP sessions only if there are not multiple interfaces between the two neighbors.

6

If you set the Local LSR ID parameter to Interface, click Select in the Interface panel and select a routing interface.

For LDP interfaces on the 7210 SAS, the Interface panel does not appear when you set the Local LSR ID parameter to Interface. The address of the interface you are configuring is used by default.

7

Save your changes and close the form.

8

Configure the protocol for the far-end device, if applicable. Use CLI for devices that are managed outside the scope of the NFM-P.

END OF STEPS

28.53 To configure an LDP targeted peer

Note: You can choose to inherit values from the global-level LDP configuration by selecting the Inherit Value parameter.

The parameters that you configure for a targeted peer take precedence over the parameters that are configured for the global-level LDP configuration.

28.53.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→LDP.

2

Right-click on the LDP icon and choose Properties.

3

Configure a targeted LDP peer.

Perform the following steps:

1. Click on the Targeted Peers tab and click on Create or choose a targeted peer and click Properties. The Targeted Peer (Create|Edit) form opens.

   Alternatively, you can right-click on LDP and choose Create Targeted Peer. To configure an existing peer, you can expand LDP→Targeted Peers and right-click on a targeted peer and choose Properties.
2. Configure the required general targeted peer parameters.

4

Configure the protocol properties of the targeted peer.

Perform the following steps:

1. Click on the Protocol Properties tab and configure the required protocol parameters.
2. Choose a local LSR ID in the Local LSR ID panel.
3. Configure the required parameters in the Timer Control panel.

   By default, these parameter values are inherited from the global-level LDP instance.

   If the Enable Hello Reduction parameter is disabled while the targeted LDP session to the peer is operationally up, the change in the hold time will take effect only when the targeted LDP session is re-established.

   The parameters are not configurable if you are modifying an auto-created LDP targeted peer that was created using a CLI template. You can determine whether a targeted peer was automatically created on the NE using a template by verifying the Creator indicator.

5

Associate LSP with a targeted peer.

Click on the LSPs tab and click on Add. Choose an LSP.

Click on the Multicast LSPs tab and click on Add. Choose an LSP.

The Multicast LSPs tab is available on enabling the Multicast Tunneling Enabled parameter in the Protocol Properties tab.

Note: If you are associating a dynamic LSP with a targeted peer, the destination of the LSP and targeted peer must be the same.

See Chapter 31, “MPLS” for information on how to create and configure LSPs.

6

Save your changes and close the form.

7

Configure the protocol for the far-end device, if applicable. Use CLI for devices that are managed outside the scope of the NFM-P.

END OF STEPS

28.54 To configure an LDP peer

28.54.1 Purpose

LDP performs the label distribution only in MPLS environments. The LDP operation begins with a hello discovery process to find LDP peers in the network. LDP peers are two LSRs that use LDP to exchange label or FEC mapping information. An LDP session is created between LDP peers. A single LDP session allows each peer to learn the label mappings (LDP is bi-directional) of the other peer and to exchange label binding information.

28.54.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→LDP.

2

Right-click on the LDP icon and choose Properties.

3

Configure an LDP peer.

Perform the following steps:

1. Click on the Peers tab and click on Create or choose a peer and click Properties.
2. Configure the required general LDP peer parameters.

   To use LDP DoD, the DoD Label Distribution parameter must be enabled for both peers. If the DoD Label Distribution parameter is not enabled for a peer, LDP DU is used.
3. Configure the required parameters in the FEC panel to specify the number of LDP FECs that the LSR accepts from a given peer and adds into the LDP label database before the NFM-P raises an alarm.

   The LDP service must be up, and a session established between peers before the NFM-P raises an alarm. If the FEC value is above the Threshold parameter value, an implicitly cleared warning is raised. When the value falls below the threshold, the alarm is removed. If the FEC value goes above the FEC Limit parameter value, the NFM-P raises an implicitly cleared major alarm. When the value falls below the limit, but remains above the threshold, the MaxFECLimitReached alarm is cleared.

   You cannot configure the Log Only parameter if the FEC Limit parameter is set to 0.

   You cannot configure the Threshold parameter if the FEC Limit parameter is set to 0.

4

To assign an authentication keychain to the LDP peer, click on the KeyChain tab and select an LDP peer keychain.

5

Click on the Authentication tab and configure the Type and Key parameters.

6

Click on the Import Policies tab and choose up to 5 import policies.

The FEC prefix Import Policy provides a mean of controlling which FEC prefixes received from other LDP and T-LDP peers should be re-distributed to this LDP peer. If no policy name is specified, the NE imports all of the FEC prefixes it learns from other LDP and T-LDP peers to this LDP peer.

7

Click on the Export Policies tab and configure up to 5 export policies.

The FEC prefix Export Policy provides a mean of controlling which FEC prefixes from this peer should be re-distributed to all other LDP and T-LDP peers. If no policy name is specified, the NE will export all FEC prefixes it learns from this LDP peer to all other LDP and T-LDP peers.

8

Save your changes and close the forms.

END OF STEPS

28.55 To configure ECMP for LDP routing

28.55.1 Purpose

The NFM-P supports LDP ECMP when LDP routing is configured on supporting NEs.

For supporting 7210 SAS NEs, resources for LDP ECMP must be allocated using the system resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

28.55.2 Steps

1

In the navigation tree Routing view, expand *NE*→Routing Instance→LDP.

2

Right-click on an LDP icon and choose Properties. The LDP Routing Instance (Edit) form opens.

3

Click on the OAM tab, then the ECMP sub-tab, and set the Tree Trace parameter to Configured.

4

Configure the required tree discovery configuration parameters.

5

Click on the Discovered FECs tab.

6

Choose an FEC and click Properties. An Autodiscovered FEC (Edit) form opens.

7

Perform one of the following:

a. Click Create LDP Tree Trace to configure and execute an LDP tree trace. See Chapter 90, “OAM diagnostic tests” for information about how to configure this test.

b. Click Create LSP Ping to configure and execute an LSP ping test. See Chapter 90, “OAM diagnostic tests” for information about how to configure this test.

c. Click Create LSP Trace to configure and execute an LSP trace test. See Chapter 90, “OAM diagnostic tests” for information about how to configure this test.

8

Click on the Discovered ECMP tab and choose an ECMP path.

9

Click Properties. An Autodiscovered Path (Edit) form opens.

10

Click Create LSP. See Chapter 90, “OAM diagnostic tests” for information about how to configure and execute an LSP trace test.

11

Click Create LSP Ping. See Chapter 90, “OAM diagnostic tests” for more information about how to configure and execute an LSP ping test.

12

Save your changes and close the form.

END OF STEPS

28.56 To view the LDP session information

28.56.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→LDP.

2

Right-click on the LDP icon and choose Properties. The LDP - Routing Instance (Edit) form opens.

3

Click on the Sessions tab.

4

To view the information for a specific session, choose the session entry and click Properties. The Session (View) form opens.

5

Click on the Session tab to view specific session and connection information.

6

View information about the addresses exchanged in LDP peer sessions.

Perform the following steps:

1. Click on the Addresses tab. A list of the addresses received from peers is displayed on the Received tab.
2. To view the information about a received address, choose the address entry and click Properties. The Received Addresses (View) form opens.
3. To view the information about the addresses sent to peers, click on the Sent tab.
4. To view the information about a sent address, choose the address entry and click Properties. The Sent Addresses (View) form opens.

7

Save your changes and close the form.

END OF STEPS

IS-IS configuration workflow and procedures

28.57 IS-IS configuration overview

28.57.1 Overview

Configuration planning is essential to organize devices in level 1, level 2 and level 1 and 2 areas. IS-IS provides defaults for basic protocol operability. Use the NFM-P to configure the IS-IS parameters. The following figure shows the IS-IS view in the navigation tree

Figure 28-7 IS-IS view

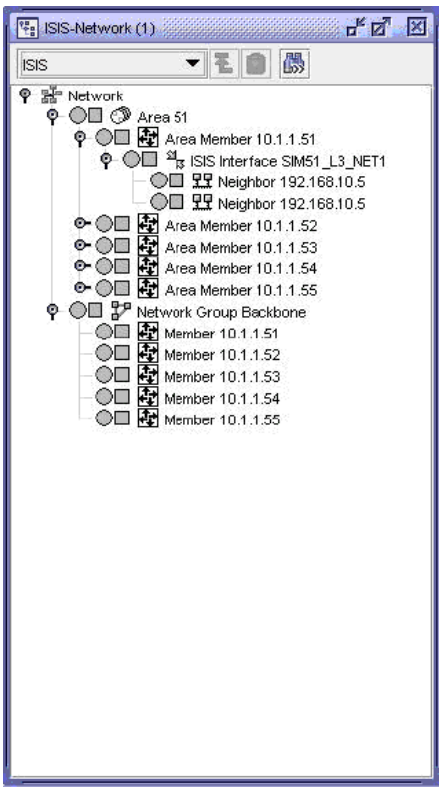

The IS-IS information that is displayed includes:

- backbone area and backbone-area devices
- a list of areas and the participating routers in each area
- a list of devices and associated interfaces that shows the IP address and the configured level (1, 2, or 1 and 2) of each interface
- the IS-IS adjacencies for each interface

The following procedures describe how to configure IS-IS.

28.58 Workflow to configure IS-IS

28.58.1 Stages

1

Enable IS-IS on a routing instance; see 28.59 “To enable IS-IS on a routing instance” (p. 952) .

2

Configure global-level IS-IS on a routing instance; see 28.60 “To configure IS-IS on a routing instance” (p. 953) .

3

Configure IS-IS segment routing as required; see 28.79 “To configure IS-IS segment routing” (p. 978) .

4

Configure an IS-IS link group on a routing instance as required; see 28.61 “To configure an IS-IS link group on a routing instance” (p. 956) .

5

Configure at least one IS-IS NET address; see 28.62 “To configure an IS-IS NET address” (p. 957) .

6

Configure at least one IS-IS interface; see 28.63 “To configure an IS-IS interface” (p. 958) .

7

If LDP over RSVP is enabled for IS-IS, configure an operational LSP between routers; see Chapter 31, “MPLS” .

28.59 To enable IS-IS on a routing instance

28.59.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Protocols tab and select the IS-IS Enabled parameter.

4

Click OK.

5

Click Yes. IS-IS is displayed in the list of enabled protocols.

6

Save your changes and close the form. The ISIS Instance object is displayed in the navigation tree.

END OF STEPS

28.60 To configure IS-IS on a routing instance

28.60.1 Purpose

Routing instance IS-IS parameter values can differ from the IS-IS interface parameter values that are configured in 28.63 “To configure an IS-IS interface” (p. 958) . Interface capabilities are compared to the router-wide capabilities to determine the type of level 1, level 2, and level 1 and 2 adjacencies that are set up between routers to exchange IS-IS routing information.

28.60.2 Steps

1

Enable the IS-IS protocol on a routing instance, as described in 28.59 “To enable IS-IS on a routing instance” (p. 952) .

2

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances.

3

Right-click on the ISIS Instances object and choose Create ISIS Instance, or right-click on an ISIS Instance object and choose Properties. The ISIS Site, Routing Instance, ISIS Instance (Create|Edit) form opens.

4

Configure the required general IS-IS instance parameters.

A level 1 adjacency can be established when there is at least one area ID shared by this router and a neighboring router. A level 2 adjacency is established when another router is configured as a level 2 or a level 1 and 2 router with interfaces configured as level 2 or level 1 and 2. A level 1 and 2 adjacency is created when the neighboring router is also configured as a level 1 and 2 router and the routers have at least one area ID in common.

5

Click on the Behavior tab and configure the required parameters in the General panel.

Note the following:

- You can configure the IPv6 Unicast Multi-Topology parameter only when the Multi-Topology parameter is set to true.
- If you plan to enable IPv6 on the IS-IS instance (see Step 8), you must set the Multi-Topology and IPv6 Unicast Multi-Topology parameters to true.
- You can configure the IID TLV parameter only when the L1 MAC Address and L2 MAC Address parameters are modified from the default values.

6

Click on the LFA tab to configure the required Loop-free Alternates parameters.

Remote loop-free alternate uses segment routing by default. To configure segment routing, see 28.79 “To configure IS-IS segment routing” (p. 978).

7

Configure the required parameters on the Graceful Restart tab.

8

Configure the required parameters to specify the allowed IP versions for the IS-IS instance on the IP Versions tab.

9

Click Apply.

10

Configure the authentication behavior:

Perform the following steps:

1. Click on the Authentication tab and enable authentication.
2. Specify the type of authentication and a key.
3. Select an authentication key chain. See the *NSP System Administrator Guide* for information about how to create a key chain.

11

Configure route preferences and authentication for level 1 or level 2:

Perform the following steps:

1. Click on the Level 1 or Level 2 tab and configure the required parameters.
2. Click on the Authentication tab and configure the required parameters.

12

Select up to five routing policy statements on the Export Policies tab.

You can specify multiple routing policies, in order of preference, to determine the routes that are exported from the routing table to IS-IS. When multiple policies are specified, the policies are evaluated in numerical order.

13

Select up to five routing policy statements on the Import Policies tab.

You can specify multiple routing policies, in order of preference, to determine the routes that are imported to IS-IS. When multiple policies are specified, the policies are evaluated in numerical order.

14

Select up to five loop-free alternate policy statements on the LFA tab.

15

Configure participation in one or more flexible algorithms as needed:

1. Click on the Flexible Algorithms tab.
2. Click Search to populate a list of flexible algorithms.
3. Choose a flexible algorithm and click Properties, or click Create. The Flexible Algorithm (Create|Edit) form opens.
4. Configure the Algorithm ID and Participate parameters.
5. Select a local Flexible Algorithm Definition if needed.
6. Click OK.

Repeat this step to add additional flexible algorithms as needed.

16

Configure one or more IGP shortcuts as needed:

1. Click on the IGP Shortcut tab.
2. Choose a hop and click Properties. The IGP Shortcut Tunnel Next Hop (Edit) form opens.
3. Configure the parameters.
4. Click OK to close the form.

17

Configure a NET address:

Perform the following steps:

1. Click on the NET Addresses tab.
2. Click Create and configure the Area ID parameter on the Area ID, ISIS Instance (Create) form.

3. Save your changes and close the form.

18

Configure a route summarization. IS-IS route summaries allow users to create aggregate IPv4 or IPv6 addresses that include multiple groups of IPv4 or IPv6 addresses for a specific IS-IS summary level. This can help reduce the size of the link state database and the routing table.

Perform the following steps:

1. Click on the Route Summarizations tab and click Create. The Route Summarization (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

19

To add an IS-IS interface:

Perform the following steps:

1. Click on the Interfaces tab and click Create. The ISIS Interface (Create) form opens.
2. Select an interface in the Interface panel.
3. Click OK and confirm.
4. Perform Step 5 to Step 11 of 28.63 “To configure an IS-IS interface” (p. 958) .

20

Save your changes and close the form.

END OF STEPS

28.61 To configure an IS-IS link group on a routing instance

28.61.1 Purpose

You can create an IS-IS link group to bundle multiple member interfaces that are treated as a single virtual link for ECMP. See 28.60 “To configure IS-IS on a routing instance” (p. 953) and 28.63 “To configure an IS-IS interface” (p. 958) for more information about how to create and configure IS-IS interfaces.

28.61.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances→ISIS Instance.

2

Right-click on an ISIS Instance object and choose Properties. The ISIS Routing Instance, ISIS Instance (Edit) form opens.

3

Create a link group for an ISIS routing instance:

Perform the following steps:

1. Click on the Link Group tab and click Create. The Link Group Config (Create) form opens.
2. Configure the required parameters.

4

Configure Level 1 or Level 2 for the link group:

Perform the following steps:

1. Click on the Level 1 or Level 2 tab. The General sub-tab is displayed.
2. Configure the required parameters.

5

To add a member to the Level 1 or Level 2 link group:

Perform the following steps:

1. Click on the Members tab and click Add. The Select... form opens.
2. Choose an interface and click OK.

6

Save your changes and close the form.

END OF STEPS

28.62 To configure an IS-IS NET address

28.62.1 Steps

1

Enable the IS-IS protocol on a device, as described in 28.59 “To enable IS-IS on a routing instance” (p. 952) .

2

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances→ISIS Instance.

3

Right-click on the ISIS Instance object and choose Add NET Address. The Area ID (Create) form opens.

4

Configure the Area ID parameter.

Note: The Area ID forms part of the NET address. NET addresses are built from some non-configurable elements, including the device system ID and the network selector ID.The NET address is exchanged in hello and link-state PDUs. Level 1 interfaces must have at least one area ID in common. Level 2 interfaces can have different area IDs. If all of the interfaces have different area IDs, they are considered level 2 interfaces only.

5

Save your changes and close the form.

END OF STEPS

28.63 To configure an IS-IS interface

28.63.1 Purpose

The IS-IS interface parameters can differ from the IS-IS routing instance parameters configured in 28.60 “To configure IS-IS on a routing instance” (p. 953) . Interface parameters specify the interface routing levels. Interface level capabilities are compared to the router-wide level capabilities to determine the type of level 1, level 2, and level 1 and 2 adjacencies that are created between devices.

28.63.2 Steps

1

Enable the IS-IS protocol on a routing instance, as described in 28.59 “To enable IS-IS on a routing instance” (p. 952) . If no IS-IS instance exists, you must create one by performing 28.60 “To configure IS-IS on a routing instance” (p. 953) .

2

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances→ISIS Instance.

3

Right-click on the ISIS Instance object and choose Create Interface or right-click on an ISIS Interface object and choose Properties. The ISIS Interface (Create|Edit) form opens.

4

If you are creating an IS-IS interface, select an interface in the Interface panel.

5

Configure the required parameters on the General tab.

6

Click on the Behavior tab and configure the required parameters.

To create a mesh group, set the Mesh Group Status parameter to Enabled and specify the same mesh group value for all interfaces. The mesh group parameters specify the assigned mesh group for the interface. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised in an area.

The TE Metric parameter is a read-only field that displays the TE metric used for LSP path computation by CSPF.

The Default Instance parameter is configurable on interfaces on non-zero IS-IS instances.

7

Click on the Segment Routing tab and configure the required parameters.

You cannot configure the IPv4 or IPv6 SID Type or SID Value parameter on a network interface that is not loopback enabled or that is on a service interface. In the navigation tree Routing view, expand Network→NE→Routing Instance→Interface, right-click on the interface and choose Properties. Configure the Loopback Enabled parameter on the Network Interface (Edit) form.

The IPv4 or IPv6 SID Value range depends on the Prefix Start Label and Prefix Map Index configured in the IS-IS instance, see 28.79 “To configure IS-IS segment routing” (p. 978). The IPv6 SID value cannot be the same as the IPv4 SID.

8

Configure the required parameters on the Flexible Algorithms tab.

You cannot configure the IPv4 or IPv6 SID Type or SID Value parameter on a network interface that is not loopback enabled or that is on a service interface. In the navigation tree Routing view, expand Network→NE→Routing Instance→Interface, right-click on the interface and choose Properties. Configure the Loopback Enabled parameter on the Network Interface (Edit) form.

The IPv6 SID value cannot be the same as the IPv4 SID.

1. Click Search to populate a list of flexible algorithms.
2. Choose a flexible algorithm and click Properties, or click Create. The Interface Flexible Algorithm (Create|Edit) form opens.
3. Configure the Flexible Algorithm ID parameter.
4. Configure the IPv4 and IPv6 SID parameters as needed.
5. Click OK.

Repeat this step to add additional flexible algorithms as needed.

9

Configure the required parameters on the Authentication and Policies tabs.

10

Configure the route preferences and authentication for level 1 or level 2:

Perform the following steps:

1. Click on the Level 1 or Level 2 tab and configure the required parameters on the General sub-tab.
2. Click on the Authentication sub-tab and configure the required parameters.

11

Save your changes and close the form.

END OF STEPS

OSPF configuration workflow and procedures

28.64 OSPF configuration overview

28.64.1 Overview

The NFM-P supports the configuration of OSPFv2 and OSPFv3 on supporting NEs.

Configuration planning is essential to organize OSPF areas, interfaces, and virtual links. OSPF provides defaults for basic protocol operability. OSPF configuration requires, as a minimum, that:

- you create a single OSPF backbone area that contains the area border routers.
- for larger networks, you create several areas that contain the other routers.
- for smaller networks, you place all routers in the OSPF backbone area.

Use the NFM-P navigation tree to configure the OSPF parameters. The following figure shows the OSPF view in the navigation tree.

Figure 28-8 OSPF view

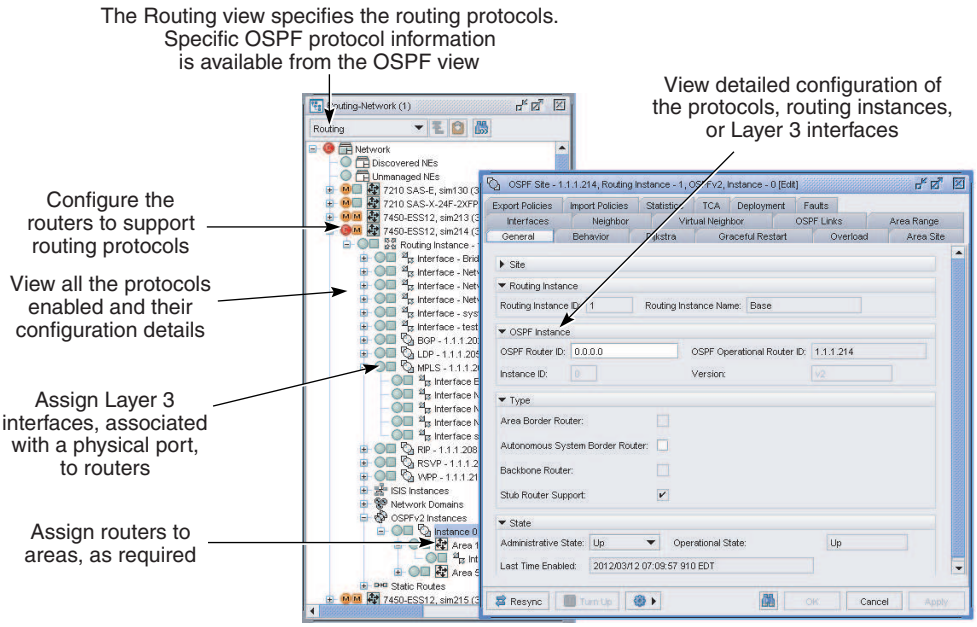

22868

The OSPF parameters that are required for OSPF deployment are:

- Router ID — Each device that runs OSPF must be configured with a unique router ID. The router ID is used by the OSPF and BGP routing protocols in the routing table manager. When you configure a new router ID, protocols are not automatically restarted with the new router ID. You must shut down and restart the protocol to initialize the new router ID.

- An area — At least one OSPF area must be created. An interface must be assigned to each OSPF area. The types of OSPF areas include a backbone area, stub area, and NSSA.
- Layer 3 interfaces — A Layer 3 interface is the logical IP connection between a router and one of its attached networks. A physical interface is associated with the Layer 3 interface to provide the cabled connected to another device. A Layer 3 interface has state information from the underlying lower-level protocols and the routing protocol. A network interface has an associated IP address and mask that combine to create an IP prefix, unless the interface is in an unnumbered, point-to-point network.

The following procedures describe how to configure OSPF.

28.65 Workflow to configure OSPFv2 and OSPFv3

28.65.1 Stages

1

Enable OSPF on a routing instance; see 28.66 “To enable OSPF on a routing instance” (p. 963) .

2

Create at least one OSPFv2 or OSPFv3 area; see 28.67 “To create an OSPF area” (p. 963) .

3

Create an OSPF neighbor on an OmniSwitch as required; see 28.68 “To create an OSPF neighbor on an OmniSwitch” (p. 965) .

4

Assign Layer 3 interfaces to the routers in the OSPFv2 or OSPFv3 area; see 28.69 “To add a Layer 3 interface to an OSPF router” (p. 965) .

5

Create an OSPF area range as required; see 28.70 “To create an OSPF area range” (p. 967) .

6

Create a virtual link between the backbone OSPF area and a remote OSPF area that does not advertise the OSPF topology as required; see 28.71 “To create a virtual link between OSPF areas” (p. 968) .

7

Configure OSPF on a default routing instance or a VRF routing instance; see 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969) .

8

Configure OSPF segment routing as required; see 28.80 “To configure OSPF segment routing” (p. 980).

9

Assign routers to the OSPFv2 or OSPFv3 area; see 28.74 “To add a router to an OSPF area” (p. 972) .

28.66 To enable OSPF on a routing instance

28.66.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Protocols tab.

4

Select the OSPFv2 Enabled parameter to enable OSPFv2, if required.

5

Select the OSPFv3 Enabled parameter to enable OSPFv3, if required.

6

Save your changes and close the form.

END OF STEPS

28.67 To create an OSPF area

28.67.1 Steps

1

In the navigation tree OSPF view, right-click on the Network icon and choose Create Area. The Area (Create) form opens.

2

Configure the required general parameters and click Apply. The form refreshes with additional tabs.

3

Click on the Area Site tab and click Create to add a routing instance to the area. The Area Site (Create) form opens.

4

Click Select to choose a routing instance.

5

Configure the required parameters and click Apply. Additional configurable tabs appear, and an icon for the new area appears in the navigation tree.

Configure the area as a blackhole range to avoid routing loops.

The Type parameter is not configurable for a backbone-area routing instance.

6

Depending on the Type parameter value, the Stub/NSSA tab is configurable. Configure stub or NSSA functionality, if required.

The Default Cost parameter is configurable only when the Type parameter value in Step 5 is Stub (No Type 5 External) or Totally Stub (No Summaries).

1. Click on the Stub/NSSA tab.
2. Configure the required parameters.

   The Redistribute External Routes, Originate Default Route, and Adjacency Check parameters are configurable only when the Type parameter value from Step 5 is NSSA (No Type 5 External) or NSSA (No Summaries).

   The Adjacency Check parameter is configurable only when the Originate Default Route is set to Originate Type 3 or Originate Type 7. Enable the Adjacency Check parameter to enforce an ABR to have an adjacency in area 0 before the default route or routes can be advertised into the NSSA.

7

Click on the Import Policies and Export Policies tabs and configure the required parameters. Routing policies determine the routes that are advertised to peers, and are created as routing policy statements. See 54.5 “To configure a routing policy statement” (p. 1745) .

8

If the area is a backbone area that requires links to remote areas that do not advertise OSPF topology, configure a virtual link.

Perform the following steps:

1. Click on the Virtual Link tab.
2. Perform Step 4 to Step 7 of 28.71 “To create a virtual link between OSPF areas” (p. 968) .
3. Save your changes. The new virtual link entry appears in the list.

9

Create an area range:

Perform the following steps:

1. Click on the Area Range tab and click Create. The Area Range (Create) form opens.
2. Perform Step 3 and Step 4 of 28.70 “To create an OSPF area range” (p. 967) .
3. Save your changes. The new area range entry appears in the list.

10

Save your changes and close the form.

END OF STEPS

28.68 To create an OSPF neighbor on an OmniSwitch

28.68.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→OSPFv2.

2

Right-click on the OSPFv2 icon and choose Create Neighbor. The OspfNeighbor (Create) form opens.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

28.69 To add a Layer 3 interface to an OSPF router

28.69.1 Purpose

Perform this procedure to allow an OSPF-enabled router to participate in area discovery and share routing information with other area members.

Note: This action assigns an existing Layer 3 interface to the router in the OSPF area rather than creating a new Layer 3 interface.

28.69.2 Steps

1

In the navigation tree OSPF view, expand Network→Instance.

2

Right-click on an Instance icon and choose Create Interface. The OSPF Interface (Create) form opens.

3

Configure the required general parameters.

4

Choose an interface in the Interface panel.

5

Click on the Protocol Properties tab and configure the required parameters.

6

Perform one of the following steps to configure authentication for the interface. Otherwise, go to Step 10 .

a. Click on the Authentication tab, and go to Step 7 .

b. For OSPFv3 interfaces, go to Step 8 .

7

Configure the Authentication Type parameter.

a. Choose MD5-based Authentication.

Perform the following steps:

1. Click Create to create an MD5 authentication key. The Md5Key (Create) form opens.
2. Configure the required parameters and save your changes.

b. Choose Simple Password.

Perform the following steps:

1. Click Change Password to enter a password. The Password (Create) form opens.
2. Configure the required parameters and save your changes.

8

Click on the IPsec Static SA tab and configure the required parameters.

9

Click on the Policies tab and configure the required parameters.

10

Click on the Segment Routing tab and configure the required parameters.

You cannot configure the IPv4 SID Type or SID Value parameter on a network interface that is not loopback enabled or that is on a service interface. In the navigation tree Routing view, expand Network→NE→Routing Instance→Interface, right-click on the interface and choose Properties. Configure the Loopback Enabled parameter on the Network Interface (Edit) form.

The IPv4 SID Value range depends on the Prefix Start Label and Prefix Map Index configured in the OSPF instance, see 28.80 “To configure OSPF segment routing” (p. 980) .

11

Save your changes and close the form.

END OF STEPS

28.70 To create an OSPF area range

28.70.1 Purpose

An area range summarizes a range of IP addresses in an LSA to minimize the number of flooded advertisements in the LSA.

28.70.2 Steps

1

In the navigation tree OSPF view, expand Network→Instance.

2

Right-click on an Instance icon and choose Create Area Range. The Area Range (Create) form opens.

3

Configure the required parameters.

4

Save your changes and close the form. An icon for the new area range appears in the navigation tree below the OSPF area.

END OF STEPS

28.71 To create a virtual link between OSPF areas

28.71.1 Purpose

Perform this procedure to create a virtual link between a backbone OSPF area and a remote OSPF area that does not advertise the OSPF topology. When you create a virtual link, a virtual neighbor is automatically created.

28.71.2 Steps

1

On the NFM-P navigation tree, choose OSPF from the view selector.

2

Click on the backbone area icon in the navigation tree to display the area routers. The backbone area ID is 0.0.0.0.

3

Right-click on a backbone-area router instance icon and choose Create Virtual Link. The Virtual Link (Create) form opens.

4

Configure the general parameters.

5

Click on the Protocol Properties tab and configure the required parameters.

6

Perform one of the following steps to configure authentication for the virtual link. Otherwise, go to Step 9 .

a. Click on the Authentication tab, and go to Step 7 .

b. For OSPFv3 interfaces, go to Step 8 .

7

Configure the Authentication Type parameter.

a. Choose MD5-based Authentication.

   Perform the following steps:

   1. Click Create to create an MD5 authentication key. The Md 5 Key (Create) form opens.
   2. Configure the required parameters.
   3. Save your changes and go to Step 9 .

b. Choose Simple Password.

   Perform the following steps:

   1. Click Change Password to enter a password. The Password (Create) form opens.
   2. Configure the required parameters.
   3. Save your changes and go to Step 9 .

8

Click on the IPsec Static SA tab and configure the required parameters.

9

Save your changes and close the form. The Virtual Link (Create) form closes, and the NFM-P displays two new icons: one for the virtual link and one for the virtual neighbor.

10

As required, review the properties of the virtual link or virtual neighbor, by clicking on the appropriate icon and selecting Properties.

Note: You can delete a virtual link by selecting the virtual link icon and clicking Delete. To delete a virtual neighbor, you must select the virtual link icon and click Delete.

END OF STEPS

28.72 To configure OSPF on a default routing instance or a VRF routing instance

28.72.1 Steps

1

In the navigation tree Routing view, perform one of the following:

a. For NEs that support multiple OSPF instances, expand Network→*NE*→Routing Instance→OSPFv*N* Instances, where *N* is the OSPF version number; for example, 2 or 3.

b. For all other supported devices, expand Network→*NE*→Routing Instance→OSPFv*N*, where *N* is the OSPF version number; for example, 2 or 3.

2

Perform one of the following:

a. For NEs that support multiple OSPF instances, right-click on the appropriate OSPFv*N* Instances icon and choose Create OSPFv*N* Instance. The OSPF Site (Create) form opens. Go to Step 3 .

b. For all other supported devices, right-click on the appropriate OSPFv*N* icon and choose Properties. The OSPF Site, Routing instance, OSPFv*N* instance (Edit) form opens. Go to Step 3 .

3

Configure the required general parameters.

4

Click on the Behavior tab and configure the required parameters.

Some parameters are configurable for OSPFv2 only.

Remote loop-free alternate uses segment routing by default. To configure segment routing, see 28.80 “To configure OSPF segment routing” (p. 980).

5

Click on the Dijkstra tab and configure the required parameters:

For SPF:

- SPF Max Wait (milliseconds)
- Redistribute Delay (milliseconds)
- Initial Wait (milliseconds)
- Second Wait (milliseconds)
- Incremental SPF Wait (milliseconds)

For LSA:

- LSA Generate Max Wait (milliseconds)
- LSA Arrival Wait (milliseconds)
- Initial Wait (milliseconds)
- Second Wait (milliseconds)
- LSA Accumulate (milliseconds)

6

Click on the Graceful Restart tab and configure the required parameters.

7

Depending on the OSPF version, the Overload tab is configurable. Click on the Overload tab and configure the required parameters.

The Overload Interval (seconds) parameter is configurable when the Overload Enabled parameter is selected.

The Boot Overload Interval (seconds) parameter is configurable when the Boot Overload Enabled parameter is enabled.

8

Click on the LFA tab and configure the required parameters.

Additional parameters are available when you select the Loop-free Alternate parameter.

9

Configure participation in one or more flexible algorithms as needed:

1. Click the Flexible Algorithms tab.
2. Click Search to populate a list of flexible algorithms.
3. Choose a flexible algorithm and click Properties, or click Create. The Flexible Algorithm (Create|Edit) form opens.
4. Configure the Algorithm ID and Participate parameters.
5. Select a local Flexible Algorithm Definition if needed.
6. Click OK.

Repeat this step to add additional flexible algorithms as needed.

10

If the OSPF version is OSPFv2, the OSPF Super-Backbone tab is configurable. Click on the OSPF Super-Backbone tab and configure the required parameters.

The following parameters apply only to VPRN instances of OSPF:

- VPN Domain Type
- VPN Domain ID (hex)
- VPN Tag
- Super-Backbone

11

Click on the Import Policies and Export Policies tabs and configure the required parameters. Routing policies determine the routes that are advertised to peers, and are created as routing policy statements. See 54.5 “To configure a routing policy statement” (p. 1745) .

12

Select up to five loop-free alternate policy statements on the LFA Policies tab.

13

To configure segment routing, see 28.80 “To configure OSPF segment routing” (p. 980)

14

Click on the following tabs to add optional OSPF configuration items. Click Search to list previously configured items or click Create to create the item, if supported.

- Area Sites; see 28.67 “To create an OSPF area” (p. 963)
- Interfaces (Layer 3 interfaces); see 28.69 “To add a Layer 3 interface to an OSPF router” (p. 966)
- Area Range; see 28.70 “To create an OSPF area range” (p. 967)

- Virtual Neighbor; see 28.71 “To create a virtual link between OSPF areas” (p. 968)
- OSPF Links (configurable only in a backbone area); see 28.71 “To create a virtual link between OSPF areas” (p. 968)
- IGP Shortcut; see 28.73 “To configure an IGP shortcut on an OSPF instance” (p. 971)

15

Save your changes and close the form.

END OF STEPS

28.73 To configure an IGP shortcut on an OSPF instance

28.73.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→OSPFv*N* Instances, where *N* is the OSPF version number; for example, 2 or 3.

2

Right-click on the appropriate OSPFv*N* Instance icon and choose Properties. The OSPF Site (Edit) form opens.

3

Click on the IGP Shortcut tab.

4

Choose a hop and click Properties. The IGP Shortcut Tunnel Next Hop (Edit) form opens.

5

Configure the parameters.

6

Save your changes and close the forms.

END OF STEPS

28.74 To add a router to an OSPF area

28.74.1 Steps

1

In the navigation tree OSPF view, expand Network→Routing→OSPF Area.

2

Right-click on the OSPF area icon and choose Add OSPF Instance. The Area Site (Create) form opens.

3

Configure the required parameters.

Configure the area as a blackhole range to avoid routing loops.

4

Select a Routing Instance in the Routing Instance panel.

5

Depending on the Type parameter value from Step 3 , the Stub/NSSA tab is configurable. Click on the Stub/NSSA tab and configure the required parameters.

The Default Cost parameter is configurable only when the Type parameter value from Step 3 is Stub (No Type 5 External) or Totally Stub (No Summaries).

The Redistribute External Routes, Originate Default Route, and Adjacency Check parameters are configurable only when the Type parameter value from Step 3 is NSSA (No Type 5 External) or NSSA (No Summaries).

The Adjacency Check parameter is configurable only when the Originate Default Route is set to Originate Type 3 or Originate Type 7. Enable the Adjacency Check parameter to enforce an ABR to have an adjacency in area 0 before the default route or routes can be advertised into the NSSA.

6

Save your changes and close the form. An icon for the router appears in the navigation tree under the OSPF Area icon.

END OF STEPS

28.75 To configure an OSPF interface

28.75.1 Purpose

The OSPF interface parameters can differ from the OSPF routing instance parameters configured in 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969) . Interface parameters specify the interface routing levels. Interface level capabilities are compared to the router-wide level capabilities to determine the type of level 1, level 2, and level 1 and 2 adjacencies that are created between devices.

28.75.2 Steps

1

Enable the OSPF protocol on a routing instance, as described in 28.66 “To enable OSPF on a routing instance” (p. 963) . If no OSPF instance exists, you must create one by performing 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969) .

2

In the navigation tree Routing view, perform one of the following:

a. For NEs that support multiple OSPF instances, expand Network→*NE*→Routing Instance→OSPFv*N* Instances, where *N* is the OSPF version number; for example, 2 or 3.

b. For all other supported devices, expand Network→*NE*→Routing Instance→OSPFv*N*, where *N* is the OSPF version number; for example, 2 or 3.

3

Perform one of the following:

a. For NEs that support multiple OSPF instances, right-click on the appropriate OSPFvN Instances icon and choose Create OSPFvN Interface. The OSPF Interface (Create) form opens. Go to Step 4.

b. For all other supported devices, right-click on the appropriate OSPFvN icon and choose Properties. The OSPF Interface, Routing instance, OSPFvN instance (Edit) form opens. Go to Step 4.

4

Configure the required parameters on the Segment Routing tab.

You cannot configure the SID Type or SID Value parameter on a network interface that is not loopback enabled or that is on a service interface. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→Interface, right-click on the interface and choose Properties. Configure the Loopback Enabled parameter on the Network Interface (Edit) form.

5

Configure the required parameters on the Flexible Algorithms tab.

You cannot configure the SID Type or SID Value parameter on a network interface that is not loopback enabled or that is on a service interface. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→Interface, right-click on the interface and choose Properties. Configure the Loopback Enabled parameter on the Network Interface (Edit) form.

1. Click Search to populate a list of flexible algorithms.
2. Choose a flexible algorithm and click Properties, or click Create. The Interface Flexible Algorithm (Create|Edit) form opens.
3. Configure the Flexible Algorithm ID parameter.
4. Click OK.

Repeat this step to add additional flexible algorithms as needed.

6

Configure the required parameters on the Authentication and Policies tabs.

7

Save your changes and close the form.

END OF STEPS

Segment routing configuration workflow and procedures

28.76 Workflow to configure segment routing

28.76.1 Stages

1

Configure a segment routing policy; see 28.77 “To create a segment routing policy” (p. 976)

2

Configure BGP sites to learn a segment routing policy as needed; see 28.78 “To enable SR policy support on a BGP site, peer, or peer group” (p. 977)

3

Configure segment routing for ISIS and OSPF instances as needed:

- 28.79 “To configure IS-IS segment routing” (p. 978)
- 28.80 “To configure OSPF segment routing” (p. 980)

4

Configure segment routing trees as needed; see 28.81 “To create a segment routing tree” (p. 981)

28.77 To create a segment routing policy

28.77.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance and select Properties.

2

Click on the Segment Routing tab.

3

Select the SR policy entry in the SR Policy tab and click Properties. The Segment Routing Policy form opens.

4

Configure the required parameters in the General tab.

5

To enable ingress statistics collection, perform the following:

1. In the Statistics Configuration panel, click Create. The Segment Routing Policy Statistics Configuration form opens.
2. Set the Administrative State to Up, and configure all other required parameters.
3. Click OK and confirm to close the form. A statistics configuration entry appears in the table.

Note: Statistics are not collected unless there are at least two NEs using the policy.

6

Click on the Static Policies tab. Create segment routing static policies by performing the following:

1. Click Create. The Segment Routing Static Policy form opens.
2. Configure the parameters in the General tab.
3. If you will need to use Seamless BFD, click Select in the Maintenance Policy panel to add a maintenance policy; see 54.18 “To configure a maintenance policy” (p. 1764).
4. Click on the Segment List tab and click Create. The Segment List form opens.
5. Configure the parameters on the General tab.
6. Click on the Segment List Entries tab and click Create. The Segment List Entry form opens.
7. Configure the parameters then click OK and confirm to close the Segment List Entry form. A segment list entry appears in the table.
8. Click OK and confirm to close the Segment List form. A segment list appears in the table.
9. Click on the Segment Routing SRv6 tab, and then on the Binding SID sub-tab.
10. Click Create. The Binding SID form opens.
11. Configure the parameters on the General tab.
12. Click OK and confirm to close the Binding SID form.
13. Click OK and confirm to close the Segment Routing Static Policy form. A static policy appears in the table.

7

Click OK and confirm to close the Segment Routing Policy form.

END OF STEPS

28.78 To enable SR policy support on a BGP site, peer, or peer group

28.78.1 Steps

1

Click on the VPN tab on a BGP site, peer, or peer group configuration form.

2

Enable the SR Policy IPv4 parameter.

3

Click on the Next-Hop Resolution tab.

4

Enable SR policies on tunnel interfaces by performing the following:

1. Select an instance in the Shortcut Tunnel or Label Route Transport Tunnel and click Properties.
2. Enable the SR Policy parameter.
3. Repeat for each tunnel family on which you need to enable SR policies.

5

Click on the Behavior tab and enable the Import Segment Routing Policy parameter.

6

Click OK and confirm to close the form.

END OF STEPS

28.79 To configure IS-IS segment routing

28.79.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances.

2

Right-click on an ISIS Instance object and choose Properties. The ISIS Site, Routing Instance, ISIS Instance (Edit) form opens.

3

Click on the Behavior tab.

4

In the General panel, set the Advertise Router Capability parameter to AS or to Area.

5

Click on the Segment Routing tab.

6

Configure the parameters on the SID Statistics panel.

7

Set the Administrative State parameter to Up and configure the required parameters.

The Prefix SID Start Label is configurable only when the Prefix SID Type parameter is set to Local.

The minimum value you can enter for the Prefix SID Start Label parameter depends on the value of the Start Segment Routing Label parameter on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

8

To support dual adjacency SIDs, configure the parameters in the Adjacency SID panel.

9

Configure the parameters in the Multi-Topology panel.

10

To configure IS-IS/LDP segment routing interworking, perform the following:

1. Set the Export Tunnel Table Protocol attribute to LDP.
2. Click on the Mapping Prefix sub-tab.
3. Click on the Create button. The Mapping Server form opens.
4. Configure the required parameters.
5. Save your changes and close the form.

11

To configure an adjacency set and SID, perform the following:

1. In the Reserved Label Block panel, click on the Select button and choose a reserved label block. See 31.35 “To create a reserved label block” (p. 1163) for information about creating reserved label blocks.
2. In the Adjacency Set panel, click on the Create button. The Segment Routing Adjacency Set form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

12

Save your changes and close the form.

END OF STEPS

28.80 To configure OSPF segment routing

28.80.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→OSPF Instances.

2

Right-click on an OSPF Instance object and choose Properties. The OSPF Site, Routing Instance, OSPF Instance (Edit) form opens.

3

Click on the Behavior tab.

4

In the General panel, set the Advertise Router Capability parameter to AS or to Area.

Note: For OSPFv3, the parameter must be set to Area.

5

Click on the Segment Routing tab.

6

Set the Administrative State parameter to Up and configure the required parameters.

The Prefix SID Start Label is configurable only when the Prefix SID Type parameter is set to Local.

The minimum value you can enter for the Prefix SID Start Label parameter depends on the value of the Start Segment Routing Label parameter on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

7

To support dual adjacency SIDs, configure the parameters in the Adjacency SID panel.

8

To configure OSPF/LDP segment routing interworking, perform the following:

1. Set the Export Tunnel Table Protocol attribute to LDP.
2. Click on the Mapping Prefix sub-tab.
3. Click on the Create button. The Mapping Server form opens.
4. Configure the required parameters.
5. Save your changes and close the form.

9

To configure an adjacency set and SID, perform the following:

1. In the Reserved Label Block panel, click on the Select button and choose a reserved label block. See 31.35 “To create a reserved label block” (p. 1163) for information about creating reserved label blocks.
2. In the Adjacency Set panel, click on the Create button. The Segment Routing Adjacency Set form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

10

Save your changes and close the form.

END OF STEPS

28.81 To create a segment routing tree

28.81.1 Purpose

Use this procedure to create a segment routing multicast tree for P2MP traffic engineering. The segment routing tree must be created on a routing instance of the NE that will serve as the root node.

A reserved label block is required; see 31.35 “To create a reserved label block” (p. 1163).

28.81.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance and select Properties.

2

Click on the Segment Routing tab.

3

Click on the SR-Tree tab and click Create. The Segment Routing Tree form opens.

4

Configure the required parameters in the General tab.

5

Click on the SR-Tree Replication Policies tab. Create segment routing tree replication policies by performing the following:

1. Click Create. The Replication Segment Policy form opens.
2. Configure the parameters in the General tab.
3. Click on the Next Hop Entries tab and click Create. The Replication Segment Policy Next Hop form opens.
4. Configure the parameters.
5. Click OK to close the Replication Segment Policy Next Hop form. A next hop entry appears in the table.
6. Create additional hops as needed.
7. Click OK to close the Replication Segment Policy form. An SR-Tree Replication policy appears in the table.

6

Click on the SR-Tree Policies tab. Create segment routing tree policies by performing the following:

1. Click Create. The Segment Routing Tree Policy form opens.
2. Configure the parameters in the General tab.
3. Click on the Candidate Paths tab and click Create. The Candidate Path form opens.
4. Configure the parameters.

   The Active Instance parameter indicates the replication policy instance ID to be used.
5. Click OK to close the Candidate Path form. A candidate path appears in the table.
6. Create additional candidate paths as needed. The candidate path with the higher preference becomes the active candidate path.
7. Click OK to close the Segment Routing Tree Policy form. An SR-Tree policy appears in the table.

7

Click OK and confirm to close the Segment Routing Tree form.

END OF STEPS

28.82 To configure segment routing with IPv6

28.82.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance and select Properties.

2

Click on the Segment Routing tab.

3

Create or select an entry in the SRv6 tab and click Properties. The Segment Routing V6 (Create/Edit) form opens.

4

Create or select an entry in the Locators tab and click Properties. The SRv6 Locator (Create/Edit) form opens.

5

Configure the required parameters in the General tab.

1. Associate FPE ID under FPE block. See 12.41 “To create an FPE” (p. 374) for more information on FPE.

   Note: The associated Termination FPE must have SRv6 Type set as Termination. Similarly, Origination FPE must be configured with SRv6 Type as Origination.
2. Under Prefix panel, specify the network address for the Prefix Address parameter.
3. Associate a label block under Static Function panel.

   Note: The configuration of FPE ID, Prefix Address, and parameters under Static Function panel are required to set the Administrative State parameter to Up.
4. Save your changes and close the form.

6

On the Segment Routing V6 (Create/Edit) form, configure the required parameters in the General tab.

7

Create or select an entry in the Base Routing Instance panel and click Properties. The SRv6 Function (Create/Edit) form opens.

8

Configure the required parameters in the General tab.

1. Associate a locator under Locator panel.
2. Configure parameters under End Function panel.

   Supported End Function types are Dynamic and Static.

9

Configure the End Function, the End-X Function, and the End -X Allocate Function in respective tabs.

Note: End (End DT4, End DT6, End DT46, End-X) Functions range configuration is dependent on the max entries configured in static function.

10

Save your changes and close the forms.

END OF STEPS

28.83 To configure IS-IS segment routing with IPv6

28.83.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→ISIS Instances.

2

Right-click on an ISIS Instance object and choose Properties. The ISIS Site, Routing Instance, ISIS Instance (Edit) form opens.

3

Click on the Behavior tab.

4

In the General panel, set the Advertise Router Capability parameter to Area.

5

Click on the Segment Routing V6 tab.

6

Create or select an entry in the Locator tab and click Properties. The Segment Routing V6 Locator (Create/Edit) form opens.

1. Associate a locator under Locator panel.
2. Save your changes and close the form.

7

Configure the required parameters in the General tab. Set the Administrative State parameter to Up.

8

Save your changes and close the forms.

END OF STEPS

RSVP configuration workflow and procedures

28.84 Workflow to configure RSVP

28.84.1 Stages

1

Configure MPLS, LSP, MPLS path, and LSP path parameters as required; see Chapter 31, “MPLS” .

2

Configure RSVP on a routing instance; see 28.85 “To configure RSVP on a routing instance” (p. 985) .

3

Configure an RSVP interface; see 28.86 “To configure an RSVP interface” (p. 986) .

28.85 To configure RSVP on a routing instance

28.85.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→RSVP.

2

Right-click on the RSVP icon and choose Properties. The RSVP form opens.

3

Configure the required parameters.

4

To configure Diff-Serv TE, MPLS must be shutdown. Otherwise, go to Step 5 .

Perform the following steps:

1. Configure the required parameters in the Diff-Serv TE panel.

   Note:

   Enabling a Diff Serv Model only takes effect if you have already enabled traffic engineering at one or both of the following routing protocol levels:

   - IS-IS, as described in 28.60 “To configure IS-IS on a routing instance” (p. 953)
   - OSPF, as described in 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969)

A dialog box appears if IS-IS and/or OSPF TE is not enabled, but you can save your configuration if IS-IS and/or OSPF TE is not enabled.

The Class Type BW Percent parameters are configured globally and the parameters are applied to all RSVP interfaces in the system. After the parameters are configured, the parameters can only be changed after you shut down the MPLS protocol. The total bandwidth specified for the eight classes should not exceed 100%.

2. Configure the parameters in the TE Update panel. The TE Update parameters are only configurable if the TE Threshold Update Enabled parameter is enabled in 1 .
3. Configure the Include Node in RRO parameter.
4. Configure the required parameters in the Merge Point Abort Timer and GR Helper Time panels.
5. Configure the Authentication Over Bypass Enabled parameter to support authentication per RSVP interface.
6. Click on the TE Classes tab, select one of the TE classes, and click Properties. The TE Class (Edit) form opens.

   Note:

   The TE Class Definition is used to map all TE Classes (up to a maximum of 8) to Class Type and LSP setup priority. There is no default TE Class after Diff-Serv is enabled. You must explicitly define each TE Class. If Diff-Serv is disabled, the default CT (CT0) and eight pre-emption priorities are used internally.
7. Configure the required parameters and click Apply.
8. Click on the ForwardingClassMaps tab and click Create. The Forwarding Class Map (Create) form opens.
9. Configure the required parameters and click Apply.
10. Click on the TE Thresholds tab and configure the required parameters.

    You can also click Reset TE Up/Down Thresholds to Default to set the Up or Down Thresholds to their default values, respectively.

5

Save your changes and close the form.

END OF STEPS

28.86 To configure an RSVP interface

28.86.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→RSVP→Interfaces→Interface.

2

Right-click on the Interface icon and choose Properties. The RSVP Interface (Edit) form opens.

3

Configure the required parameters.

4

Click on the Protocol Properties tab and configure the required parameters.

The Class Type BW Percent parameters can only be configured if the Inherit NFM-P Class Type BW parameter is disabled. If you enter values at the RSVP interface level, they override the global values configured in 28.85 “To configure RSVP on a routing instance” (p. 985) for the RSVP routing instance on the NE.

5

Click on the Authentication tab and configure the Key parameter.

6

Select an RSVP Keychain.

7

Click on the TE Thresholds tab and configure the required parameters.

You can only configure the Up Threshold (%) and Down Threshold (%) parameters if the Inherit TE Up Thresholds and Inherit TE Down Thresholds parameters are disabled, respectively.

8

Save your changes and close the form.

END OF STEPS

L2TP configuration workflow and procedures

28.87 L2TP configuration

28.87.1 Overview

The NFM-P supports the configuration and management of L2TP sites, tunnel group profiles, and tunnel profiles. L2TPv2 and L2TPv3 tunnels are available. Both can be configured at the same time.

A typical L2TPv2 configuration is applied to two NEs; one NE performs the LAC role and the other NE performs the LNS role. An NE can perform both LAC and LNS roles. At least one ISA-LNS group must be configured for the LNS NE. See 13.1.5 “Working with ISA-LNS groups” (p. 407) for information about ISA-LNS groups. See 13.12 “To configure an ISA-LNS group” (p. 426) for information about creating and configuring an ISA-LNS group member.

L2TPv3 tunnels pass Ethernet traffic. The tunnels are symmetric and do not have LAC and LNS ends.

Note: An NE must be in chassis mode B or higher to support L2TP.
You can also enable L2TP on a VPRN site. See Chapter 79, “VPRN service management” .

The following procedures describe how to configure L2TP.

28.88 Workflow to configure L2TP

28.88.1 Stages

1

Configure L2TP on a routing instance; see 28.89 “To configure L2TP on a routing instance” (p. 989) .

2

Update tunnel instance endpoints on an L2TP site as required; see 28.90 “To update tunnel instance endpoints on an L2TP site” (p. 992) .

3

Create and configure an LNS group and group member for an LNS site. If the site is an LAC, this configuration is not required; see 13.12 “To configure an ISA-LNS group” (p. 426) .

4

Configure the ESM profiles for LNS L2TP termination of the following as required:

a. VPRN; see 79.37 “To configure a group interface on a VPRN” (p. 2586)

b. IES; see 78.19 “To configure a group interface on an IES” (p. 2449)

5

Configure L2TP on each local user database PPPoE host to be forwarded to the LNS (applicable for the LAC and the LNS, where RADIUS is not used) as required; see 74.9 “To configure a local user database for subscriber host authentication” (p. 2025) .

6

View the following L2TP related objects, as required:

- L2TP tunnels and tunnel endpoints; see 28.91 “To view L2TP tunnels and tunnel endpoints” (p. 993) .
- L2TP tunnel instance endpoints associated with a subscriber host; see 28.92 “To view L2TP tunnel instance endpoints on a subscriber instance” (p. 994) .
- L2TP sessions; see 28.93 “To view L2TP sessions” (p. 994) .
- PPP sessions on an LAC or LNS; see 28.94 “To view PPP sessions” (p. 995) .

28.89 To configure L2TP on a routing instance

Note: For most parameters in this procedure, you can specify that the parameter value is inherited from the parent L2TP configuration using the Inherit Value parameter.

If you disable value inheritance for a parameter, the available options are restricted, based on the parent parameter value and the protocol functionality. For example, if a parameter in the global-level L2TP configuration is set to True, the only available option for the same parameter in the tunnel group profile-level L2TP configuration is False, unless a value of False violates a protocol rule, in which case the only available option is True.

The parameters that you configure for a L2TP tunnel group profile take precedence over the parameters that are configured for global-level L2TP.

28.89.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→L2TP.

2

Right-click on the L2TP icon and choose Properties. The L2TP Site (Edit) form opens.

3

Configure the required parameters.

4

Select an L2TP Accounting Policy next to the RADIUS Accounting Policy field. See 57.5 “To configure an L2TP RADIUS accounting policy” (p. 1796) for information on creating an L2TP RADIUS Accounting Policy.

5

On the Failover Information panel, configure the failover recovery parameters, as required.

All tunnel group profiles and tunnel profiles configured on the L2TP site can be configured to inherit the failover recovery information specified here, or they can be configured individually with local failover recovery information.

6

Click on the Tunnel Selection Blacklist tab and configure the required parameters.

7

Click on the L2TPv3 tab to configure L2TPv3 support on the L2TP site.

Configure the required parameters.

8

To create a tunnel group profile:

If you are modifying an existing tunnel group profile, you can manually update the Tunnel Instance Endpoints list for the L2TP site after saving your changes. See Step 11 .

1. Click on the Tunnel Group Profiles tab and click Create. The L2TP Tunnel Group Profile (Create) form opens.
2. Select the L2TP Protocol Version in the Identification panel. Different tabs and parameters are available based on the L2TP protocol version.
3. Configure the required L2TPv2 parameters.
   a. To generate operational L2TPv2 tunnels, a start operation must be performed on the L2TP tunnel on the LAC NE. You can also perform a start operation from a tunnel profile on the LNS.

      Failover recovery information can be configured locally or inherited from the parent L2TP site.
   b. Select an LNS group in the LNS Group ID panel.

      You must choose an LNS group for the tunnel group profile or tunnel profile of the site that acts as the LNS.
   c. Select a RADIUS accounting policy in the L2TP Radius Accounting Policy panel.
4. Configure the required parameters in the L2TPv3 tab if L2TPv3 is chosen.

   When a tunnel group is established, the L2TPv3 Ethernet Sessions tab will list a single item.

9

To configure PPP on an L2TPv2 tunnel group profile:

PPP is required only for LNS L2TP tunnel group profiles and L2TP tunnel profiles.

1. Click on the PPP tab and configure the required parameters.
2. Configure the CHAP challenge length parameters, as required.
3. Select an authentication policy in the Authentication Policy panel.

4. If you did not choose an authentication policy in 3 , select a user database in the User Database panel.
5. Select a default service ID in the Default Service ID panel.
6. Select a default group interface in the Default Group Interface panel.

10

To configure a tunnel profile on the L2TP tunnel group profile:

If you are modifying an existing tunnel profile, you can manually update the Tunnel Instance Endpoints list for the L2TP site after saving your changes. See Step 11 .

1. Click on the Tunnel Profiles tab and click Create. The L2TP Tunnel Profile (Create) form opens.

   Failover recovery information can be configured locally or inherited from the parent tunnel group profile.
2. Configure the v2 parameters as required:
   a. Select an LNS group in the LNS Group ID panel.

      You must choose an LNS group for the tunnel group profile or tunnel profile of the site that acts as the LNS.
   b. Click Select in the L2TP Radius Accounting Policy panel and choose a RADIUS accounting policy from the Select L2TP RADIUS Accounting Policy form.
   c. Click on the PPP tab and repeat Step 9 to configure PPP on the L2TP tunnel profile, as required.
   d. You only configure PPP for LNS L2TP tunnel group profiles and L2TP tunnel profiles.

      The values for the MTU, Keep-Alive Interval, Keep-Alive Multiplier, and CHAP Challenge Length Min and Max parameters can be inherited from the parent tunnel group profile.
   e. Click on the MLPPP tab and configure the required parameters.

      When Inherit from Tunnel Group Profile beside a parameter is enabled, the parameter value from the tunnel group profile is used. You must disable Inherit from Tunnel Group Profile beside a parameter before you can configure the parameter.
3. Save your changes and close the forms.

11

If the tunnel group profiles or tunnel profiles for the L2TP site have been modified, you can update the Tunnel Instance Endpoints list for the L2TP site.

Perform the following steps:

1. Click on the Tunnel Instance Endpoints tab.
2. Click Resync Tunnel Instance Endpoints. A warning message window opens.
3. Click Yes.

12

To configure an L2TP failover SRRP peer site:

Perform the following steps:

1. Click on the Failover SRRP Peers tab.
2. Click Create or select an existing L2TP peer and click Properties. The L2TP Failover SRRP Peer (Create|Edit) form opens.
3. Select a peer address.
4. Select a Track SRRP Instance.
5. Configure the Sync Tag parameter.
6. Save your changes and close the form.

13

Save your changes and close the form.

END OF STEPS

28.90 To update tunnel instance endpoints on an L2TP site

28.90.1 Purpose

If a tunnel group profile or tunnel profile is modified on an L2TP site, the list of tunnel instance endpoints for the site is not updated automatically. Perform this procedure to manually update the content of the Tunnel Instance Endpoints tab for the L2TP site.

28.90.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→L2TP.

2

Right-click on the L2TP icon and choose Properties. The L2TP Site (Edit) form opens.

3

Click on the Tunnel Instance Endpoints tab and click Resync Tunnel Instances Endpoints.

4

Save your changes and close the form.

END OF STEPS

28.91 To view L2TP tunnels and tunnel endpoints

28.91.1 Steps

1

Choose Manage→ISA Functions→ISA-L2TP from the NFM-P main menu. The Manage ISA-L2TP form opens.

2

Choose L2TP tunnel (L2TP).

3

Click Search. A list of L2TP tunnels is displayed.

4

Choose an entry in the list and click Properties. The L2TP Tunnel - Endpoint A - Endpoint B form opens.

5

View the information for Tunnel Endpoint A and Tunnel Endpoint B.

6

Click on the L2TP Tunnel Endpoints tab. The two L2TP tunnel endpoints are displayed.

7

Choose an entry in the list and click Properties. The L2TP Tunnel Endpoint (View) form opens.

8

View the information. You can also click Properties to view additional information for the following:

- Site ID
- Tunnel Instance Endpoint
- Tunnel Profile
- Tunnel Group Profile
- Peer

9

Close the forms.

END OF STEPS

28.92 To view L2TP tunnel instance endpoints on a subscriber instance

28.92.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) from the Select Object Type drop-down menu.

3

Choose a subscriber instance and click Properties. The Residential Subscriber Instance (Edit) form opens.

4

Click on the L2TP Tunnel Instance Endpoints tab.

5

Choose an L2TP tunnel instance endpoint and click Properties. The L2TP Tunnel Instance Endpoint (View) form opens.

6

View the information, as required. You can also click on the Sessions tab to view session information.

7

Close the forms.

END OF STEPS

28.93 To view L2TP sessions

28.93.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose L2TP Session (L2TP) from the Select Object Type drop-down menu.

3

Select a site ID and a routing instance site.

4

Configure the Connection ID parameter.

5

Choose a session and click Properties. The L2TP Session (View) form opens.

6

View the information in the following tabs, as available:

Note: Different tabs and panels will be displayed based on whether the session is using L2TPv2 or L2TPv3.

- General tab:
  - Session
  - Tunnel
  - Operational Information
  - LNS PPP

    Note:

    The LNS PPP panel is applicable only for an LNS.
  - L2TPv3 Information
- L2TPv3 Ethernet Tunnel Stats tab

7

Close the forms.

END OF STEPS

28.94 To view PPP sessions

28.94.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose PPP Session (PPP Policy and Session) from the Select Object Type drop-down menu.

3

Select a site ID and a service site.

4

Choose a session in the list and click Properties. The PPP Session form opens.

5

View the information in the following panels:

- PPP General
- Residential Subscriber
- Business Subscriber
- DHCP

6

Click on the following tabs to view the information:

- Authentication Protocols
- DNS and NBNS IP Addresses
- QoS Overrides
- Deployment

7

Close the forms.

END OF STEPS

PIM configuration workflows and procedures

28.95 Workflow to configure PIM

28.95.1 Stages

1

Enable PIM on a routing instance; see 28.97 “To enable PIM on a routing instance” (p. 998) .

2

Configure PIM on a routing instance or a VPRN routing instance; see 28.98 “To configure PIM on a routing instance” (p. 998) and 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003) .

3

Configure Anycast RP for PIM on a router as required; see 28.100 “To configure Anycast PIM on a router” (p. 1011) .

4

Configure a PIM interface; see 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013) .

28.96 Workflow to configure VRRP-aware PIM

28.96.1 Stages

1

Create an operational group. Do not enable BDF Interface monitoring. See 12.8 “To create an operational group” (p. 345).

2

Configure a VRRP instance on either a base router, IES interface, or VPRN interface. In the Operational Group panel, add the operational group created in Stage 1.

See one of the following:

- 78.54 “To create a VRRP instance on an IES L3 access interface for a virtual router” (p. 2500)
- 79.110 “To create a VRRP instance on a VPRN L3 access interface for a virtual router” (p. 2685)
- 37.4 “To create and configure a VRRP instance” (p. 1283)

3

Configure a PIM interface on the instance. Configure the Designated Router parameters in the Behavior tab to ensure that the Operational DR Priority is the highest priority on the interface if the monitored group is up. For example, if Delta Priority Action is set to Add, Operational DR Priority is equal to DR Priority plus Operational Group Delta Priority.

See one of the following:

- 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013)
- 78.9 “To add a PIM interface to an IES” (p. 2435)
- 79.45 “To add a PIM interface to a VPRN” (p. 2606)

28.97 To enable PIM on a routing instance

28.97.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance and choose Properties. The Routing Instance (Edit) form opens.

3

Enable PIM:

Perform the following steps:

1. Click on the Multicast tab.
2. Select the PIM Enabled parameter.

4

Save your changes and close the form. The PIM icon appears in the navigation tree below the routing instance.

END OF STEPS

28.98 To configure PIM on a routing instance

28.98.1 Purpose

Note: If you are creating a PIM site on a VPRN routing instance, perform 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003) . See 79.22 “To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance” (p. 2555) for more information about configuring PIM on a VPRN routing instance.

The tabs and parameters that are configurable vary, depending on the NE on which the PIM site resides.

28.98.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→PIM.

2

Right-click on the PIM icon and choose Properties. The PIM Site (Edit) form opens.

3

Configure the Administrative State parameter in the States panel.

4

Configure the required parameters in the PIM General Configurations and ECMP Balancing panels.

5

In the Lag Usage panel, configure the Lag Usage Optimization parameter.

6

Click on the RP Behavior tab. The IPv4 tab is displayed.

7

If you are configuring for IPv4, click on the IPv4 tab and go to Step 8 . Otherwise click on the IPv6 tab and go to Step 12 .

8

Configure the IPv4 Administrative State parameter in the States panel.

9

Configure the required parameters in the PIM General Configurations panel.

10

Configure the required parameters in the Candidate Bootstrap Router and Candidate Rendezvous Point panels.

11

Go to Step 15 .

12

Configure the IPv6 Administrative State parameter in the States panel.

13

Configure the IPv6 RPF Lookup Sequence parameter in the PIM General Configurations panel.

14

Configure the required parameters in the Candidate Bootstrap Router, Candidate Rendezvous Point, and Embedded-RP panels.

15

Add one or more C-RP groups.

Perform the following steps:

1. Click on the Candidate-RP Groups tab.
2. Click Create to add a new entry. The C-RP Group Prefix (Create) form opens.
3. Configure the required parameters.
4. Save your changes.

16

Add one or more embedded RP groups.

The Embedded-RP Groups tab is only selectable if you enabled the Enable Embedded-RP parameter.

Perform the following steps:

1. Click on the Embedded-RP Groups tab.
2. Click Create to add a new entry. The E-RP Group Range (Create) form opens.
3. Configure the required parameters.
4. Save your changes.

17

Click on the Group To RP tab. The Static RP tab is displayed.

18

Click Create. The Static RP (Create) form opens.

19

Configure the required parameters.

20

Click on the Static Group-To-RP tab.

21

Click Create. The Static Group To RP (Create) form opens.

22

Configure the required parameters.

23

Click on the Import Policies tab, then on the BootStrap Import Policies tab.

24

Configure the required parameters to filter bootstrap messages and to control the flow of bootstrap messages to the routing instance.

25

Click on the Join/Prune Import Policies tab.

26

Configure the required parameters to filter join/prune messages.

27

Click on the Register Import Policies tab.

28

Configure the required parameters to filter PIM register messages.

29

Click on the Export Policies tab, then on the BootStrap Export Policies tab.

30

Configure the required parameters to filter PIM-related export messages.

31

Depending on the device release, the Anycast RP tab is configurable. Click on the Anycast RP tab to configure Anycast RP.

Note: You can use the PIM configuration form to add peers to an anycast RP. The complete configuration of anycast RP for PIM requires additional supporting components, such as the configuration of loopback interfaces and static RPs. Due to the complex configuration dependencies, Nokia recommends that you use the Virtual Anycast manager. The Virtual Anycast manager automatically configures many of the supporting requirements for the protocol, which reduces operator configuration errors and assists in troubleshooting activities. See 28.100 “To configure Anycast PIM on a router” (p. 1011) for more information about using the Virtual Anycast manager.

32

Click on the General tab.

33

Configure the RP IP Address parameter.

34

Click on the Anycast Peer tab.

35

Click Create. The Anycast Peer (Create) form opens.

36

Configure the Peer IP Address parameter.

37

Save your changes.

38

Click on the SSM Groups tab.

39

Click Create.

40

Configure the required parameters.

41

Save your changes.

42

Click on the Interfaces tab to create an interface.

43

Click Create; see 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013)

44

Click on the SPT Switch Threshold tab to configure SPT switchover, if required.

45

Click Create. The Spt Switch Over Threshold (Create) form opens.

46

Configure the required parameters.

47

Save your changes.

48

Click on the following tabs to view and edit information.

- Groups
- Neighbor
- Statistics
- Faults

The Statistics and Faults tabs are unavailable for PIM in VPRN.

49

Save your changes and close the form.

END OF STEPS

28.99 To create a PIM site on a VPRN routing instance

28.99.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search, choose a VPRN service, and click Properties. The VPRN Service (Edit) form opens.

3

On the navigation tree, expand Sites →Routing Instance - NE System ID→Routing Instance.

4

Right-click on Protocols and choose Create PIM Site. The PIM Site (Create) form opens.

5

Configure the required parameters.

6

Click on the RP Behavior tab, then on the IPv4 and IPv6 sub-tabs and configure the required parameters.

7

Click on the MVPN tab.

8

Configure the required parameters.

Note: You must configure the Inclusive Tunnel Type parameter if you are creating an inclusive tunnel.

9

Configure redundant MVPN sources.

Perform the following steps:

1. Click on the Redundant Sources sub-tab and click Create. The Redundant Source (Create) form opens.
2. Configure the required parameters and click OK.

10

Configure primary and standby UMH addresses.

Perform the following steps:

1. Click on the UMH PE Backup sub-tab and click Create. The MVPN UMH PE Backup (Create) form opens.
2. Configure the required parameters and click OK.

11

Configure RPF select groups.

Perform the following steps:

1. Click on the RPF Select sub-tab and click Create. The RPF Select (Create) form opens.
2. Configure the required parameters and click OK.

12

Click on the MVPN Inclusive Tunnel tab. Configure UMH rate monitoring triggers for PMSI redundancy, if required.

Note: The MVPN Inclusive Tunnel tab appears only if you set the Inclusive Tunnel Type parameter to an option other than None in Step 8 .
The tab that appears on the MVPN Inclusive Tunnel tab depends on the Inclusive Tunnel Type parameter option that you configured in Step 8 .

13

Perform one of the following:

a. Create a PIM inclusive tunnel.

   Perform the following steps:

   1. Click on the PIM sub-tab.
   2. Configure the required parameters.
   3. Go to Step 14 .

b. Create an RSVP inclusive tunnel.

   Perform the following steps:

   1. Click on the RSVP tab.
   2. Configure the P2MP Administrative State parameter. You cannot set the P2MP Administrative State parameter to Up until you choose an MVPN LSP template. To create an MVPN LSP template, see 31.29 “To create an LSP template MVPN policy” (p. 1156) .
   3. Click Select. The Select MVPN Lsp Template - RSVP MVPN Provider Tunnel form opens.
   4. Choose an MVPN LSP template and click OK. The MVPN LSP template is displayed.
   5. Configure the required parameters.

c. Create an MLDP inclusive tunnel.

   Perform the following steps:

   1. Click on the MLDP tab.
   2. Configure the P2MP Administrative State parameter.
   3. Go to Step 14 .

d. Create a BIER inclusive tunnel.

   Perform the following steps:

   1. Click on the BIER tab.
   2. Configure the P2MP Administrative State parameter.
   3. Enter the BIER Subdomain number.
   4. Go to Step 14 .

e. Create a Segment Routing inclusive tunnel.

   Perform the following steps:

   1. Click on the Segment Routing tab.
   2. Configure the P2MP Administrative State parameter.
   3. Click Select and choose an SR-Tree policy.
   4. Go to Step 14 .

14

Click on the MVPN Selective Tunnel tab.

15

Configure the Delay Interval (seconds) parameter in the Data MDT panel.

16

Create UMH rate monitoring groups as required, by performing the following:

1. In the UMH Rate Monitoring panel, click Create. The UMH Rate Monitoring Group (Create) form is displayed.
2. Configure the required parameters, and click OK. The new rate monitoring group is added to the list.

Note: You can create multiple entries with the same group address and different source addresses.

17

Configure a data MVPN threshold entry by performing one of the following:

a. Choose an existing data MVPN threshold entry.

Perform the following steps:

1. Click Search in the Data MVPN Threshold tab. The Data MVPN Threshold (Create) form is displayed.
2. Choose a Data MVPN Threshold entry, and click Apply.
3. Go to Step 18 .

b. Create and choose a data MVPN threshold entry.

Perform the following steps:

1. Click Create in the Data MVPN Threshold tab. The Data MVPN Threshold (Create) form is displayed.
2. Configure the required parameters.
3. Click OK, choose the Data MVPN Threshold entry that you created, and click Apply.

18

Configure a data MT interface entry by performing one of the following:

Note: A data Multicast Tunnel (MT) is a connector between a set of PE routers forming a multicast domain. In the context of a VPN-specific PIM instance, a multicast tunnel is a single multi-access interface. The data MT Interface can be created or selected when the PIM address is configured with a value other than null (0.0.0.0).

a. Choose an existing data MT interface entry.

Perform the following steps:

1. Click Search in the Data MT Interface tab.
2. Choose a Data MT Interface entry.
3. Save your changes.
4. Go to Step 19 .

b. Create and choose a data MT interface.

Perform the following steps:

1. Click Add in the Data MT Interface tab. The Data MT Interface (Create) form is displayed.
2. Configure the required parameters.
3. Click OK, choose the data MT interface that you created, and click Apply.

19

Configure the required parameters on the PIM sub-tab.

20

Configure the required parameters on the RSVP sub-tab.

If the RSVP sub-tab is not available, go to Step 23.

21

Click Select in the RSVP panel to choose an MVPN Lsp Template. The Select MVPN Lsp Template form is displayed.

22

Perform one of the following:

a. Choose an existing LSP template for MVPN.

Perform the following steps:

1. Click Search. All of the available MVPN p2mp LSP templates for the local NE (the NE contains this VPRN/PIM) are displayed.
2. Choose the required template and click OK.
3. Go to Step 23 .

b. Create an LSP template for MVPN.

Perform the following steps:

1. Click Create and perform 31.29 “To create an LSP template MVPN policy” (p. 1156) .
2. Choose the newly-created template in the Select MVPN Lsp Template form and click OK.

23

Configure the required parameters on the MDLP sub-tab.

24

Configure the required parameters on the BIER sub-tab.

25

Configure the required parameters on the Segment Routing sub-tab.

26

Perform the following steps as needed to create a Selective PMSI:

1. Choose the MultiStream S-PMSI sub-tab.
2. Click Create. The MultiStream S-PMSI, Routing Instance (Create) form opens.
3. Click Select in the Routing Instance panel to choose an MVPN Lsp Template. The Select MVPN Lsp Template form is displayed.

27

Perform one of the following. Available options depend on the Selective Tunnel Type parameter.

a. Choose an existing LSP template for MVPN.

   Perform the following steps:

   1. Click Search. All of the available MVPN p2mp LSP templates for the local NE (the NE that contains this VPRN/PIM) are displayed.
   2. Choose the required template and click OK.
   3. Complete the configuration and close the forms.

b. Create an LSP template for MVPN.

   Perform the following steps:

   1. Click Create and perform 31.29 “To create an LSP template MVPN policy” (p. 1156) .
   2. Choose the newly-created template in the Select MVPN Lsp Template form and click OK.
   3. Complete the configuration and close the forms.

c. Select an SR-Tree policy.

   Perform the following steps:

   1. Click Search. All of the available SR-Tree policies for the local NE (the NE that contains this VPRN/PIM) are displayed.
   2. Choose the required policy and click OK.
   3. Complete the configuration and close the forms.

Go to Step 28.

28

Configure import policies.

Perform the following steps:

1. Click on the Import Policies tab.
2. Click on the BootStrap Import Policies, Join/Prune Import Policies, and Register Import Policies sub-tabs and configure the message filter parameters.

29

Configure export policies.

Perform the following steps:

1. Click on the Export Policies tab.
2. Click on the BootStrap Export Policies sub-tab and configure the required parameters to filter PIM-related export messages.

30

Create an interface. Click on the Interfaces tab, click Create, and perform 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013) .

31

Add one or more C-RP groups.

Perform the following steps:

1. Click on the Candidate-RP Groups tab.
2. Click Create to add a new entry. The C-RP Group Prefix (Create) form opens.
3. Configure the required parameters and click OK.

32

Add one or more static RPs.

Perform the following steps:

1. Click on the Group To RP tab.
2. Click Create. The Static RP (Create) form opens.
3. Configure the required parameters.
4. Click on the Static Group-To-RP tab.
5. Click Create. The Static Group To RP (Create) form opens.
6. Configure the required parameters.
7. Save your changes and close the form.

33

Create one or more SSM groups.

Perform the following steps:

1. Click on the SSM Groups tab and click Create.
2. Configure the required parameters and click OK.

34

Create one or more SPT switch thresholds.

Perform the following steps:

1. Click on the SPT Switch Threshold tab and click Create.
2. Configure the required parameters and click OK.

35

Configure GRT extranet group prefixes.

Perform the following steps:

1. Click on the GRT Extranet tab and configure the required parameters.
2. Click on the Group Prefix sub-tab and click Create.
3. Configure the required parameters and click OK.

36

Configure Anycast RP, as required.

Note: You can use the PIM configuration form to add peers to an anycast RP. The complete configuration of anycast RP for PIM requires additional supporting components, such as the configuration of loopback interfaces and static RPs. Due to the complex configuration dependencies, Nokia recommends that you use the Virtual Anycast manager.

The Virtual Anycast manager automatically configures many of the supporting requirements for the protocol, which reduces operator configuration errors and assists in troubleshooting activities. See 28.100 “To configure Anycast PIM on a router” (p. 1011) for more information about using the Virtual Anycast manager.

1. Click on the Anycast RP tab and click Create.
2. Configure the required parameters.
3. Click on the Anycast Peer tab and click Create.
4. Configure the required parameters and click OK.
5. Save your changes and close the form.

37

Click on the following tabs to view and edit information.

- Groups
- Neighbor
- Statistics
- Faults

The Statistics and Faults tabs are unavailable for PIM in VPRN.

38

Save your changes and close the form.

END OF STEPS

28.100 To configure Anycast PIM on a router

28.100.1 Purpose

This procedure describes how to use the Virtual Anycast Manager to configure anycast PIM on a router. The Virtual Anycast Manager automatically configures many of the supporting requirements for the protocol, which reduces operator configuration errors and assists in troubleshooting activities.

Note: The NFM-P does not automatically delete the supporting requirements for the protocol when you delete an anycast RP member.

You can also manually add peers to an anycast RP by navigating to the Anycast RP tab of the PIM configuration form. Ensure that PIM is enabled on the router. See 28.98 “To configure PIM on a routing instance” (p. 998) .

The NFM-P raises an alarm when VRFs from different VPRNs are added to a single anycast RP set. You can only mismatch VRFs using methods that are not associated with the Virtual Anycast Manager; for example, by using the CLI or the PIM configuration form.

28.100.2 Steps

1

Choose Manage→Networking→Virtual Anycast RP from the NFM-P main menu. The Manage Virtual Anycast RP form opens.

2

Click Create. The Virtual Anycast RP form opens.

3

Configure the required parameters.

4

If the Anycast RP Type parameter value is VPRN, choose a VPRN service for the anycast RP.

5

Click Apply. The Status panel appears. The Status panel is used in Step 9 of this procedure to identify configuration problems.

6

Click on the Components tab.

7

Add virtual Anycast RP members to the PIM configuration.

Perform the following steps:

1. Right-click on the Virtual Anycast RP Members icon and choose Add Anycast RP Member. The Enter Interface Name form opens.
2. Enter the interface name for the loopback interface. The Auto Created Interface Name parameter value appears, if configured.
3. Click OK. The Enter Interface Name form closes, and the Select Local Address form opens.
4. Configure the filter criteria.
5. Choose the local IP address that is used to communicate with the other RP members in the virtual anycast set and click OK. The new local address entry appears on the Components tab under the Virtual Anycast RP Members icon. The new local address entry also appears under the Static RP icon.

   Note:

   The local address is typically the system address.

   The NFM-P automatically performs the following virtual anycast RP configuration tasks.
   - creates a PIM-enabled loopback interface, if not present, with the RP address
   - adds a local address to the anycast RP peer set
   - updates the peer sets that participate in the anycast RP so that all peer sets contain the same members
   - enables PIM on the interface that contains the local IP address
   - creates a static RP that uses the anycast RP address. Groups must be manually created using the Components tab
6. Repeat Step 7 1 to 5 for each member in the virtual anycast RP.

8

Add a static group-to-RP mapping for the anycast configuration.

Perform the following steps:

1. Right-click on a member in the Static RP list and choose Create Static RP. The Static RP (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.
4. Repeat Step 8 1 to 3 for each static group-to-RP mapping that you want to add to the virtual anycast RP configuration. Each member needs a static RP and at least one static group-to-RP mapping. The mappings typically use the same configuration values for each member in the group.

9

Verify the status of the anycast RP for PIM configuration.

Perform the following steps:

1. Click on the General tab. The Status panel displays the status of the anycast RP configuration. The NFM-P updates the status when you add a member to or delete a member from the virtual anycast RP. You can also use the Update Status button to manually update the Status panel.

   The following table lists the status check boxes that identify potential configuration problems.

Check box	Description
Missing Group Range(s) Configurations	At least one anycast RP member does not have a group range for the anycast RP address.
Inconsistent Peer Sets	Peer sets of all the anycast RP members are not equal.
PIM not enabled on loopback interface(s)	PIM is not enabled on all loopback interfaces that are configured in the virtual anycast RP.
Only one peer configured	Only one peer is configured for anycast RP. The configuration requires two or more peers for correct functionality.
Loopback interface(s) not properly configured	At least one anycast RP member does not have a loopback interface configured with the anycast RP address.
Missing Static RP configuration(s)	A member of the virtual anycast RP does not have a static RP. Each member requires a static RP.

2. Save your changes and close the form.

END OF STEPS

28.101 To create a PIM interface on a base routing instance or VPRN routing instance

28.101.1 Steps

1

Perform one of the following:

a. Create a PIM interface on a base routing instance.

   Perform the following steps:

   1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→PIM.
   2. Right-click on the PIM icon and choose Create Interface. The PIM Interface (Create) form opens.

b. Create a PIM interface on a VPRN routing instance.

Perform the following steps:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
3. Expand VPRN Service→Sites→Routing Instance - NE System ID→Routing Instance→Routing Instance: *VPRNinstance*→Protocols→PIM.
4. Right-click on the PIM instance and choose Create PIM Interface. The PIM Interface, Routing Instance (Create) form opens.

2

Click Select, choose a Layer 3 interface, and click OK. The Select Interface - PIM Interface form closes, and the interface name appears in the Interface panel.

3

Configure the required parameters.

Note: If you are configuring IPv6 BFD on a VPRN routing instance, you need to enable the IPv6 Allowed parameter on the associated L3 access interface and set the BFD IPv6 Admin State parameter to Up on the BFD tab. See 79.106 “To configure BFD for a VPRN L3 access interface” (p. 2680) .

4

Click on the Behavior tab, then on the General sub-tab.

5

Configure the required parameters.

Note: The Sticky DR Priority parameter is configurable when the Sticky DR parameter is enabled.

6

Click on the Neighbor tab to view and edit information. The Neighbor tab is configurable only when a neighbor PIM interface exists.

7

Click on the IPv4 or IPv6 sub-tab.

8

Configure the IPv4 Administrative State or IPv6 Administrative State parameter in the States panel.

9

Click on the Multicast CAC tab to add a multicast CAC policy.

Perform the following steps:

1. Select a multicast CAC policy.
2. Configure the required parameters.
3. Click on the Levels tab.
4. Click Create. The PIM Interface Multicast CAC level form opens.
5. Configure the required parameters.
6. Click on the LAG Port Down tab.
7. Click Create. The PIM Interface Multicast CAC LAG Port Down form opens.
8. Configure the required parameters.

10

Save your changes and close the form.

END OF STEPS

IGMP configuration workflow and procedures

28.102 Workflow to configure IGMP

28.102.1 Stages

1

Enable IGMP on a routing instance; see 28.103 “To enable IGMP on a routing instance” (p. 1016) .

2

Configure IGMP on a routing instance; see 28.104 “To configure an IGMP site on a router” (p. 1017) .

3

Configure IGMP on an OmniSwitch; see 28.105 “To configure IGMP on an OmniSwitch” (p. 1019) .

4

Configure an IGMP interface; see 28.106 “To configure an IGMP interface” (p. 1019) .

5

Turn up or shut down an IGMP interface as required; see 28.107 “To turn up or shut down an IGMP interface” (p. 1021) .

6

View IGMP multicast reporting statistics or IGMP source statistics as required; see 28.108 “To view IGMP multicast reporting statistics for an IGMP site” (p. 1022) or 28.109 “To view IGMP source statistics” (p. 1022) .

28.103 To enable IGMP on a routing instance

28.103.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) configuration form opens.

3

Click on the Multicast tab.

4

Select the IGMP Enabled check box.

5

Save your changes and close the form.

END OF STEPS

28.104 To configure an IGMP site on a router

28.104.1 Purpose

Before you can configure an IGMP site on a router, you must enable IGMP on a routing instance. See 28.103 “To enable IGMP on a routing instance” (p. 1016) .

28.104.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→IGMP.

2

Right-click on the IGMP icon and choose Properties. The IGMP Site (Edit) form opens.

3

Configure the required parameters.

4

Click on the SSM Translation tab to configure SSM Translation.

Perform the following steps:

1. Click Create. The SSM Translation (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form. The new SSM Translation entry appears on the form.

5

Click on the Interfaces tab to add an IGMP interface; see 28.106 “To configure an IGMP interface” (p. 1019) for information about configuring an IGMP interface.

6

Click on the Group Interfaces tab to view IES group interfaces added to IGMP.

7

Click on the LDP Tunnel Interfaces tab to configure an LDP tunnel interface entry.

Perform the following steps:

1. Click Create. The LDP IGMP Tunnel Interface (Create) form opens.
2. Select an LDP tunnel interface.
3. Click on the Static Group/Source tab to configure a static multicast entry.
4. Click Create. The StaticGrpSrc (Create) form opens.
5. Configure the required parameters.
6. Save your changes and close the forms.

8

Click on the RSVP Tunnel Interfaces tab to configure an RSVP tunnel interface.

Perform the following steps:

1. Click Create. The RSVP IGMP Tunnel Interface (Create) form opens.
2. Select an RSVP IGMP tunnel interface.
3. Click on the Static Group/Source tab to configure a static multicast entry.
4. Click Create. The StaticGrpSrc (Create) form opens.
5. Configure the required parameters.
6. Save your changes and close the forms.

9

Click on the Static Range Group tab to configure static ranges. Perform the following:

1. Click Create to add a static range group. The Static Grp Range Src form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

10

Click on the following tabs to view and edit information.

- Multicast Group/Source
- Statistics
- Faults

11

Save your changes and close the forms.

END OF STEPS

28.105 To configure IGMP on an OmniSwitch

Note: IGMP parameters can be configured globally or for each VLAN site. The VLAN site IGMP configuration overrides the routing instance IGMP settings. You can only configure VLAN site IGMP parameters after you create the VLAN. See 75.16 “To configure IGMP on an OmniSwitch VLAN site” (p. 2086) for information about configuring IGMP on an OmniSwitch VLAN site.

28.105.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→IGMP.

2

Right-click on the IGMP icon and choose Properties. The IGMP (Edit) form opens.

3

Configure the required parameters.

Note: When the Administrative State parameter is set to Up, the IP multicast switching is enabled on the OmniSwitch.

4

Save your changes and close the form.

END OF STEPS

28.106 To configure an IGMP interface

28.106.1 Purpose

Before an IGMP interface can be configured, the following prerequisites must be met:

- The host routing instance must be configured for IGMP; see 28.103 “To enable IGMP on a routing instance” (p. 1016) .
- The host routing instance must be configured with an L3 interface; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) .

Note: IGMP IPv6 interfaces are supported on the 7950 XRS, on the 7750 SR in chassis mode A or B with mixed mode enabled, and on the 7750 SR in chassis mode C or D.

The tabs and parameters that are configurable vary, depending on the NE on which the IGMP site resides.

28.106.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→IGMP.

2

Right-click on the IGMP icon and choose Create Interface, or right-click on an existing IGMP interface and choose Properties. The IGMP Interface (Create|Edit) form opens.

3

Select a Layer 3 interface in the Interface panel.

4

Configure the required parameters.

5

Click on the Behavior tab.

6

Select an import policy, and configure the required parameters.

7

To configure multicast CAC:

Perform the following steps:

1. Click on the Multicast CAC tab
2. Select a multicast CAC policy in the Multicast CAC Policy panel.
3. Configure the required parameters.
4. Click on the Levels tab.
5. Click Create or select an existing multicast CAC level entry and click Properties. The Multicast CAC Level (Create|Edit) form opens.
6. Configure the required parameters.
7. Save your changes and close the form.
8. Click on the LAG Port Down tab.
9. Click Create or select an existing LAG port down entry and click Properties. The LAG Port Down (Create|Edit) form opens.
10. Configure the required parameters.
11. Save your changes and close the form.

8

Click on the SSM Translation tab to configure SSM translation entries.

Perform the following steps:

1. Click Create or select an existing SSM translation entry and click Properties. The SSM Translation (Create|Edit) form appears.
2. Configure the required parameters.
3. Save your changes and close the form.

9

Click on the Static Group/Source tab to create a static multicast entry.

Perform the following steps:

1. Click Create or select an existing static multicast entry and click Properties. The Static Grp Src (Create|Edit) form appears.
2. Configure the required parameters.
3. Save your changes and close the form.

10

Save your changes and close the form.

END OF STEPS

28.107 To turn up or shut down an IGMP interface

28.107.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→IGMP.

2

Right-click on the IGMP icon and choose Properties. The IGMP Site (Edit) form opens.

3

Click on the Interfaces tab.

4

Turn up or shut down the IGMP interface:

a. Click Turn Up to activate the interface.

b. Click Shut Down to deactivate the interface.

5

Save your changes and close the forms.

END OF STEPS

28.108 To view IGMP multicast reporting statistics for an IGMP site

28.108.1 Purpose

Use the following procedure to view IGMP multicast reporting statistics for an IGMP site. See the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting statistics.

28.108.2 Steps

1

Choose Manage→Networking→Routing Instances from the NFM-P main menu. The Manage Routing Instances form opens.

2

Choose an IGMP Site (IGMP) from the Select Object Type menu.

3

Click Search, choose an IGMP site and click Properties. The IGMP Site form opens.

4

Click Show IGMP Multicast Reporting Statistics. The IGMP IGMP Stats Show IGMP Multicast Reporting Statistics form opens.

5

Configure the IP address.

6

Save your changes and close the form.

END OF STEPS

28.109 To view IGMP source statistics

28.109.1 Purpose

Use the following procedure to view IGMP source statistics. See the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting statistics.

28.109.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The Service (Edit) form opens.

3

Click on the Sites tab to view a list of sites.

4

Choose a site and click Properties. The Site (Edit) form opens.

5

Click on the Multicast tab.

6

Click on the Group Interfaces tab to display a list of group interfaces.

7

Choose a group interface and click Properties. The IGMP Group Interface (Edit) form opens.

8

Click on the Group Interface SAP tab.

9

Choose a group interface SAP and click Properties. The Group Interface SAP (Edit) form opens.

10

Click on the Group tab.

11

Choose a group and click Properties. The Group Interface SAP Group (View) form opens.

12

Click on the Source tab.

13

Click Search to display a list of sources.

14

Close the Group Interface SAP Group (View) form.

15

Click on the Statistics tab.

16

Click Collect to collect additional statistics.

17

Choose a statistics record and click Properties. The Statistics Record - Group Interface SAP Stats form opens.

18

Review the statistics and close the Statistics Record - Group Interface SAP Stats form.

19

Close the forms.

END OF STEPS

MSDP configuration workflow and procedures

28.110 MSDP configuration overview

28.110.1 Overview

The MSDP command hierarchy consists of three levels:

- global level
- peer level and group level
- group peer level

MSDP parameters are initially applied at the global level. The parameters are inherited by the group and peer levels. Parameters can be modified and overridden on a level-specific basis.

The following procedures describe how to configure MSDP.

28.111 Workflow to configure MSDP

28.111.1 Stages

1

Enable MSDP on a routing instance; see 28.112 “To enable MSDP on a routing instance” (p. 1026) .

2

Configure global MSDP on a routing instance; see 28.113 “To configure global-level MSDP” (p. 1026) .

3

Create at least one group-level MSDP peer or peer-level MSDP; see 28.114 “To configure group-level MSDP” (p. 1027) or 28.115 “To configure peer-level MSDP” (p. 1029) .

4

Configure group-peer-level MSDP as required; see 28.116 “To configure group-peer-level MSDP” (p. 1030) .

5

As required, configure an MSDP source; see 28.117 “To configure an MSDP source” (p. 1031) .

6

As required, enable or disable MSDP peering; see 28.118 “To enable or disable MSDP peering” (p. 1032) .

28.112 To enable MSDP on a routing instance

28.112.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Enable MSDP.

Perform the following steps:

1. Click on the Multicast tab.
2. Select the MSDP Enabled parameter.

4

Save your changes and close the form.

END OF STEPS

28.113 To configure global-level MSDP

28.113.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP.

2

Right-click on the MSDP icon and choose Properties. The MSDP (Edit) form opens.

3

Configure the required parameters.

4

Click on the Group tab to add a group. See 28.114 “To configure group-level MSDP” (p. 1027) for more information about how to configure a group-level MSDP.

5

Click on the Peer tab to add a peer. See 28.115 “To configure peer-level MSDP” (p. 1029) for more information about how to configure a peer-level MSDP.

6

Click on the Source tab to add a source. See 28.117 “To configure an MSDP source” (p. 1031) for more information about how to configure an MSDP source.

7

Click on the Import Policies tab.

8

Configure the required parameters.

Configure the import route policy to determine which routes are accepted from peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

9

Click on the Export Policies tab.

10

Configure the required parameters.

Configure the export route policy to determine which routes are advertised to peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . A router performs no validation to ensure the policies match.

11

Click on the following tabs to view information:

- Data Source Active
- Data Source Active Rejected
- Statistics
- Faults

12

Save your changes and close the form.

END OF STEPS

28.114 To configure group-level MSDP

28.114.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP.

2

Right-click on the MSDP icon and choose Create Group. The MSDP Peer Group (Create) form opens.

3

Configure the required parameters.

4

Click on the Peer tab to add a peer.

5

Click Create. The MSDP Peer Group (Create) form opens.

6

Configure the required parameters.

7

Configure the import and export policies for the MSDP peer.

Perform the following steps:

1. Click on the Import Policies tab.
2. Configure the required parameters.

   Configure the import route policy to determine which routes are accepted from peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.
3. Click on the Export Policies tab.
4. Configure the required parameters.

   Configure the export route policy to determine which routes are advertised to peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

8

Click on the Authentication tab.

9

Configure the required parameters.

10

Save your changes and close the form.

11

Configure the import and export policies for the MSDP group.

Perform the following steps:

1. Click on the Import Policies tab.
2. Configure the required parameters.

   Configure the import route policy to determine which routes are accepted from peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.
3. Click on the Export Policies tab.
4. Configure the required parameters.

   Configure the export route policy to determine which routes are advertised to peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

12

Save your changes and close the form.

END OF STEPS

28.115 To configure peer-level MSDP

28.115.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP.

2

Right-click on the MSDP icon and choose Create Peer. The MSDP Peer (Create) form opens.

3

Configure the required parameters.

4

Configure the import and export policies.

Perform the following steps:

1. Click on the Import Policies tab.
2. Configure the required parameters.

   Configure the import route policy to determine which routes are accepted from peers. The policies should match the policies you set when configuring the Routing Policy Manager, as

described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

3. Click on the Export Policies tab.
4. Configure the required parameters.

   Configure the export route policy to determine which routes are advertised to peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

5

Click on the Authentication tab.

6

Configure the required parameters.

7

Save your changes and close the form.

END OF STEPS

28.116 To configure group-peer-level MSDP

Note: The parameters that you configure for a MSDP peer take precedence over the parameters that are configured for the MSDP peer group.

28.116.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP.

2

Right-click on the MSDP icon and choose Create Group. The MSDP Group Peer, MSDP Peer Group (Create) form opens.

3

Configure the required parameters.

4

Configure the import and export policies.

Perform the following steps:

1. Click on the Import Policies tab.
2. Configure the required parameters.

Configure the import route policy to determine which routes are accepted from peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

3. Click on the Export Policies tab.
4. Configure the required parameters.

   Configure the export route policy to determine which routes are advertised to peers. The policies should match the policies you set when configuring the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . There is no validation performed by the router to ensure the policies match.

5

Click on the Authentication tab.

6

Configure the required parameters.

7

Save your changes and close the form.

END OF STEPS

28.117 To configure an MSDP source

28.117.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP.

2

Right-click on the MSDP icon and choose Properties. The MSDP (Edit) form opens.

3

Click on the Source tab.

4

Click Create. The MSDP Source (Create) form opens.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

28.118 To enable or disable MSDP peering

28.118.1 Steps

1

Perform one of the following:

a. Enable or disable MSDP peering in a peer group.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP→MSDP Peer Group.
2. Right-click on an MSDP Peer Group icon.

b. Enable or disable MSDP peering in a group peer.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MSDP→MSDP Peer Group→MSDP Group Peer.
2. Right-click on an MSDP Group Peer icon.

2

Choose one of the following menu options:

a. Turn Up to activate

b. Shut Down to deactivate

3

Save your changes and close the form. The state information beside the appropriate MSDP icon changes accordingly.

END OF STEPS

MLD configuration workflow and procedures

28.119 MLD configuration overview

28.119.1 Overview

MLD is supported on the 7950 XRS, on a chassis mode C or D, and on the 7750 SR in chassis mode A or B with mixed mode enabled.

MLD is supported on a VPRN or IES routing instance on the 7950 XRS, the 7750 SR in chassis mode D, and the 7450 ESS in chassis mode D with mixed mode enabled.

The following procedures describe how to configure MLD.

28.120 Workflow to configure MLD

28.120.1 Stages

1

Enable MLD on a routing instance; see 28.121 “To enable MLD on a base routing instance” (p. 1034) .

2

Configure MLD on a routing instance or VPRN routing instance; see 28.122 “To configure MLD on a base routing instance or VPRN routing instance” (p. 1034) .

3

Create an MLD interface on a routing instance or VPRN routing instance; see 28.123 “To configure an MLD interface on a base routing instance or VPRN routing instance” (p. 1035) . Create an MLD interface on an IES L3 access interface; see 28.124 “To configure an MLD interface on an IES L3 access interface” (p. 1037) .

4

Configure an MLD policy; see 64.25 “To configure an MLD policy” (p. 1866) . Bind the MLD policy to a subscriber profile; see 64.4 “To configure a subscriber profile” (p. 1840) .

5

Configure an MLD group interface on a base routing instance or VPRN routing instance; see 28.125 “To configure an MLD group interface on a base routing instance or VPRN routing instance” (p. 1039) . Configure an MLD group interface on an IES site; see 28.126 “To configure an MLD group interface on an IES site” (p. 1041) .

6

Configure an MC peer group for MLD; see 40.4 “To configure an MC peer group” (p. 1330) .

28.121 To enable MLD on a base routing instance

28.121.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the Routing Instance icon and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the Multicast tab and enable the MLD Enabled check box.

4

Save your changes and close the form. An MLD icon appears in the navigation tree below the routing instance.

END OF STEPS

28.122 To configure MLD on a base routing instance or VPRN routing instance

28.122.1 Purpose

Perform this procedure to configure MLD on a base routing instance or on a VPRN routing instance.

Note: You must first enable MLD on the base routing instance of the VPRN site. See 28.121 “To enable MLD on a base routing instance” (p. 1034) .

You must ensure that the VPRN routing instance has an associated L3 access interface, and that the IPv6 Allowed parameter on the interface is enabled. You must add an IPv6 address to the L3 access interface. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) and 79.102 “To assign an IP address to a VPRN L3 access interface” (p. 2677) .

MLD is supported on a VPRN routing instance on the 7950 XRS, the 7750 SR in chassis mode D, and the 7450 ESS in chassis mode D with mixed mode enabled.

28.122.2 Steps

1

Perform one of the following:

a. Configure MLD on a base routing instance.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MLD.

2. Right-click on the MLD icon and choose Properties. The MLD, Routing Instance (Edit) form opens.

b. Configure MLD on a VPRN routing instance.

Perform the following steps:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
3. On the VPRN Service tree, click on the site to which you need to configure MLD.
4. Expand Routing Instance - NE System ID→Routing Instance→Routing Instance *name*→Protocols.
5. Right-click on the Protocols icon and choose Create MLD Site. The MLD Site (Create) form opens.

2

Configure the required parameters.

3

Click Apply to save the changes.

4

Click on the SSM Translation tab, and click Create. The SSM Translation, Routing Instance (Create) form opens.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

28.123 To configure an MLD interface on a base routing instance or VPRN routing instance

28.123.1 Steps

1

Perform one of the following:

a. Create an MLD interface on a base routing instance.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MLD.

2. Right-click on the MLD icon and choose Create Interface. The MLD Interface, Routing Instance (Create) form opens.

b. Create an MLD interface on a VPRN routing instance.

Perform the following steps:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.
3. On the VPRN Service tree, click on the site to which you need to add the routing instance.
4. Expand Routing Instance - NE System ID→Routing Instance→Routing Instance *name*→Protocols.
5. Right-click on the MLD instance and choose Create MLD Interface. The MLD Interface, Routing Instance (Create) form opens.

2

Configure the required parameters on the General tab.

In the Interface panel, make sure you choose a network interface that has IPv6 enabled.

3

Click on the Behavior tab and configure the required parameters.

4

Click on the Import Policy tab and choose an import policy.

5

Click on the Multicast CAC tab to configure multicast CAC for the MLD interface.

Perform the following steps:

1. On the General tab, select a multicast CAC policy and configure the required parameters.

   The Unconstrained Bandwidth parameter value must be greater than or equal to the Mandatory Bandwidth parameter value.
2. Click on the Levels tab to configure multicast CAC levels.
3. Click Create or select an existing multicast CAC level item and click Properties. The Multicast CAC Level (Create|Edit) form opens.
4. Configure the Level ID and Bandwidth parameters.

   The Level ID parameter can only be configured upon multicast CAC level creation.

   The Lower Level Bandwidth parameter value must be greater than the Higher Level Bandwidth parameter value.
5. Save your changes and close the form.
6. Click on the Lag Port Down tab to configure LAG port down entries.

7. Click Create or select an existing LAG port down entry and click Properties. The LAG Port Down (Create|Edit) form opens.
8. Configure the Number Of Ports Down and Level parameters.

   The Number Of Ports Down parameter can only be set upon LAG port down object creation.
9. Save your changes and close the form.

6

Click on the Static/Group Source tab and create an MLD Interface Group Source Static.

7

Click on the Static Range Group tab and create a static group range, if required.

8

Click on the SSM Translation tab and create an MLD Interface Source Specific Multicast.

9

Save your changes and close the form.

END OF STEPS

28.124 To configure an MLD interface on an IES L3 access interface

28.124.1 Purpose

Perform this procedure to create an MLD interface on an IES L3 access interface. You must ensure that MLD is enabled on the site NE. MLD is supported on an IES routing instance on the 7950 XRS, the 7750 SR in chassis mode D, and the 7450 ESS in chassis mode D with mixed mode enabled. See 28.121 “To enable MLD on a base routing instance” (p. 1034) .

28.124.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the IES Service tree, expand Sites→Site *name*→L3 Access Interfaces.

4

Right-click on an L3 access interface icon and choose Properties. The IES L3 Access Interface (Edit) form opens.

5

Ensure that the IPv6 Allowed parameter is enabled.

6

Click on the Addresses tab and perform the following tasks:

Perform the following steps:

1. Click Create. The IP Address (Create) form opens.
2. Configure the Address ID parameter with a value greater than 1.
3. Enter an IPv6 address for the IP Address parameter.
4. Configure the Prefix Length parameter.
5. Configure the required parameters in the IPv6 panel.
6. Save your changes and close the form.

7

Click on the Multicast tab and perform the following tasks:

Perform the following steps:

1. Click Add. The Create Interface form opens.
2. Choose MLD and click OK. The Create Interface form closes and the MLD Interface (Create) form opens.
3. Configure the Description parameter.
4. Click Select in the Interface panel. The Select Interface - MLD Interface - Routing Instance form opens.
5. Choose a network interface that has IPv6 enabled and click OK. The Select Interface - MLD Interface - Routing Instance form closes.
6. Configure the required parameters in the States panel.

8

Click on the Behavior tab and configure the required parameters.

9

Click on the Import Policy tab and choose an import policy.

10

Click on the Static/Group Source tab and perform the following tasks:

Perform the following steps:

1. Click Create. The MLD Interface Group Source Static (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

11

Click on the SSM Translation tab and perform the following tasks:

Perform the following steps:

1. Click Create. The MLD Interface Source Specific Multicast (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

12

Save your changes and close the form.

END OF STEPS

28.125 To configure an MLD group interface on a base routing instance or VPRN routing instance

28.125.1 Steps

1

Perform one of the following:

a. Create an MLD group interface on a base routing instance.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance.
2. Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form appears.
3. Click on the Multicast tab, select an MLD site and click Properties. The MLD Site (Edit) form appears.
4. Click on the Group Interfaces tab.
5. Click Create or select an existing group interface and click Properties. The MLD Group Interface (Create|Edit) form appears.

b. Create an MLD group interface on a VPRN routing instance.

Perform the following steps:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPRN service and click Properties. The VPRN Service (Edit) form appears.
3. On the VPRN Service tree, click on the site to which you need to add the MLD group interface.
4. Expand Routing Instance - NE System ID→Routing Instance→Routing Instance *name*→Protocols.
5. Right-click on the MLD icon and choose Properties. The MLD Site form appears.
6. Click on the Group Interfaces tab.
7. Click Create or select an existing group interface and click Properties. The MLD Group Interface (Create|Edit) form appears.

2

Configure the required parameters.

3

Select a forwarding service.

4

Select an interface ID.

5

Click on the Behavior tab and configure the required parameters.

If any (or all) of the three query interval parameters are configured, the following validation rules apply:

1. The Query Interval parameter value must be greater than the Last Listener Query Interval and Response Query Interval parameter values.
2. Setting any (or all) of the three Query Interval parameter values to zero on the MLD group interface has the same effect as enabling the Inherit From Site check box for the parameter: the corresponding query interval value configured on the MLD site is used for the MLD group interface.
3. In cases where a mix of MLD site-configured and MLD group interface-configured query interval parameters are in use, validation rule 1 still applies.

   For example, if the Query Interval parameter is NOT configured on the MLD group interface but the Last Listener Query Interval and Response Query Interval parameters ARE configured on the MLD group interface, the MLD site-configured Query Interval parameter value must be greater than the MLD group interface-configured Last Listener Query Interval and Response Query Interval parameter values.

6

Click on the Multicast CAC tab to configure multicast CAC for the MLD group interface.

Perform the following steps:

1. Select a multicast CAC policy.
2. Configure the required parameters.

7

Save your changes and close the form.

END OF STEPS

28.126 To configure an MLD group interface on an IES site

28.126.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form appears.

3

On the navigation tree, expand the Sites icon and click on the site to which you want to add an MLD group interface. The IES Service form is updated with the IES site parameters.

4

Click on the Multicast tab.

5

Click on the Group Interfaces tab and click Create. The MLD Group Interface (Create) form appears.

6

Configure the required parameters.

7

Click Select next to the Interface ID field and choose an interface from the Select Group Interface form.

8

Click on the Behavior tab and configure the required parameters.

9

Save your changes and close the form.

END OF STEPS

Bridging configuration workflow and procedures

28.127 Workflow to configure bridging on an OmniSwitch

28.127.1 Stages

1

Configure bridging on an OmniSwitch; see 28.128 “To configure bridging on an OmniSwitch” (p. 1043) .

2

Release a violated OmniSwitch LSP port if required; see 28.129 “To release a violated OmniSwitch LPS port” (p. 1047) .

28.128 To configure bridging on an OmniSwitch

CAUTION

Service Disruption

Changing bridge parameter values may affect the spanning tree calculations and trigger a topology change in the network.

Note: Not all steps described in this procedure are supported by all OmniSwitches. If a step in this procedure does not apply, proceed to the next applicable step for your OmniSwitch type.

28.128.1 Steps

1

In the navigation tree Routing view, expand Network→*OmniSwitch NE*→Bridge Instance.

2

Right-click on the Bridge Instance Object icon and choose Properties. The Bridge Instance (Edit) form opens.

3

Configure the required parameters.

4

Perform the following to configure the OmniSwitch with a dynamically learned MAC address to restrict a port's ingress traffic, otherwise go to Step 7 .

Perform the following steps:

1. Click on the Learned Port Security tab and configure the required parameters.

   Note:

   The Status parameter does not appear until you enter a non-zero value for the Learning Time Window (minutes) parameter.

   Click Restart Timer to restart the MAC source learning timer if you need to restart dynamic MAC address learning on a port.
2. Click on the Ports tab and click Create to select ports on which you want to enable learned port security. The Select Port - Bridge form opens.
3. Choose one or more ports and click OK. If you need to apply the same LPS properties to multiple ports but do not need to configure static MAC addresses on the ports, go to Step 6 .

5

Optionally, configure static MAC addresses on the ports.

Note: Static MAC addresses can only be added to LPS-enabled ports individually.
A port must be LPS-enabled and belong to a VLAN before you can add static MAC addresses.
If traffic containing MAC addresses outside of the allowable MAC address range attempts to access an LPS port, the switch either restricts access to the port for that traffic or shuts the port down to all traffic. When this happens, the port is in an operationally violated state and an alarm is raised. See 28.129 “To release a violated OmniSwitch LPS port” (p. 1047) for information about how to release a violated LPS port.

Perform the following steps:

1. Choose a port and click Properties. The Learned Port Security (Edit) form opens. See 28.130 “To add MAC address range entries to an OmniSwitch LPS port” (p. 1048)for information about how to add static MAC range entries.
2. Click Properties. The Physical Port (Edit) form opens.
3. Configure the required parameters.
4. Click Select in the VLAN Site panel. The Select VLAN Site form opens.
5. Choose a VLAN site and click OK.
6. Save the changes and close the form. The Learned Port Security (Edit) form reappears.
7. Optionally choose a port and click Convert to Static to stop the aging out of dynamic MAC addresses on LPS ports.
8. Repeat Step 5 if you need to add another static MAC address to the port.

6

Optionally, apply the same LPS properties to multiple ports.

Perform the following steps:

1. Choose multiple ports and click Properties. The Learned Port Security (Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the forms. The Bridge Instance (Edit) form reappears.

7

Click on the Spanning Tree tab to configure the STP modes to the bridge. This prevents bridge loops and the broadcast radiation that results from them.

a. To apply a STP Flat Mode to the bridge; go to Step 8 .

b. To apply a STP 1x1 Mode to the bridge; go to Step 10 .

8

CAUTION

Service Disruption

Service Disruption

Changing the Protocol parameter to MSTP resets the flat bridge priority and path.

Click on the STP Flat Mode tab and choose a 1x1 instance and click Properties. The VLAN STP Instance (Edit) form opens.

1. Configure the required parameters.
2. Click on the Port tab and choose one or more ports.
3. Click Properties. The CIST Instance Ports (Edit) form opens.
4. Configure the required parameters.
5. Save your changes.

9

Optionally, click on the MSTI tab to configure a multiple spanning tree instance to represent represents a group of VLANs, otherwise go to Step 11 .

Note: In order to configure MSTI, the Protocol parameter must be set to MSTP.

Perform the following steps:

1. Click Create. The MST instance form opens.
2. Configure the required parameters.
3. Click on the VLAN tab and click Create to assign a VLAN to the MSTI. The Select VLAN Sites - MST Instance form opens.

4. Choose one or more VLANs and click OK.
5. Save your changes and close the forms. The Bridge Instance (Edit) form reappears.
6. Click on the MST Region tab and configure the required parameters. Go to Step 11 .

10

CAUTION

Service Disruption

Changing the STP 1x1 Mode configuration may affect the STP calculations for this instance of the VLAN and trigger a topology change in the network.

Click on the STP 1x1 Mode tab and choose a 1x1 instance and click Properties. The VLAN STP Instance (Edit) form opens.

1. Configure the required parameters.
2. Click on the Port tab and choose one or more ports from the list of ports that have been assigned to VLANs and click Properties. The VLAN STP Instance Ports (Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form. The Bridge Instance (Edit) form reappears.

11

Click on the TLS tab and configure the TLS Mode parameter.

12

Click on the QoS tab and configure the required parameters.

Note: The applied values for the Default Bridged Disposition, Default Routed Disposition, and Default IGMP Disposition parameters should be the same as the configured values.

13

Click on the MVRP tab and configure the required parameters.

14

Click on the IGMP Port Group Limit tab and perform one the following:

a. Configure active (administratively Up) ports:

Perform the following steps:

1. Choose one or more ports and click Properties. The IGMP Port Group Limit (Edit) or IGMP Port Group Limit - (Multiple Instances) (Edit) form opens.

   Note:

   Only ports that are active (administratively Up) appear in the list. You can view the configuration of the inactive ports using the OmniSwitch CLI.

2. Configure the required parameters.
3. Save your changes and close the form. The Bridge Instance (Edit) form reappears.

b. Configure inactive (administratively Down) ports.

Perform the following steps:

1. Click Create. The Select Port - Bridge form opens.
2. Choose an inactive port and click Properties. The Physical Port (Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form. The Bridge Instance (Edit) form reappears.

Note:

As required, click on the Multicast VLAN Port tab in the multicast VLAN IGMP site properties to view the Port Max Group and Port Action parameter information for a multicast VLAN port. You can also view the number of IGMP groups dynamically learned by the port.

15

Save your changes and close the form.

END OF STEPS

28.129 To release a violated OmniSwitch LPS port

28.129.1 Purpose

Perform the following procedure to release a violated LPS port.

Releasing a violated port restores the port to the same operational state it was in before the violation. When a violated port is released, all MAC addresses known to the port are flushed from the switch MAC address table.

28.129.2 Steps

1

In the navigation tree Routing view, expand Network→*OmniSwitch NE*→Bridge Instance.

2

Right-click on the Bridge Instance Object icon and choose Properties. The Bridge Instance (Edit) form opens.

3

Click on the Learned Port Security tab.

4

Click on the Port tab.

5

Select one or more violated LPS ports. Violated LPS ports are highlighted in orange.

6

Click Release Violated Port to release the selected ports.

7

Save your changes and close the form.

END OF STEPS

28.130 To add MAC address range entries to an OmniSwitch LPS port

28.130.1 Overview

Perform the following procedure to add a static MAC address to an OmniSwitch LPS port.

You can add up to eight MAC address range entries. The default entry with Low MAC Range value 00-00-00-00-00-00, and High MAC Range value FF-FF-FF-FF-FF-FF, is auto-created under LPS Learned MAC Ranges tab when a port is added as an LPS port. Overlapping of Low MAC range or High MAC range is not allowed.

28.130.2 Steps

1

In the navigation tree Routing view, expand Network→*OmniSwitch NE*→Bridge Instance.

2

Right-click on the Bridge Instance Object icon and choose Properties. The Bridge Instance (Edit) form opens.

3

Choose a port and click Properties. The Learned Port Security (Edit) form opens.

4

Click on the LPS Learned MAC Ranges tab. By default, an entry of Low MAC range and High MAC range is displayed.

5

Click Create. The MAC Range (Create) form opens.

6

Enter the Low MAC Range and High MAC Range values.

7

Click OK to save your changes.

8

Repeat steps Step 5 to Step 7 to add up to eight MAC address range entries.

9

Click Apply to save the changes and close the forms.

END OF STEPS

28.131 To modify MAC address range entries in an OmniSwitch LPS port

28.131.1 Steps

1

In the navigation tree Routing view, expand Network→*OmniSwitch NE*→Bridge Instance.

2

Right-click on the Bridge Instance Object icon and choose Properties. The Bridge Instance (Edit) form opens.

3

Choose a port and click Properties. The Learned Port Security (Edit) form opens.

4

Click on the LPS Learned MAC Ranges tab. A list of Low and High MAC address ranges added on the LPS port is displayed.

5

Select the MAC address range entry to modify. You may either click Properties button or double-click the MAC range entry.

6

You can edit the MAC address range entries to meet the following conditions:

- Overlapping Low or High MAC Ranges are not allowed.
- When the MAC address range entry must contain the Low MAC Range starting from 00-00-00-00-00-00, click on the default entry 00-00-00-00-00-00 to FF-FF-FF-FF-FF-FF to modify the High MAC Range
- To expand the LPS MAC range, keep the Low MAC Range the same and modify the High MAC Range to a higher value than the current High MAC Range value
- To shrink the LPS MAC range, keep the Low MAC Range the same and modify the High MAC Range to a lower value than the current High MAC Range value
- Overlapping low or high MAC range is not allowed.

7

Click OK to save your changes.

8

If required, repeat steps Step 5 to Step 7 to modify the MAC range entries.

9

If required, repeat steps Step 3 to Step 7 to modify the MAC range entries for other LPS ports.

10

Click Apply to save the changes and close the form.

END OF STEPS

28.132 To Delete MAC address range from an OmniSwitch LPS port

28.132.1 Steps

1

In the navigation tree Routing view, expand Network→*OmniSwitch NE*→Bridge Instance.

2

Right-click on the Bridge Instance Object icon and choose Properties. The Bridge Instance (Edit) form opens.

3

Choose a port and click Properties. The Learned Port Security (Edit) form opens.

4

Click on the LPS Learned MAC Ranges tab. A list of Low and High MAC address ranges added on the LPS port is displayed.

5

Select one or multiple MAC address range entries. Click Delete.

Note: You can delete a newly added MAC address range entry that is not in the NFM-P database.

6

Confirm your action. Click OK to save the change and close the form.

END OF STEPS

WPP configuration workflow and procedures

28.133 Workflow to configure WPP

28.133.1 Stages

1

Enable WPP on the following routing instances as required:

a. Base routing instance; see 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

b. VPRN routing instance; see 79.21 “To enable routing protocols on a VPRN site” (p. 2554) .

2

Create a Web portal routing instances; see 28.134 “To create a web portal routing instance” (p. 1051) .

3

If required, configure a local user database for WPP-based host authentication. Bind the local user database to a WPP-enabled base routing instance or VPRN routing instance, and a Web portal; see 74.9 “To configure a local user database for subscriber host authentication” (p. 2025) .

4

Bind a web portal to a group interface as follows:

a. IES group interface; see 78.19 “To configure a group interface on an IES” (p. 2449) .

b. VPRN group interface; see 79.37 “To configure a group interface on a VPRN” (p. 2586) .

28.134 To create a web portal routing instance

28.134.1 Steps

1

Enable WPP on a routing instance as described in 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

2

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→WPP.

3

Right-click on the WPP icon and choose Create Web Portal. The Web Portal - Routing Instance (Create) form appears.

4

Configure the required parameters.

5

Click on the WPP Hosts tab to view WPP hosts connected to the Web portal, and to manually clear hanging WPP hosts in the logout state.

Perform the following steps:

1. Select a host item in the list and click Clear Logging Out WPP Host.
2. To clear the entire list, click Clear All Logging Out WPP Hosts.

6

Save your changes and close the form.

END OF STEPS

BIER configuration workflow and procedures

28.135 Workflow to configure BIER

28.135.1 Stages

1

Configure a BIER template on a routing instance; see 28.136 “To configure a BIER template on a routing instance” (p. 1053)

2

Apply the BIER template as needed:

a. Bind the template to an IS-IS site in the Level 1 or Level 2 tab; see 28.63 “To configure an IS-IS interface” (p. 958) .

b. Bind the template to an OSPFv2 site on a routing instance; see 28.67 “To create an OSPF area” (p. 963) .

c. Select BIER as a tunnel type on a VPRN interface; see 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003) .

28.136 To configure a BIER template on a routing instance

28.136.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance and choose Properties. The Routing Instance (Edit) form opens.

3

Configure the Fast Reroute and the BFD Enabled parameter.

4

From the BIER tab, click Create. The BIER (Create) form opens.

5

From the BIER form, perform the following to configure BIER:

1. Click Create to create a template
2. In the BIER Template (Create) form, enter a template name and click Create to create a subdomain.

3. Configure the parameters in the BIER Sub-Domain (Create) form and click OK.
4. Create and configure additional sub-domains if needed.

6

Click OK in each form to save your changes and close the form.

END OF STEPS

IPSec configuration workflow

28.137 Workflow to configure IPSec

28.137.1 Stages

1

Configure an IPSec security policy on the NE; see Step 35 in 27.2 “To configure a routing instance or a VRF instance” (p. 826).

2

Create an IPSec instance on a routing instance.

Perform the following steps:

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance.
2. Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.
3. Click on the IPSec tab and click Create. The Ipsec, (New Instance) (Create) form opens.
4. Configure the required parameters.

3

Configure IPsec Tunnels on the IPSec instance.

Perform Step 6 to Step 19 of 34.21 “To configure an IPsec tunnel on a VPRN tunnel interface” (p. 1253).

29 OpenFlow

OpenFlow overview

29.1 OpenFlow overview

29.1.1 Overview

OpenFlow is a communication protocol for managing multiple NE forwarding planes using an external controller. The direct external control enables traffic management that is superior to the use of local ACL filters and routing information alone.

As an OpenFlow controller, the NFM-P uses a flow table on each switch to manage flows of routing information. Each entry in a flow table is a set of match criteria, packet counters, and the action to perform on matching packets.

29.2 OpenFlow switches

29.2.1 Overview

An NFM-P-managed OpenFlow switch includes the following:

- one flow table
- one communication channel for each OpenFlow controller of the switch
- the OpenFlow protocol

The NFM-P uses SNMP to create and enable an OpenFlow switch instance on an NE. After the initial configuration, the controller and switch communicate using the OpenFlow protocol over an in-band or out-of-band connection. One NE can have up to eight OpenFlow switch instances.

The following figure shows an NFM-P system that manages two OpenFlow switches on separate NEs.

Figure 29-1 Basic OpenFlow management topology

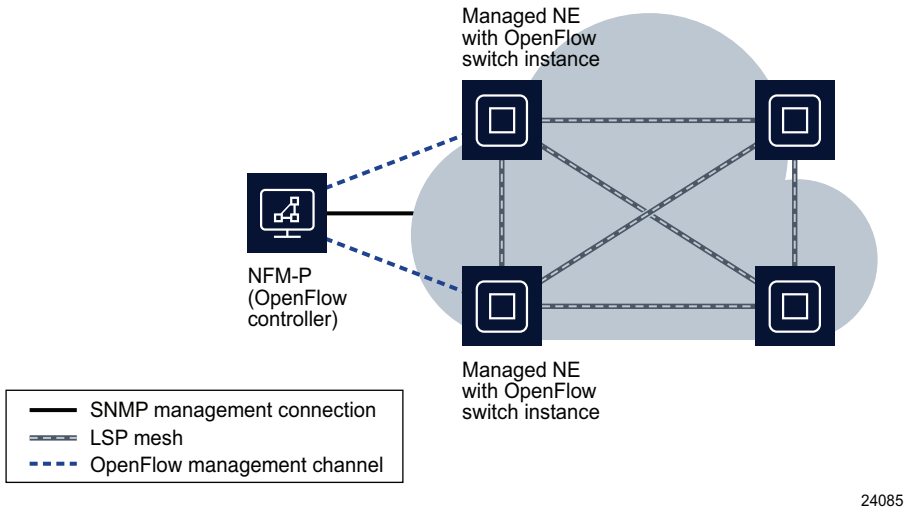

29.3 Configuration

29.3.1 Overview

You can use an NFM-P GUI or OSS client to configure an OpenFlow switch on a 7450 ESS, 7750 SR, or 7950 XRS that has an IOM3.

As part of the OpenFlow switch configuration, an NFM-P operator defines one or more controllers using the IPv4 address of the NFM-P OpenFlow communication interface. The operator also specifies the ports that the switch uses to forward traffic; each port connects to an MPLS-TP or RSVP-TE LSP. IPv4 or IPv6 ACL filters that the operator creates are embedded in the switch logic, and flow entries specify the action to perform on matching packets.

Note the following:

- An IPv6 OpenFlow switch is configurable only on an NE in chassis mode D that has mixed mode enabled.
- SR-TE LSP is not supported for OpenFlow.

29.4 Operation and management

29.4.1 Operation

The OpenFlow protocol uses unsecured TCP as the transport. Using OpenFlow, a controller adds, updates, and deletes flow entries in response to changing network conditions, and retrieves OpenFlow performance and forwarding statistics.

29.4.2 Management

The Manage OpenFlow form is the interface for all GUI management of OpenFlow. Using the form, you can display and manage the following:

- OpenFlow channel states
- flow tables and flow entries
- ACL filter bindings to switch instances
- LSPs associated with a switch
- SNMP statistics for OpenFlow ports
- OpenFlow operation and forwarding statistics

The properties form of an OpenFlow switch displays the following:

- controller communication parameters
- OpenFlow host device type and software release
- OpenFlow configuration information

Note: OpenFlow information is displayed only when there is an active OpenFlow session between the switch and controller.

From an IPv4 or IPv6 ACL filter properties form, you can list and view the associated OpenFlow filter.

29.4.3 Multiple controllers

You can create up to two OpenFlow controller definitions for each switch instance. The controllers have no awareness of each other, but have access to all switch information, including the flow information associated with the other controller.

The following figure shows two NFM-P systems that act as independent OpenFlow controllers of multiple switch instances on one NE.

Figure 29-2 Dual OpenFlow controllers

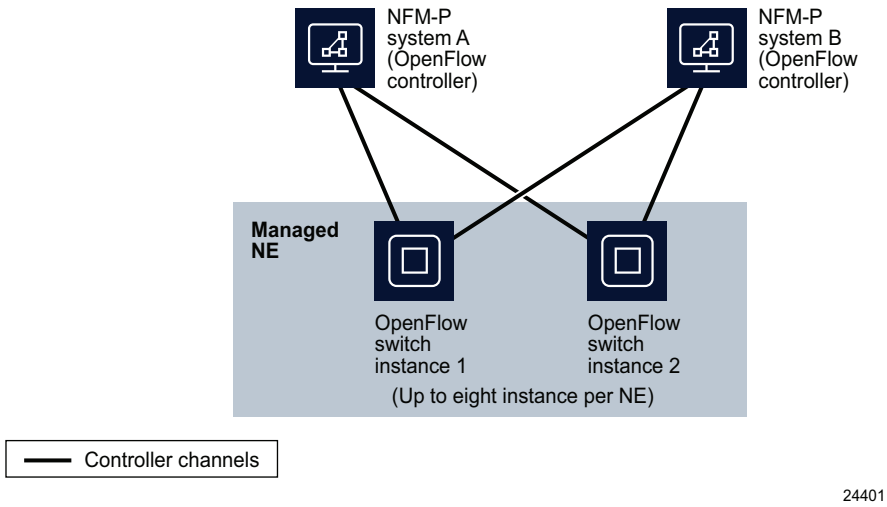

29.4.4 OpenFlow statistics

The NFM-P collects the following statistics from each OpenFlow switch:

- port statistics collected using SNMP
- aggregated packet counts per flow table
- aggregated packet and byte counts per flow entry

OpenFlow configuration and management procedures

29.5 OpenFlow configuration and management workflow

29.5.1 Stages

1

Create the IPv4 and IPv6 ACL filters that are to be embedded in the switch logic. See Chapter 51, “Filter policies” for information about creating ACL filters.

2

Configure and enable OpenFlow switch instances on NEs, as required. See 29.6 “To configure an OpenFlow switch” (p. 1061) .

3

Configure additional flow table entries, as required. See 29.7 “To configure an OpenFlow flow table entry” (p. 1063) .

4

Monitor OpenFlow, as required.

a. View the controller channel status and SNMP statistics. See 29.8 “To display the OpenFlow controller channel status and SNMP statistics” (p. 1065) .

b. View the OpenFlow port status and SNMP statistics. See 29.9 “To display the ports and port SNMP statistics of an OpenFlow switch” (p. 1066) .

c. View the OpenFlow flow table statistics. See 29.10 “To display aggregate flow table statistics” (p. 1067) .

d. View the OpenFlow flow table entry statistics. See 29.11 “To display aggregate flow table entry statistics” (p. 1067) .

29.6 To configure an OpenFlow switch

29.6.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Perform one of the following.

a. Create an OpenFlow switch.

Perform the following steps.

1. Click Create→OpenFlow switch. The Select Network Element form opens.

2. Select an NE and click OK. The Select Network Element form closes, and the OpenFlow Switch (Create) form opens.
3. Configure the Displayed Name parameter.
4. Click Apply. The form displays additional tabs and status indicators, and the form name changes to OpenFlow Switch (Edit).

b. Configure an existing OpenFlow switch. Perform the following steps.

Perform the following steps.

1. Configure the filter criteria and click Search. A list of OpenFlow switches is displayed.
2. Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

3

Configure the required parameters.

4

Define an OpenFlow controller.

Perform the following steps.

1. Click on the Controller tab.
2. Click Create. The OpenFlow Controller (Create) form opens.
3. Configure the required parameters.
4. Click OK to save your changes and close the forms.

5

Repeat Step 4 to define an additional controller for the switch.

6

Configure a flow table.

Perform the following steps.

1. Click on the Flow Table tab.
2. Click Create, or select an existing flow table and click Properties. The Flow Table (Edit) form opens.
3. Configure the required parameters.
4. Click OK to save your changes and close the forms.
5. To add one or more entries to the flow table, perform Step 5 to Step 13 of 29.7 “To configure an OpenFlow flow table entry” (p. 1063) .

7

To configure an embedded IPv4 or IPv6 filter, perform the following steps.

Perform the following steps.

1. Click on the Embedded Filter tab.
2. Click Create, or select an existing filter and click Properties. The OpenFlow Embedded IP Filter (Create | Edit) form opens.
3. Configure the required parameters.
4. Choose an ACL IP or IPv6 filter.
5. Click OK to save your changes and close the form.

8

Close the open forms, as required.

END OF STEPS

29.7 To configure an OpenFlow flow table entry

29.7.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Click Search. A list of OpenFlow switches is displayed.

3

Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

4

Click on the Flow Table tab.

5

Select the flow table and click Properties. The Flow Table (Edit) form opens.

6

Perform one of the following.

a. Add a flow table entry. Perform the following steps.

Perform the following steps.

1. Click Create. The OpenFlow Flow Entry (Create) form opens.
2. Configure the Auto-Assign ID and Flow ID parameters.

3. Click Create to add an action. The OpenFlow Flow Action (Create) form opens.
4. Configure the Action Type parameter.
5. If you choose Output Action as the Action Type parameter value, click Select to choose a port.
6. If you choose Redirect to Next Hop as the Action Type parameter value, configure the Next Hop IP Address parameter.
7. Click OK to save the changes and close the OpenFlow Flow Action (Create) form.

b. Configure an existing flow table entry. Perform the following steps.

Perform the following steps:

1. Configure the filter criteria and click Search. A list of flow table entries is displayed.
2. Select an entry and click Properties. The OpenFlow Flow Entry (Edit) form opens.

7

Configure the required parameters:

8

Click on the Match Criteria tab.

Note: You must select the check box beside a parameter on the tab before you can configure the parameter.

9

Configure the Ethernet Type parameter.

Note: The parameter setting must match the IP address type configured in the actions associated with the flow entry.

10

Configure the following parameters:

- Ethernet Type
- IP Source Address
- IP Dest Address
- DSCP
- Protocol Type
- Prefix
- Prefix

11

If the Protocol Type parameter is set to TCP, configure the TCP Source Port and TCP Dest Port parameters:

12

If the Protocol Type parameter is set to ICMPv4 or ICMPv6, configure the ICMP Type and ICMP Code parameters:

13

Click OK to save the changes and close the forms, as required.

14

Close the Manage OpenFlow form.

END OF STEPS

29.8 To display the OpenFlow controller channel status and SNMP statistics

29.8.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

3

Click on the Controller tab.

4

Select the controller and click Properties. The OpenFlow Controller (Edit) form opens.

5

Click on the Channel tab.

6

Select a channel and click Properties. The OpenFlow Controller Channel Info form opens.

7

View the channel status information, as required.

8

Click on the Statistics tab.

9

Click Collect. A statistics record is listed.

10

Select the statistics record and click Properties. The statistics record form opens.

11

View the statistics, as required.

12

Close the open forms.

END OF STEPS

29.9 To display the ports and port SNMP statistics of an OpenFlow switch

29.9.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

3

Click on the Ports tab. A list of ports is displayed.

4

To view a port entry, select the entry and click Properties. The OpenFlow Ports (View) form opens.

5

To view the aggregate statistics of all ports in use by the switch.

Perform the following steps.

1. Click on the Statistics tab.
2. Click Collect to perform an on-demand statistics collection, or the Collect All to retrieve all collected statistics records.
3. Select a statistics record and click Properties. The Statistics Record form opens.
4. View the statistics counter data, as required.

6

Close the open forms.

END OF STEPS

29.10 To display aggregate flow table statistics

29.10.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

3

Click on the Flow Table tab.

4

Select the flow table and click Properties. The Flow Table (Edit) form opens.

5

Click on the Flow Table Counts tab.

6

View the statistics counter data, as required.

Note: OpenFlow statistics information is displayed only when there is an active OpenFlow session between the switch and controller.

7

Close the open forms.

END OF STEPS

29.11 To display aggregate flow table entry statistics

29.11.1 Steps

1

Choose Manage→OpenFlow from the NFM-P main menu. The Manage OpenFlow form opens.

2

Select an OpenFlow switch and click Properties. The OpenFlow Switch (Edit) form opens.

3

Click on the Flow Table tab.

4

Select the flow table and click Properties. The Flow Table (Edit) form opens.

5

Select a flow table entry and click Properties. The Flow Table Entry (Edit) form opens.

6

Click on the Statistics tab.

7

Click Collect to perform an on-demand statistics collection, or the Collect All to retrieve all collected statistics records.

8

Select a statistics record and click Properties. The Statistics Record form opens.

9

Click on the Flow Stats tab.

10

View the statistics counter data, as required.

11

Close the open forms.

12

View the statistics counter data, as required.

Note: OpenFlow statistics information is displayed only when there is an active OpenFlow session between the switch and controller.

13

Close the open forms.

END OF STEPS

29.12 To list the OpenFlow bindings of an IP filter

29.12.1 Steps

1

Choose Policies→Filter→ACL IP Filter or Policies→Filter→ACL IPv6 Filter from the NFM-P main menu. The ACL IP Filter Policies or ACL IPv6 Filter Policies form opens.

2

Click on the Embedded Filters tab.

3

Click on the OpenFlow Embedded tab. A list of embedded OpenFlow filters is displayed.

4

View the list entries, as required.

5

Close the ACL IP Filter Policies or ACL IPv6 Filter Policies form.

END OF STEPS

30 NAT

30.1 Network Address Translation

30.1.1 Overview

Network Address Translation (NAT) rewrites address information in IP packets that travel between private and public networks. NAT effectively extends the public IPv4 address space; multiple end users can share one IP address. NAT provides security by preventing an internal IP address, such as an end-user address, from entering a public network.

NAT translates internal, or private, host IP address and TCP port values to external, or public, values. When NAT receives a packet from an internal host, it replaces the source address information in the packet with public address information, and forwards the packet to the destination. NAT assigns private IP addresses and ports dynamically using values from an allocated pool, but can also use values that are statically assigned to internal hosts, depending on the type of NAT deployment.

The NFM-P NAT support includes:

- policy-based management
- subscriber-specific implementation
- configurable port usage limits per subscriber, address range, or policy
- reserved port ranges for specified forwarding classes that are exempt from port usage limits
- configurable protocol timeout periods for efficient resource management

The NAT configuration on a routing instance includes:

- one or more NAT address pools that are associated with local ISA-NAT groups
- a NAT policy that specifies port ranges and operational parameters, and optionally, an IPFIX export policy that specifies IP traffic flow export information
- host address match criteria
- one or more NAT destinations

An NFM-P operator assigns port ranges in an address pool configuration. Port ranges specify the number of ports that are allocated to a subscriber for mapping to host sessions. If all ports for a subscriber are in use, additional port assignments cannot occur and host sessions are rejected. This function helps to prevent the flooding of NAT by a virus attack or multiple peer-to-peer file transfers. You can override a port range limit by configuring a range of reserved ports that are assigned based on the traffic forwarding class.

Note: You cannot use ports 0 to 1023, which are called the well-known ports or privileged ports, in a NAT configuration.

To provide equitable and timely NAT resource allocation to hosts, you can specify timeout values in a NAT policy for protocols such as ICMP, TCP, and UDP. NAT ends a host session after the timeout period, for example, when a TCP handshake takes an excessive amount of time.

The NAT drain function is a mechanism that is used to gracefully remove the host sessions associated with an address range in a NAT address pool. When an address range is in the draining state, NAT drops new session requests for the address range. After an existing session associated with a draining address range closes, a new session for the same host is created using a different address range. The drain function removes an address range only when there are no sessions associated with the address range. See 30.13 “To start or stop a NAT address-pool drain operation” (p. 1094) for information about using the NAT drain function.

You can view general NAT statistics, for example, address pool allocation and session counts for different protocols on the Statistics tab of an ISA-NAT group member properties form.

Note: The NFM-P does not record detailed NAT session information such as session duration or packet counts. You can obtain this information from a request to a participating NE.

The NFM-P supports the following NAT deployment types:

- L2-aware—Hosts of the same subscriber can share a private address, and are assigned public addresses from a pool in a base routing instance.
- large-scale—Each host can have a unique static or dynamic private address, and is assigned a public address from a pool in a VPRN routing instance.

30.1.2 L2-aware NAT

In L2-aware NAT, a NAT policy is associated with a subscriber profile that is applied to an IES SAP or a VPRN routing instance. The policy specifies an address pool on the base routing instance. When DHCP assigns a private IP address that is in an L2-aware NAT address range, NAT assigns a public IP address to the host packets. See Chapter 78, “IES management” for information about configuring NAT in an IES. See Chapter 79, “VPRN service management” for information about configuring NAT on a VPRN.

Note: A NAT configuration on a base routing instance applies to each IES site or VPRN routing instance on the NE.

A NAT policy that is specified in a subscriber profile redirects the IPv4 traffic for the subscriber to NAT.

External address assignment

The External Assignment option on L2–aware NAT pools allows the assignment of outside IP addresses to a subscriber by an external system (for example, a RADIUS server). The External Assignment configuration ensures that all traffic to or from a particular L2-aware subscriber is received on the same ISA card when outside IP addresses are allocated by an external system.

30.1.3 Large-scale NAT

Large-scale NAT, or LSN, is required when each host in a customer VPRN service has a unique private IP address and requires a unique public IP address, for example, in mobile network deployments. Large-scale NAT is used between routing instances. A customer VPRN routing instance provides host access and forwards packets through NAT to a VPRN or through an IES to an NE routing instance, which provides public network access.

A VPRN routing instance requires an ACL IP filter or static route to transmit host traffic through the NAT function, and uses an address pool in the VPRN routing instance. You can configure large-scale NAT to statically and dynamically assign private addresses. See Chapter 79, “VPRN service management” for information about configuring NAT in a VPRN service. See 30.1.4 “Static port forwarding” (p. 1072) for information about configuring NAT to use static private addresses.

You can use the NFM-P Statistics Plotter to plot real-time LSN subscriber host statistics on a base or VPRN routing instance. See 30.16 “To plot LSN subscriber host statistics” (p. 1097) .

30.1.4 Static port forwarding

In a large-scale NAT deployment, you can configure NAT to assign static private addresses to subscriber hosts using static port forwarding. Static port forwarding ensures that an internal host uses the same private IP address and port each time they connect to the network. You can use the NFM-P to configure static port forwarding for TCP and UDP independently.

A standby routing instance can be specified to create static port forwards for pools in a redundant configuration. The NFM-P identifies the active pool from the two specified NEs and creates a static port forward on the active pool. In the event of a wildcard static port forward request, if there is a switch-over of the active pool, the NFM-P attempts to create the static port forward on the other pool. If there is a second pool switch-over, the static port forward is not created.

See Chapter 79, “VPRN service management” for information about configuring static port forwarding in a VPRN service. See Chapter 74, “Residential subscriber management” for information about configuring L2Aware static port forwarding on a residential subscriber instance.

30.1.5 Base router and VPRN static one-to-one NAT

With static one-to-one NAT, NAT is performed on packets traveling from an inside (private) interface to an outside (public) interface or from an outside interface to an inside interface. Static one-to-one NAT can be applied to a single IP address or a subnet of IP addresses and is performed on the IP header of a packet, not on the UDP/TCP port.

Mapping statements, or entries, can be configured to map an IP address range to a specific IP address. The direction of the NAT mapping entry dictates whether NAT is performed on a packet source IP address or subnet, or on a packet destination IP address or subnet. The 7705 SAR supports inside mapping entries that map an inside IP address range to an outside IP address range sequentially.

With an inside mapping entry, consider the following:

- Packets that originate from an inside interface and are destined for an inside interface are forwarded without any NAT being applied.
- Packets that originate from an outside interface and are destined for an outside interface are forwarded without any NAT being applied.
- If there is a matching one-to-one NAT mapping entry, packets that originate from an inside interface and are destined for an outside interface undergo static one-to-one NAT where NAT changes the source IP address of the packet IP header. The packet is forwarded whether or not a NAT mapping entry is found unless the Drop Packets without NAT parameter is enabled. When a mapping entry is not found and the Drop Packets without NAT parameter is enabled, the packet is not forwarded.
- If there is a matching one-to-one NAT mapping entry, packets that originate from an outside

interface and are destined for an inside interface undergo static one-to-one NAT where NAT changes the destination IP address of the packet IP header. The packet is forwarded whether or not a NAT mapping entry is found unless the Drop Packets without NAT parameter is enabled. When a mapping entry is not found and the Drop Packets without NAT parameter is enabled, the packet is not forwarded.

Static one-to-one NAT is performed on packets that transit the node and match the mapping entry. These packets include IPSec packets, GRE packets, and IP packets. NAT can be performed on packets from a single inside interface or multiple inside interfaces that are traveling to a single outside interface or multiple outside interfaces.

Static one-to-one NAT is not performed on packets that are destined for the node, on self-generated traffic, or on routing protocols. The 7705 SAR blocks static one-to-one NAT to a public prefix that has the same IP subnet as a local interface.

Static one-to-one NAT is supported on the GRT and on VPRNs.

The following table lists the types of outside and inside interfaces that are supported on the GRT for static one-to-one NAT.

Table 30-1 GRT interfaces supported for static one-to-one NAT

GRT interface type	Outside	Inside
Network interface	Yes	No
IES interface	Yes	Yes
IES R-VPLS interface	Yes	Yes
IES Layer 3 spoke SDP interface	Yes	Yes
IPSec public interface	n/a	n/a

VPRN static one-to-one NAT

For VPRNs, static one-to-one NAT can be configured between an inside interface and a outside MP-BGP MPLS transport tunnel interface. See the 7705 SAR NE documentation for more information about VPRNs and one-to-one NAT.

The following table lists the types of outside and inside interfaces that are supported on a VPRN for one-to-one NAT.

Table 30-2 VPRN interfaces supported for static one-to-one NAT

VPRN interface type	Outside	Inside
SAP interface	Yes	Yes
R-VPLS interface	Yes	Yes
Layer 3 spoke SDP interface	Yes	Yes
IPSec private interface	Yes	Yes

Table 30-2 VPRN interfaces supported for static one-to-one NAT (continued)

VPRN interface type	Outside	Inside
Auto-bind GRE/MPLS (MP-BGP), where MPLS includes segment routing, LDP, and RSVP	Yes	No

30.1.6 Dynamic block reservation

To prevent starvation of dynamic port blocks for subscribers that use static port forwarding, dynamic port blocks can be reserved for the lifetime of the static port forward. Dynamic port blocks are associated with the same inside IP addresses that are associated with the static port forward. A log entry is not generated until the dynamic port block is used.

At the time of static port forward creation:

- If the corresponding dynamic port block mapping (same inside IP address) does not exist, dynamic port block(s) are reserved and associated with the same inside IP addresses used in the static port forward. No log entries are generated for the reserved dynamic port blocks.
- If corresponding dynamic mappings (relative to the inside IP address) exist, the existing port block is reserved (or kept) after the last mapping within it has expired. The reserved dynamic port block(s) continue to be associated with the same inside IP address until the static port forward is deleted. A log entry is generated when the last mapping in the dynamic port block expires (even if the dynamic port block continues to be reserved).

30.1.7 NAT pools

A NAT pool is a range of public IP addresses that are translated to inside private IP addresses in a Large Scale or L2 Aware NAT configuration. A NAT pool is required in order to configure a NAT policy, as well as static port forwarding, DS Lite, and NAT 64.

In a dual-homing redundant NAT configuration, a group of NAT pools is configured for switch-over. The group is called a Pool Fate Sharing Group (PFSG), in which there is one lead pool, and a group of subordinate pools that follow the lead pool. The subordinate pools inherit the lead pool activity and state. The lead pool has its own export route which must match the monitoring route of all the pools in the corresponding PFSG on the peering node. All pools in a PFSG should belong to the same ISA-NAT group.

30.1.8 Application-agnostic NAT

If a NAT pool is created as application-agnostic, IP addresses are translated in 1:1 fashion, regardless of protocol. Ports are not translated for TCP or UDP traffic. Traffic through the NAT pool can be initiated from inside or outside. When a NAT pool is configured as application-agnostic, certain parameters in the pool are pre-configured and cannot be changed:

- Subscriber Limit: 1
- Deterministic NAT Port Reservation: 65536
- Pool Mode: One-to-One
- Port Reservation Type: Blocks

- Port Reservation Value: 1
- Port Forward Range End: 0

An application agnostic NAT pool is used to configure static 1:1 NAT, where the operator has control of mapping between the inside and outside IP addresses.

ALG for TCP/UDP are supported in a protocol-agnostic NAT pool.

30.1.9 NAT policies

A NAT policy defines general NAT properties and associates a NAT address pool with an ISA-NAT group on the same NE. A NAT policy is associated with a subscriber profile for L2-aware NAT in an IES or VPRN service, or for large-scale NAT in a VPRN service.

A NAT pool must be assigned to the local NAT policy before the NAT policy can be associated with the NAT configuration. A NAT policy is required in order to configure DS Lite, NAT 64, and NAT subscriber identification.

Individual NAT policies can be configured on various objects on an inside NAT configuration in order to create flexible mapping of inside traffic to outside pools, based on traffic criteria. Subscribers are mapped to different NAT pools, based on their source IP.

- A local NAT policy can be configured for each destination prefix. Destination prefixes point to the global NAT policy for a NAT configuration by default.
- A local NAT policy can be configured for each static port forward entry. Static port forwards point to the global NAT policy for a NAT configuration by default.
- ACL IP and IPv6 filter policies can be configured to divert traffic to NAT based on filter match criteria. Each filter entry can point to an individual NAT policy. For IPv6 filter entries, you can specify the NAT type (DS-LITE or NAT-64).
- A local NAT policy is configured on a NAT prefix list for L2-aware NAT configurations. The NAT policy must be configured with a NAT pool.

The subscriber retention timeout configuration specifies the time a NAT subscriber and its associated IP address is kept after all hosts and associated port blocks have expired. If a NAT subscriber host appears before the retention timeout has elapsed, it is assigned the same outside IP address.

30.1.10 NAT firewall policies

A firewall policy is associated with a subscriber profile for IPv6 firewall configurations. The policy is configured with an NE, on which downstream forwarding occurs. Incoming network traffic to the NE is disallowed if there is no matching traffic flow, or if port forwarding does not exist. IPv6 domains are provisioned, containing IPv6 prefixes. Each domain points to a NAT group and the IPv6 prefixes are micronetted over the NAT group members. The IPv6 address is assigned externally, using DHCPv6 relay.

Several components must be configured to function with the firewall policy:

- a local DHCPv6 server
- a firewall domain
- IPv6 prefixes

All of these components can be configured on both base and VPRN routing instances.

30.1.11 NAT prefix lists

A NAT prefix list is used to create a set of mappings between destination IP prefixes and a specific NAT policy for an L2-aware NAT configuration. The NAT prefix list is associated with a subscriber profile, and with a NAT policy.

30.1.12 NAT classifiers

The NAT classifier is referenced by the NAT policy in DNAT configurations. It determines the destination IP address and the type of traffic that is subject to DNAT.

30.1.13 IPFIX export policies

An IPFIX export policy defines how IP traffic flow information is formatted and transferred from an exporter to a collector to facilitate services such as measurement, accounting, and billing. You can specify the router instance type, IPFIX Collector address, MTU, and refresh timeouts. You can add an IPFIX export policy to a NAT policy.

30.1.14 ISA-NAT groups

An ISA-NAT group provides a redundant NAT function for routing instances using ISA Broadband Applications MDAs or ESA VMs.

See the NE documentation for more information about NAT deployment.

30.1.15 NAT Destinations

A NAT destination prefix applies NAT to all traffic that matches a given route. A NAT Destination cannot be configured in conjunction with a NAT policy that points to a NAT pool that resides in the same service. Such a configuration would result in a routing loop.

30.1.16 NAT 64

NAT 64 allows IPv6 hosts to connect to IPv4-only servers, mapping their inside IPv6 addresses to outside IPv4 addresses. NAT 64 configuration requires support for IPv6 (chassis mixed mode) on the NE. You must assign a NAT policy to the NAT configuration before enabling NAT 64. IP fragmentation can be configured to allow IPv6 packet fragmentation before transmission.

30.1.17 DS Lite

DS Lite is intended for network operators that have an IPv6-only access network, but IPv4-only CPEs. The inside IPv4 host information is encapsulated into a 4to6 tunnel that runs over the IPv6 network. At the other end, the IPv4 address is de-encapsulated and translated to an outside IPv4 address via CGN. You must assign a NAT policy to the NAT configuration before enabling DS Lite. IP fragmentation can be configured to allow IPv6 packet fragmentation before transmission if the packets are larger than the tunnel MTU for a given DS Lite address.

30.1.18 Deterministic NAT

Deterministic NAT is a mode of operation in which the inside IP address and source ports are deterministically mapped to the outside IP address and port range at the time of configuration. Each inside IP address is permanently mapped to an outside IP address and a dedicated port block. The dedicated port block is referred to as the deterministic port block. Deterministic mapping can be automatically extended by a dynamic port block in cases where a deterministic port range runs out of ports. Because a simple formula is used to determine the inside IP address based on the outside IP address and port, there is no need for NAT logging.

The NAT deterministic script generates an offline calculation of the deterministic NAT map entries. You can save the script results to a folder on a remote server by specifying a URL for the server in the local NE properties.

In the NFM-P, inside IP addresses are configured in the form of prefixes. Any inside prefix on any routing instance can be mapped to any deterministic pool on any routing instance, including the deterministic pool in which the inside prefix is defined. The mapping between the inside prefix and the deterministic pool is achieved through a NAT policy that is referenced for each individual prefix. IP addresses from the prefixes on the inside are distributed over the IP addresses defined in the outside pool referenced by the NAT policy.

A port block is a collection of ports that is assigned to a subscriber. A deterministic LSN subscriber can have only one deterministic port block that can be extended by multiple dynamic port blocks. All port blocks for an LSN subscriber must be allocated from a single outside IP address.

A port range is a collection of ports that can spawn multiple port blocks of the same type. For example, a deterministic port range includes all ports that are reserved for deterministic consumption. A dynamic port range is a total collection of ports that can be allocated in the form of dynamic port blocks.

30.1.19 Destination NAT

Destination NAT (DNAT) is used for traffic steering when the destination IP address is rewritten. A typical use of DNAT is the redirection of unauthenticated traffic to an authentication server. Once authenticated, DNAT would be removed by means of a RADIUS CoA request.

DNAT is used as part of large-scale NAT (deterministic and non-determistic) and L2–aware NAT configurations. DNAT-only configuration is applicable only to large-scale NAT. When a NAT policy is configured for use in DNAT-only configurations, it must not be configured with port limit or NAT pool parameters. When a NAT policy is configured for use in DNAT configurations, it must reference a NAT classifier, which determines the destination IP address and the type of traffic that is subject to DNAT.

30.2 Workflow to configure NAT

30.2.1 Stages

1

Configure an ISA RADIUS accounting policy that is used to configure an ISA-NAT group; see 57.12 “To configure an ISA RADIUS policy” (p. 1803).

2 Configure an ISA-NAT group; see 30.3 “To configure an ISA-NAT group” (p. 1080).

3 Configure an IPFIX export policy as required and add it to the NAT policy; see 30.4 “To configure an IPFIX export policy” (p. 1081).

4 Configure a NAT classifier, if required; see 30.8 “To configure a NAT classifier” (p. 1085).

5 Configure a NAT policy; see 30.5 “To configure a NAT policy” (p. 1082).

6 Configure a NAT prefix list, if required; see 30.7 “To configure a NAT prefix list” (p. 1084).

7 Associate a NAT policy with a subscriber profile; see 64.4 “To configure a subscriber profile” (p. 1840).

8 Configure NAT on a routing instance; see 30.9 “To configure NAT on a routing instance” (p. 1086).

9 Start or stop a NAT address-pool drain operations; see 30.13 “To start or stop a NAT address-pool drain operation” (p. 1094).

10 Configure a NAT deterministic script on a remote server; see 30.14 “To configure a NAT deterministic script on a remote server” (p. 1095).

11 Configure statistics on an ISA-NAT group; see 30.15 “To configure statistics on an ISA-NAT group” (p. 1096).

12 View LSN subscriber host statistics data in graphical form; see 30.16 “To plot LSN subscriber host statistics” (p. 1097).

13 View reserved IP address and reserved block information on ISA-NAT groups; see 30.17 “To view reserved IP address and reserved block information on an ISA-NAT group” (p. 1098).

14

View ISA-NAT object information such as Classic LSN Subscriber, ISA-NAT Group, or LSN block allocation; see 30.18 “To view ISA-NAT object information” (p. 1099).

30.3 To configure an ISA-NAT group

30.3.1 Purpose

Perform this procedure to configure an ISA-NAT group on supported 7750 SR devices. An ISA-NAT group provides a redundant NAT function for routing instances using ISA broadband applications MDAs or ESA VMs.

30.3.2 Steps

1

In the navigation tree Equipment view, expand Network→*NE*→Logical Group.

2

Right-click on the ISA-NAT Groups object and choose Create ISA-NAT Groups or select an existing ISA-NAT Group and click Properties. The ISA-NAT Group (Create|Edit) form opens.

3

Configure the parameters, as required.

The Redundancy and Failed MDA Limit parameters can be configured only when the Administrative State parameter is set to Down.

4

Select an ISA RADIUS policy. See 57.12 “To configure an ISA RADIUS policy” (p. 1803) for information about how to configure a ISA RADIUS policy.

5

Click Apply. The ISA-NAT Group form refreshes with additional tabs and parameters.

6

When the Hardware Type parameter is set to ISA, perform the following steps:

1. Click on the ISA-NAT MDA tab.
2. Click Create.The ISA-NAT Group MDA (Create) form opens.
3. Select an ISA Broadband Applications MDA. See 15.78 “To configure an MDA” (p. 536) for information about how to configure an MDA.

   NOTE: You can configure up to six ISA-NAT group members. Each group member must be an ISA Broadband Applications MDA
4. Click on the General tab and configure the Administrative State parameter.

NOTE: You can set the Administrative State parameter to Up only when the ISA-NAT group contains at least one MDA and when the Active MDA Limit parameter is set to a value greater than 0.

5. Click on the Dynamic MDA tab.
6. Click Create. The Dynamic MDA (Create) form opens.
7. Configure the required parameters.

7

When the Hardware Type parameter is set to ESA VM, perform the following steps:

1. Click on the NAT ESA VM tab.
2. Click Create. The NAT ESA VM (Create) form opens.
3. Select an ESA ID and an ESA VM ID.

8

Save your changes and close the forms.

END OF STEPS

30.4 To configure an IPFIX export policy

30.4.1 Steps

1

Choose Policies→IPFIX Export Policy from the NFM-P main menu. The IPFIX Export Policy form opens.

2

Click Create or select an existing IPFIX export policy and click Properties. The IPFIX Export Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the IPFIX Collectors tab and click Create. The IpfixCollector form opens.

5

Configure the required parameters.

6

If you specified a Router Instance Type of VPRN, perform the following:

1. Click Select in the VPRN panel. The Select vprnPointer form opens.
2. Select a VPRN pointer and click OK.

7

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

30.5 To configure a NAT policy

30.5.1 Purpose

Perform this procedure to configure a NAT policy for a large-scale or L2-aware NAT implementation.

30.5.2 Steps

1

Choose Policies→ISA Policies→NAT from the NFM-P main menu. The NAT Policies form opens.

2

Click Create or select an existing NAT policy and click Properties. The NAT Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

In the IPFIX Export Policy panel, select an IPFIX export policy, if required.

5

In the DNAT panel, select a NAT classifier, if required.

6

Click Apply. The NAT Policy (Create|Edit) form refreshes with additional tabs.

7

Assign a NAT pool to a local definition of the policy, as required.

Note: The policy must be distributed to the NE on which the NAT pool is configured before the local definition of the policy can be configured. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Perform the following steps:

1. Click on the Local Definitions tab and click Search. A list of NEs to which the policy has been distributed appears.
2. Click on an item in the list and click Properties. The NAT Policy, Local Policy form appears.
3. Ensure that the Distribution Mode field of the local policy is set to Local Edit Only. Click Switch Mode to change the distribution mode if required.
4. Select a NAT pool.

8

If the policy is part of a DNAT-only configuration, associate a local definitions of the policy with a routing instance and an NAT-ISA group, as required.

Note: The policy must be distributed to an NE before the local definition of the policy can be configured. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Perform the following steps:

1. Click on the Local Definitions tab and click Search. A list of NEs to which the policy has been distributed appears.
2. Click on an item in the list and click Properties. The NAT Policy, Local Policy form appears.
3. Ensure that the Distribution Mode field of the local policy is set to Local Edit Only. Click Switch Mode to change the distribution mode if required.
4. In the DNAT-only panel, select a routing instance and NAT-ISA group.

9

Save your changes and close the forms.

END OF STEPS

30.6 To configure a NAT firewall policy

30.6.1 Purpose

Perform this procedure to configure a NAT firewall policy for an IPv6 firewall configuration for vRGW hosts.

30.6.2 Steps

1

Choose Policies→ISA Policies→NAT from the NFM-P main menu. The NAT Policies form opens.

2

Select Firewall Policy (Network Address Translation) in the object drop-down.

3

Click Create or select an existing firewall policy and click Properties. The Firewall Policy (Create|Edit) form opens.

4

Configure the parameters as required.

5

Click on the Timeouts tab to configure timeout parameters as required.

For each timeout category, you must disable the Default check box before specifying non-default timeout values.

6

Click Apply. The Firewall Policy (Create|Edit) form refreshes with additional tabs.

7

Save your changes and close the forms.

End of steps

30.7 To configure a NAT prefix list

30.7.1 Purpose

Perform this procedure to configure a NAT prefix list for an L2-aware or dynamic NAT implementation.

30.7.2 Steps

1

Choose Policies→ISA Policies→NAT Prefix Lists from the NFM-P main menu. The NAT Prefix Lists form opens.

2

Click Create or select an existing NAT prefix list and click Properties. The NAT Prefix List (Create|Edit) form opens.

3

Configure the Displayed Name and Application parameters.

4

Click Apply. The NAT Prefix List form refreshes with additional tabs.

5

Click on the NAT Prefix tab to configure NAT prefixes. Up to 1024 prefixes can be configured on a prefix list.

1. Click Create or select an existing NAT prefix and click Properties. The NAT Prefix (Create|Edit) form opens.
2. Configure the Prefix Address and Prefix Length parameters.
3. Select a NAT policy if the NAT prefix list is part of an L2–aware NAT configuration.

6

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

30.8 To configure a NAT classifier

30.8.1 Purpose

Perform this procedure to configure a NAT classifier for use in a destination NAT (DNAT) configuration.

30.8.2 Steps

1

Choose Policies→ISA Policies→NAT Classifiers from the NFM-P main menu. The NAT Classifiers form opens.

2

Click Create or select an existing NAT classifier and click Properties. The NAT Classifier (Create|Edit) form opens.

3

Configure the required parameters.

If the Default Action parameter is set to DNAT, you can configure the Default Action Address parameter, the Default DNAT Address parameter, or both.

4

Click on the Classifier Entries tab to configure NAT classifier entries. You can configure up to 32 classifier entries.

1. Click Create or select an existing classifier entry and click Properties. The NAT Classifier Entry (Create|Edit) form opens.
2. Configure the required parameters.

   If the Action parameter is set to DNAT, you must configure the DNAT Address parameter, unless the Default DNAT Address parameter was configured in Step 3.
3. Save your changes and close the form.

5

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

30.9 To configure NAT on a routing instance

30.9.1 Purpose

Perform this procedure to configure NAT on a base or VPRN routing instance.

30.9.2 Steps

1

Do one of the following:

a. To configure NAT on a base routing instance, in the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

b. To configure NAT on a VPRN routing instance, choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens. On the service navigation tree, expand the Sites icon. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the NAT tab.

3

Click Create or select an existing entry and click Properties. The NAT Configuration form (Create|Edit) opens.

4

Click on the Hints button at the bottom of the form to display information about the various NAT components and options available.

5

Configure the MTU parameter.

6

Click Select to assign any of the following types of ACL IP filter policy:

- Upstream IPv4 Filter
- Downstream IPv4 Filter
- Upstream IPv6 Filter
- Downstream IPv6 Filter

See 51.5 “To configure an ACL IP filter policy” (p. 1671) for information about configuring an ACL IP filter policy.

7

Configure one or more NAT pools.

Perform the following steps:

1. Click Create or select an existing entry and click Properties. The NAT Pool form (Create|Edit) appears.
2. Configure the required parameters.

   The Deterministic NAT Port Reservation and Port Forwarding Dynamic Block Reservation parameters are mutually exclusive. If one is configured, the other is disabled. The Port Forwarding Dynamic Block Reservation parameter applies only to large scale NAT pools.

   The Application Agnostic parameter can only be configured at NAT pool creation time. When a NAT pool is configured as application-agnostic, the Subscriber Limit, Deterministic NAT Port Reservation, Port Forwarding Dynamic Block Reservation, Pool Mode, Port Reservation (Type and Value), and Port Forward Range End parameters are pre-configured and cannot be changed.

   If the NAT Pool Type parameter is set to L2–aware, you can configure the External Assignment parameter, if required.
3. If the NAT pool belongs to a PFSG, click Select in the Redundancy panel to assign a lead NAT pool for the PFSG.
4. Select an ISA-NAT group for the NAT pool. See 30.3 “To configure an ISA-NAT group” (p. 1080) for information about configuring a ISA-NAT group.
5. Click on the NAT Pool Ranges tab.
6. Click Create or select an existing entry and click Properties. The Pool Range form (Create|Edit) opens.
7. Configure the required parameters.

8. Save your changes and close the forms. The NAT Configuration form reappears.

8

Click on the Inside tab.

Perform the following steps:

1. In the Deterministic NAT panel, select a NAT policy.

   Note:

   The NAT policy you choose must be distributed to the local NE, and the local definition of the NAT policy must be configured with a NAT pool.
2. Click Select to choose a NAT policy, as required. If you select a NAT policy, the NAT 64 Enabled parameter also becomes configurable, if required.
3. Click Select to choose a Downstream IPv4 filter, as required.
4. If a RADIUS Proxy Server is in use, click Select in the RADIUS Proxy Server panel and choose the following:
   - routing instance
   - RADIUS proxy server; see 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847). The selection of the RADIUS Proxy Server Name is dependent on the selected routing instance. You can only select servers enabled for accounting.
5. In the RADIUS Vendor Information panel select a RADIUS attribute type.
6. In the DNAT Only panel, select a DNAT source prefix list.

9

Click on the L2 Aware IP Addresses tab to configure IP addresses for L2-aware NAT forwarding.

Perform the following steps:

1. Click Create or choose an existing entry and click Properties. The L2 Aware IP Address, Routing Instance (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form. The NAT Configuration form reappears.

10

Click on the DS Lite Address tab to configure DS Lite addresses.

In order to assign DS Lite Addresses, you must assigned a NAT policy in Step 8.

1. Click Create or choose an existing entry and click Properties. The Dual Stack Lite Address (Create|Edit) form opens.
2. Configure the required parameters.

   If address IP fragmentation is enabled, the Tunnel MTU parameter value must be equal to or greater than 1232 bytes.
3. Save your changes and close the form. The NAT Configuration form reappears.

11

Click on the NAT 64 tab to configure NAT 64.

In order to configure NAT 64, you must enable the NAT 64 Enabled parameter in Step 8.

1. Configure the Subscriber Prefix Length parameter.
2. If the Administrative State parameter is set to Out Of Service, you can configure the required parameters.

   Bits 64 to 71 of the NAT Destination Prefix parameter should be set to zero in order for the prefix value to be in compliance with RFC6052.

   If IP fragmentation is enabled, the IPv6 MTU parameter value must be equal to or greater than 1280 bytes.

12

Click on the Destination Prefixes tab to configure inside NAT destination addresses.

Perform the following steps:

1. Click Create or choose an existing entry and click Properties. The Destination Prefix (Create|Edit) form opens.
2. Configure the required parameters:
3. If you require a local NAT policy configuration for the destination prefix, click Select and choose a NAT policy. See 30.5 “To configure a NAT policy” (p. 1082).
4. Save your changes and close the form. The NAT Configuration form reappears.

13

Click on the Deterministic Prefixes tab to configure inside NAT deterministic prefixes.

In order to configure a NAT deterministic prefix, you must have configured one of the Deterministic NAT parameters in Step 8.

1. Click Create or choose an existing entry and click Properties. The Deterministic NAT Prefix (Create|Edit) form opens.
2. Configure the required parameters.
3. Select a NAT policy. See 30.5 “To configure a NAT policy” (p. 1082).
4. In the Deterministic Map panel, click Create or choose an existing entry and click Properties. The Deterministic NAT Map, Routing Instance (Create|Edit) form opens.
5. Configure the required parameters.
6. Save your changes and close the forms. The NAT Configuration form reappears.

14

Click on the NAT Port Forwarding tab to configure NAT port forwarding.

Perform the following steps:

1. Configure the Resync Static and Dynamic parameter, if required. If this parameter is enabled, both static and dynamic forwarders are retrieved from the NE.

   Static port forwards are persisted in the NFM-P database. Dynamic port forwards are retrieved from the NE, and are retained in the NFM-P GUI only for the duration of the current session.
2. Click on the Static tab.
3. Click Create or choose an existing entry and click Properties. The NAT Static Port Forwarding (Create|Edit) form opens.
4. Configure the required parameters.

   You cannot specify the same set of Inside Port and Protocol values in more than one static port mapping to an Inside IP Address.

   You can specify the same Outside Port value in multiple mappings to an Inside IP Address.

   The B4 Address and AFTR Address parameters apply only to DS Lite LSN configurations.
5. If you require a local NAT policy configuration for NAT port forwarding, click Select and choose a NAT policy. See 30.5 “To configure a NAT policy” (p. 1082).
6. If a redundant pool configuration is in use, click Select and choose the following for the standby pool.
   - routing instance
   - NAT policy; see 30.5 “To configure a NAT policy” (p. 1082)
7. Save your changes and close the forms. The NAT Configuration form reappears.

15

Click Resync NAT Port Forward and confirm the synchronization.

16

Save your changes and close the form.

END OF STEPS

30.10 To configure static one-to-one NAT on a 7705 SAR base routing instance or VPRN routing instance

30.10.1 Purpose

Perform this procedure to configure static one-to-one NAT on a 7705 SAR base or VPRN routing instance.

30.10.2 Steps

1

Perform one of the following:

a. To configure static one-to-one NAT on a base routing instance, in the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

b. To configure static one-to-one NAT on a VPRN routing instance, choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens. On the service navigation tree, expand the Sites icon. Right-click on a routing instance icon and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the Static NAT tab, and then on the General subtab.

3

Enable the Enable Static NAT parameter.

4

Configure the Drop Packets without NAT parameter, as required.

This parameter configures the router to drop packets that are traveling from either an inside network to an outside network or an outside network to an inside network that do not have a NAT mapping entry. By default, packets traveling from either an inside network to an outside network or an outside network to an inside network are forwarded whether or not there is a NAT mapping entry.

5

Click on the Inside Map subtab and click Create. The Static NAT Inside Map (Create) form opens.

6

Configure the required parameters.

The Inside Map maps a range of inside source IP addresses that will undergo NAT to a specified outside IP address range. The 7705 SAR will sequentially map each inside source IP address to its corresponding outside IP address.

7

Save your changes and close the form.

8

Perform one of the following:

a. Navigate to the IES L3 Access Interface Properties form, General tab, and enable the Enable

Static NAT Inside parameter.

b. Navigate to the VPRN L3 Access Interface Properties form, General tab and enable the Enable Static NAT Inside parameter.

c. Navigate to the VPRN Tunnel Interface Properties form, General tab and enable the Enable Static NAT Inside parameter.

9

Save your changes and close the form.

END OF STEPS

30.11 To configure an IPv6 firewall domain

30.11.1 Purpose

Perform this procedure to configure an IPv6 firewall domain on a NAT configuration on a base or VPRN routing instance.

30.11.2 Steps

1

Do one of the following:

a. To configure an IPv6 firewall domain on a base routing instance, in the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

b. To configure an IPv6 firewall domain on a VPRN routing instance, choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens. On the service navigation tree, expand the Sites icon. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the NAT tab.

3

Click on the Firewall Domain tab.

4

Click Create or select an existing firewall domain in the list and click Properties. The Firewall Domain (Create|Edit) form opens.

5

Configure the parameters on the General tab and select an ISA group.

6

Click on the Prefixes tab to configure firewall prefixes.

1. Click Create or select an existing firewall prefix in the list and click Properties. The Firewall Prefix (Create|Edit) form opens.
2. Configure the parameters.

7

Save your changes and close the forms.

END OF STEPS

30.12 To configure a MAP-T domain

30.12.1 Purpose

Perform this procedure to configure an MAP-T domain on a NAT configuration on a base or VPRN routing instance.

30.12.2 Steps

1

Do one of the following:

a. To configure a MAP-T domain on a base routing instance, in the navigation tree Routing view, expand Network→*NE*→Routing Instance. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

b. To configure a MAP-T domain on a VPRN routing instance, choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens. Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens. On the service navigation tree, expand the Sites icon. Right-click on a routing instance icon that has a NAT configuration and choose Properties. The Routing Instance (Edit) form opens.

2

Click on the NAT tab.

3

Click on the MAP-T Domain tab.

4

Click Create or select an existing MAP-T domain in the list and click Properties. The MAP-T Domain (Create|Edit) form opens.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

30.13 To start or stop a NAT address-pool drain operation

30.13.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on the routing instance icon that has one or more configured NAT pools and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the NAT Configuration tab.

4

Click Search, select a NAT configuration in the list and click Properties.The NAT Configuration (Edit) form opens.

5

Click on the NAT Pools tab.

6

Click Search, select the required address pool and click Properties.The NAT Pool (Edit) form opens.

Note: You can also select multiple address pools and use the Drain or Stop Drain buttons to control the draining for the selected address pools, as described in subsequent procedure steps.

7

Click Drain to start a drain operation on the address pool. The number of host sessions that are associated with the address range decreases as sessions close.

8

Click Stop Drain to stop an active drain operation on the address pool. New host sessions can use the address range.

9

Close the forms.

END OF STEPS

30.14 To configure a NAT deterministic script on a remote server

30.14.1 Steps

1

In the navigation tree Equipment view, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab.

3

Click on the Subscriber Management tab.

4

Configure the Remote URL parameter using the following syntax:

[{ftp://|tftp://}<*login*>:<*password*>@<*remote_address*>/] [<*path_file*>]

where...

<*login*> is the login name of the remote server.

<*password*> is the login password of the remote server.

<*remote_address*> is the IP address of the remote server.

<*path_filename*> is the path to the folder on the remote server where the mappings are stored.

5

Click Apply and then click Save Script to save the current deterministic NAT mappings to a file on the specified remote server.

6

If required, click Resync to update the read-only information in the Save Needed, Last Saved Result, and Last Saved Time fields.

If a change occurs on the local NE that affects the deterministic NAT mappings, the Save Needed parameter changes to Yes, indicating that you must click Save Script to update the mappings.

7

Save your changes and close the form.

END OF STEPS

30.15 To configure statistics on an ISA-NAT group

30.15.1 Steps

1

In the navigation tree Equipment view, expand Network→*NE*→Logical Groups→ISA-NAT Groups→ISA-NAT Group *n*.

2

Right-click on the ISA-NAT Group *n* icon where you want to configure statistics and choose Properties. The ISA-NAT Group (Edit) form opens.

3

Click on the ISA-NAT MDA tab to configure ISA Member statistics, ISA Member Usage statistics, or Policy statistics,

Note: You can also collect ISA NAT Member statistics or ISA NAT Member Usage statistics (but not NAT Policy statistics) from the ISA-NAT Members tab. The procedure for collecting statistics is the same as is outlined below.

4

Click Search and select an ISA-NAT MDA and click Properties. The ISA-NAT Group MDA (Edit) form opens.

5

Click on the Statistics tab.

6

In the Select Object Type area, choose one of the following from the contextual menu:

- ISA NAT Member Usage Stats
- ISA NAT Resources Stats

- NAT Policy Stats

7

Click Collect, select a record on the ISA-NAT Group MDA (Edit) form and click Properties. The Statistics Record form opens.

8

View the read-only statistics parameter values for the ISA-NAT group MDA.

9

Close the forms.

END OF STEPS

30.16 To plot LSN subscriber host statistics

30.16.1 Purpose

Perform this procedure to view LSN subscriber host statistics data in graphical form.

Note: You can plot only real-time statistics data for a subscriber host. The NFM-P does not retain historical subscriber host information.

30.16.2 Steps

1

Perform one of the following.

a. Plot the statistics data from a base routing instance.

   Perform the following steps:

   1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance.
   2. Right-click on the Routing Instance icon and choose Properties. The Routing Instance (Edit) form opens.

b. Plot the statistics data from a VPRN routing instance.

   Perform the following steps:

   1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
   2. Click Search, select a VPRN service in the list and click Properties.The VPRN Service (Edit) form opens.
   3. In the service navigation tree, expand Sites→Routing Instance.
   4. Select a routing instance in the service navigation tree and choose Properties. The Site (Edit) form opens.

2

Click on the NAT tab and perform one of the following.

a. To plot host statistics for a DS Lite subscriber, click on the NAT DS Lite Subscribers tab.

b. To plot host statistics for a NAT 64 subscriber, click on the NAT64 Subscribers tab.

c. To plot statistics for classic LSN subscribers, click on the Classic LSN Subscribers tab.

3

Click Search, select a subscriber host in the list and click Properties. The subscriber host properties form opens.

4

Choose Plotter→New Plot. The Statistics Plotter window opens.

See “Graphing statistics” in the *NSP NFM-P Statistics Management Guide* for information about using the NFM-P Statistics Plotter.

5

View the plotted statistics data as required.

Save the tabular data or plot graph, as required.

6

Close the forms.

END OF STEPS

30.17 To view reserved IP address and reserved block information on an ISA-NAT group

30.17.1 Steps

1

In the navigation tree Equipment view, expand Network→*NE*→Logical Groups→ISA-NAT Groups→ISA-NAT Group *n*.

2

Right-click on an ISA-NAT Group *n* icon and choose Properties. The ISA-NAT Group (Edit) form opens.

3

Click on the ISA-NAT Members tab and select an ISA-NAT member and click Properties. The ISA-NAT Member form opens.

4

In the Metrics panel, review the information provided in the Reserved IP address and Reserved block read-only parameters.

Note: Click Resync to update the IP Address Reserved and Blocks Reserved values to the most recent information available on the MDA.

5

Close the forms.

END OF STEPS

30.18 To view ISA-NAT object information

30.18.1 Purpose

Use this procedure to perform an on-demand retrieval of large-scale NAT entries, based on a routing instance ID. Search results can be further constrained by specifying an inside address (for a classic LSN host, DS Lite subscriber, or NAT64 subscriber) or an outside address and start port number (for a LSN block allocation).

30.18.2 Steps

1

Click Manage→ISA Functions→ISA-NAT from the NFM-P main menu. The Manage ISA-NAT form appears.

2

Select the appropriate ISA-NAT object type, for example, Classic LSN Subscriber, ISA-NAT Group, or LSN block allocation. The Manage ISA-NAT form refreshes displaying the search parameters used to locate the ISA-NAT object.

3

Configure the parameters as required to locate the ISA-NAT object.

4

Click Search, select an entry in the list and click Properties. A read-only object properties form appears.

5

Review the information as required.

6

Close the forms.

END OF STEPS

31 MPLS

MPLS overview

31.1 MPLS overview

31.1.1 Overview

The NFM-P supports the configuration and provisioning of MPLS paths and LSPs.

MPLS is a data-carrying mechanism that emulates some properties of a circuit-switched network over a packet-switched network. MPLS prepends a packet with an MPLS header that is comprised of labels in a stack. As an NE forwards a packet through an MPLS network, it examines the top label for routing instructions and pops it off the stack; the contents below the label stack are not examined. The label stack is not removed until it reaches the egress NE.

MPLS uses one or more IGPs, such as IS-IS, OSPF, or RIP, to forward packets. Regardless of the IGP chosen, the MPLS configuration is the same. An advantage of MPLS is that if more than one IGP is enabled, it continues to work when another protocol fails. See Chapter 28, “Routing protocol configuration” for more information on IS-IS, OSPF, and RIP.

MPLS can be used as the underlying transport mechanism for service tunnels. For MPLS to be used as such, an MPLS mesh and an LSP mesh must be created before the tunnel is created. Since LSPs and service tunnels are unidirectional, they must be created in both directions. See Chapter 33, “Service tunnels” for more information about service tunnels.

When MPLS is enabled on an NFM-P-managed NE routing instance, the NFM-P displays an MPLS instance below the routing instance in the network view of the navigation tree. Using this view, you can perform various functions on the MPLS instance, for example, assign an L3 interface to an MPLS instance.

If a link, NE, or path fails, MPLS can determine a redundant path by using fast reroute. Fast reroute uses an alternative NE to complete the failed LSP. Fast reroute provides the following types of route protection.

- Many-to-one—one backup route is maintained for multiple protected LSPs on an NE.
- One-to-one—a separate backup route is maintained for each protected LSP on an NE algorithm.

You can list and view MPLS objects such as MPLS and RSVP instances and interfaces, static and dynamic LSPs, cross connections, and MPLS paths using the NFM-P Manage MPLS Objects form.

31.1.2 MPLS forwarding policies

The data model of a forwarding policy represents each pair of next hops (primary and backup) as a group and models the ECMP set as the set of next-hop groups. Flows of prefixes can be switched on a per next-hop group basis from the primary next-hop, when it fails, to the backup next-hop without disturbing the flows forwarded over the other next-hop groups of the policy. The same can be performed when reverting from a backup next-hop to the restored primary next-hop of the same next-hop group.

An endpoint MPLS forwarding policy allows the user to forward unlabeled packets over a set of user-defined direct (with option to push a label stack), indirect, or LSP next-hops. Routes are bound to an endpoint policy when their next-hop matches the endpoint address of the policy.

A label-binding MPLS forwarding policy provides the same capability for labeled packets. In this case, labeled packets matching the ILM of the policy binding label are forwarded over the set of next-hops of the policy.

31.2 LSPs

31.2.1 Overview

An LSP is a path through an MPLS network that is set up based on criteria in a forwarding equivalency class (FEC). An FEC is a set of characteristics that define how NEs in an MPLS network forward packets that are bound to an MPLS label. An FEC includes specifications such as the destination IP address and QoS parameters. An FEC is associated with a specific LSP, but an LSP may be used for multiple FECs.

An LSP begins at a PE device called an LER, which is the ingress NE that prepends a label to a packet based on an FEC and sends the packets to the next transit NE in the path. The transit NE swaps the packet label for a new one, and sends the packet to the next NE. The LER that acts as the egress PE NE removes the label and forwards the packet according to the header of the next layer, for example, IP.

LSPs are unidirectional; they enable the label switching of a packet through an MPLS network from one endpoint to another. Bidirectional communication through an MPLS network requires the configuration of an LSP in the opposite direction.

The NFM-P supports the following types of LSPs:

- Static LSPs specify a static path through an MPLS network. All transit NEs require manual configuration with LSP labels and require no signaling protocol, such as LDP or RSVP.
- Bypass-only LSPs are LSPs with manually configured bypass tunnels on Point-of-Local-Repair (PLR) NEs. Such LSPs are used exclusively for the purpose of bypass protection.
- Dynamic, or signaled LSPs, use a protocol such as LDP or RSVP. The signaling protocol allows an ingress NE to dynamically assign labels to an egress NE. You must configure the ingress NE, but not the transit NEs in an LSP. LDP LSPs are not explicitly defined. The downstream unsolicited (DU) method configures them automatically through the network.
- Segment Routing TE LSPs (SR-TE LSPs) are established with traffic engineering and protection requirements based on different parameters, such as hop limit, IGP shortcut, BGP shortcut and maximum segment routing labels.
- Point-to-Multipoint LSPs allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as PIM, to be configured in the network.

You can apply a metric value to a Dynamic LSP or a Point-to-Multipoint LSP that determines which LSP the NFM-P uses when multiple LSPs lead to the same destination. The metric value for a static route is set to 1 and cannot be configured.

You can list and view detailed information on LSPs, templates, cross connections, service tunnels, and detour and bypass information on the LSP (Edit) forms.

31.2.2 Dynamic LSPs

A Dynamic LSP uses a signaling protocol such as LDP or RSVP-TE and an MPLS path between two NEs. You can then create an LSP between the two NEs and bind it to an MPLS path. An LSP-MPLS path binding is called an LSP path. You can configure the LSP path after the MPLS path and the LSP are bound.

Dynamic LSPs are categorized as follows:

- Explicit-path LSPs resemble static LSPs when all hops are strict; each hop in the LSP requires explicit configuration. MPLS uses RSVP-TE to set up an explicit-path LSP. Transit NE hops can be strict, and use a direct path between NEs, or loose, and include other NEs in the path between NEs.
- Constrained-path LSPs have dynamically assigned intermediate NE hops that rely on CSPF to find a path that satisfies the LSP constraints. CSPF is a routing algorithm that takes different LSP constraints, such as the available bandwidth and MPLS administrative groups, into account to balance the network load. When the CSPF path is found, RSVP uses the path to request the LSP setup. If Fast Reroute is enabled, the ingress NE signals the downstream NEs to set up a detour configuration for the LSP.

When an LSP is established, the reserved bandwidth is controlled by the bandwidth parameter at the primary path level, regardless of whether the LSP has auto-bandwidth enabled. When auto-bandwidth is enabled and a trigger occurs, the NE attempts to change the bandwidth of the LSP to a value between the minimum and maximum bandwidth, which are configurable at the LSP level. Automatic bandwidth allocation is supported on RSVP LSPs that have both CSPF and MBB enabled. If an RSVP LSP is configured for auto-bandwidth, the ingress LER determines, at every adjust interval, whether to attempt an auto-bandwidth adjustment. You can change the minimum bandwidth, maximum bandwidth, or threshold parameters on an operational LSP, however, the changes do not take effect until the next auto-bandwidth trigger, for example, an adjust interval expiry. If the bandwidth adjustment fails, for example, the CSPF cannot find a path, the existing LSP is maintained with its existing bandwidth reservation. See 31.11 “To create a Dynamic LSP” (p. 1126) for more information.

The NFM-P provides OAM tools for troubleshooting service and network transport issues. You can run an OAM validation test for the Dynamic LSP by clicking on the One Time Validation button. If a check mark appears beside the OAM Validation Failed state cause indicator, the test has failed. The One Time Validation Result tab on the Tests tab displays detailed information about the OAM test result. See 90.55 “To run a one-time validation test on a service” (p. 3084) for information on how to use the One Time Validation function.

You can manually switch away or switch back to the primary path for a RSVP-TE using the Manual Switch Path button on the Tunnels tab of the Dynamic LSP form. Selecting the Switch Away option will degrade the primary path, switching to a standby path, and selecting the Switch Back option will restore the primary path.

31.2.3 Static LSPs

A static LSP uses an IGP instead of a signaling protocol, such as LDP or RSVP-TE. The NFM-P attempts to derive the hop configurations based on the hop labels in the path. You must use an ingress label that is unique within an NE when you create a new static hop; otherwise, the NFM-P rejects the new hop.

You can assign static label mappings for LSP cross-connections on an MPLS interface during interface creation or modification. A static label map is used only for intervening unmanaged NEs. The NFM-P attempts to derive the interface and label values from the swap egress label of the previous hop. The Static Label Maps form lists the static label maps for an interface whether they are created using the NFM-P client GUI, an OSS client, or CLI.

The NFM-P does not raise an alarm against an unmanaged NE.

The NFM-P raises an alarm against a static LSP under the following conditions when the LSP destination is managed by the NFM-P:

- No hops are configured for the static LSP.
- The last hop does not match the LSP destination.
- The label action specified for the last hop is not a pop action.

31.2.4 SR-TE LSPs

You can configure a Segment Routing Traffic Engineered (SR-TE) LSP using the Manage→MPLS→Segment Routing LSPs option. You can associate an empty path or a path with strict or loose explicit hops with the primary path of the SR-TE LSP.

The SR-TE LSP has the following characteristics:

- Static route tunnel binding configuration allows the forwarding of packets of a static route to an indirect next-hop over an SR-TE LSP.
- BGP shortcut tunnel binding configuration allows the forwarding of packets of BGP prefixes over an SR-TE LSP.
- The BGP transport tunnel can be enabled so that the SR-TE LSP can resolve the next-hop of a BGP IPv4 or IPv6 (6PE) label route.
- An SDP sub-type of the MPLS encapsulation type allows service binding to a SR-TE LSP.
- The support of SR-TE LSP in the data path requires that the ingress LER pushes a label stack where each label represents a hop, a TE link, or an NE, in the ERO for the LSP path computed by the PCE.

An SR-TE LSP is configured on the NE, but its path can be computed by the NE or by an external Traffic Engineering controller referred to as a Path Computation Element (PCE). This feature works with the Nokia stateful PCE, which is part of the NSP. Three modes of operations can be configured for an SR-TE LSP:

- When the path of the SR-TE LSP is computed by the NE acting as a PCC, the LSP is referred to PCC-initiated and PCE-controlled.
- When the path of the SR-TE LSP is computed by the PCE at the request of the PCC, it is referred to as PCC-initiated and PCE-computed.
- When the path of the SR-TE LSP is computed and updated by the PCE following a delegation from the PCC, it is referred to as PCC-initiated and PCE-controlled.

You can manually switch away or switch back to the primary path for a SR-TE using the Manual Switch Path button on the Tunnels tab of the Segment Routing LSP form. Selecting the Switch Away option will degrade the primary path, switching to a standby path, and selecting the Switch Back option will restore the primary path.

31.2.5 Point-to-Multipoint LSPs

A Point-to-Multipoint (P2MP) MPLS LSP allows the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as PIM, to be configured in the network. A P2MP LSP tree is established in the control plane for which the path consists of a head-end NE, one or many branch NEs, and multiple leaf NEs. Packets that are injected by the head-end NE are replicated in the data plane at the branching NEs before they are delivered to the leaf NEs.

The P2MP LSP has a source IP address but does not a have a destination address. Instead, you must create S2L leaf NEs for each destination in the P2MP LSP tree. The P2MP LSPs are not supported in SDP bindings.

Each P2MP LSP object has one P2MP LSP instance object. This is the primary instance, and can contain multiple child S2L path objects. Just as the P2MP LSP can appear as a tree, each S2L path object represents a root-to-leaf (S2L) sub-LSP path for the primary instance. The S2L paths can be empty paths or can specify a list of explicit hops. The same path can be used by more than one S2L of the instance. However the destination IP address must have a unique argument per S2L, because it corresponds to the address of the egress LER NE.

PIM tunnel interfaces are associated with a P2MP LSP. The tunnel interfaces are needed at both the ingress LER and the egress LER NEs. You must also associate static multicast groups with the tunnel interfaces that are associated with the P2MP LSPs.

On the ingress side, you can create one or more tunnel interfaces in PIM and associate each with a different RSVP P2MP LSP. The tunnel interface is associated with the P2MP LSP name. You can then assign static multicast group joins to each tunnel interface using an IGMP configuration. A specific <*,G> or <S,G> can only be associated with a single tunnel interface. A multicast packet which is received on an interface and matches the <*,G> or <S,G> specified in the IGMP join for the tunnel interface is replicated and forwarded to all branches of the P2MP LSP.

On the egress side, the PIM tunnel interface is associated with both the P2MP LSP name and the system address of the ingress LER. You must define a multicast info policy to associate specific multicast groups, or a specific <S,G>, to the primary tunnel interface on each egress LER leaf NE. The multicast info policy must then be applied to the router. A multicast packet synced from a tunnel interface associated with a P2MP LSP on the egress leaf NE can then be forwarded over a PIM or IGMP interface, which can be an access interface or a network interface, including a spoke SDP-based IES interface. You may create multiple tunnel interfaces per NE, where each interface is associated with a different P2MP LSP. The P2MP LSP association to a multicast group can be applied at the bundle, channel, and/or source channel level in the policy.

You can create a P2MP, one-hop P2P, or mesh P2P LSP using the LSP template for MVPN from the Policies→MPLS→LSP Template MVPN menu option. You can configure the scope as either local or global. The global template can be modified and multiple local versions can exist on the NEs. You must assign a default MPLS path to the template on the local definition before the template is turned up. If you modify the template, you must shut down the local policy first. When you turn up the local policy, all previously created P2MP LSP parameters are synchronized with the new template. See 31.29 “To create an LSP template MVPN policy” (p. 1156) for more information.

31.2.6 Bypass LSPs and manual bypass tunnels

You can create manual bypass tunnels and dynamic (or automatic) bypass tunnels on the SR/ESS, 7705 SAR, and 7210 SAS. You can manually configure an LSP to be bypass-only on a PLR NE. The LSP is then used exclusively for the purpose of bypass protection.

You can use manual bypass alone or with dynamic bypass. When used with the dynamic bypass, the manual bypass has precedence over the dynamic bypass for the path selection.

Dynamic bypass tunnels can be disabled on a per-NE basis. They are enabled by default. If dynamic bypass is disabled on a network element while dynamic bypass tunnels are active, traffic loss occurs. Furthermore, if no suitable manual bypass LSPs are found, the protected LSP remains without protection.

A bypass-only LSP does not have all the configurable attributes of a regular dynamic LSP. The main differences include:

- Manual bypass LSPs only support primary path. Secondary or standby paths cannot be created.
- Manual bypass LSPs do not support bandwidth or fast-reroute.
- Path monitoring is not available for manual bypass LSPs.
- Manual bypass LSPs cannot be assigned to a service tunnel, or be added to an LDP targeted peer as the LDP tunneling LSP.
- There is no rule-based topology support available for manual bypass LSPs.
- There is no network-wide LSP view (the physical map LSP view) support available for manual bypass LSPs,
- You cannot create a static LSP using a tunnel template.

The following additional characteristics also apply to a manual bypass:

- The default protection level for a bypass tunnel is “node-protect”.
- CSPF is supported for the manual bypass tunnel. When CSPF is enabled, loose path and hop-less path are allowed on a manual bypass tunnel. The CSPF calculation is performed when the manual bypass tunnel is first set up.

31.2.7 Administrative groups with facility bypass backup LSPs

Administrative group constraints configured on an LSP primary path are used for signaling of bypass LSPs. The NFM-P supports P2P LSP and P2MP S2L path bypass LSPs. Administrative groups with one-to-one detour LSPs are not supported.

You can enable the use of administrative group constraints in the association of a manual or dynamic bypass LSP with the primary LSP path on the MPLS instance of the NE. When constraints are enabled, each PLR NE reads the constraints in the FAST_REROUTE object in the path message of the primary path LSP. If the FAST_REROUTE object is not included in the path message, the PLR NE reads the constraints from the session attribute object in the path message. The PLR NE uses the administrative group constraints and other constraints, such as the hop limit and SRLG, to choose a manual or dynamic bypass LSP from the LSPs that are in use.

You can enable the use of the administrative group constraints on FRR backup LSP on the P2P and the P2MP LSP configuration forms, and in the LSP template MVPN policy. The constraints apply only to new attempts to find a valid bypass.

31.2.8 Automatic ABR selection and dynamic ABR bypass protection

The ABR selection for the inter-area RSVP P2P LSP path computation by CSPF is automatic at the ingress LER. You do not need to include the ABR as a loose-hop in the LSP path definition. On the LSP path binding configuration form, you can view whether:

- the operational path is inter-area or intra-area
- the current CSPF hop is an ABR (on the CSPF Path tab)

31.2.9 Dynamic bypass LSP for ABR NE protection

The NFM-P provides dynamic bypass computation, signaling, and association with the primary path of an inter-area P2P LSP to provide ABR NE protection. For a PLR NE within the local area of the ingress LER to provide ABR NE protection, the NE must dynamically signal a bypass LSP and associate it with the primary path of the inter-area LSP. You can view the address of the NE that must be bypassed along the path in the RSVP session tab on the Dynamic LSP properties form.

31.2.10 FRR protection

Fast reroute protection is a MPLS and IP resiliency technology that provides fast traffic recovery upon link or router failures for mission critical services. FRR allows a user to provide local protection for an LDP FEC by pre-computing and downloading to the IOM both primary and backup next hops for the FEC.

FRR provides for the use of the Loop-free Alternate (LFA) backup next-hop for forwarding in-transit and CPM generated IP packets when the primary next-hop is not available. IP FRR is supported on IPv4 and IPv6 prefixes forwarded in the base router instance over a network IP interface, or an IES SAP or spoke interface. It is also supported for VPRN VPN-IPv4 prefixes and VPN-IPv6 prefixes forwarded to a VPRN SAP or spoke interface.

FRR requires SPF computation of an LFA next-hop (in addition to the primary next-hop) for all prefixes used by LDP to resolve FECs. The LFA next-hop is populated into the routing table along with the primary next-hop for each prefix.

LFA is configured on IS-IS routing instances, enabling LFA computation by SPF under the IS-IS routing protocol. The LFA next hop can be excluded individually on the IS-IS Level 1 and IS-IS Level 2 objects, as well as on IS-IS interface objects.

LFA is configured on OSPF routing instances, enabling LFA computation by SPF under the OSPF routing protocol. The LFA next hop can be excluded individually on the OSPF area objects, as well as on OSPF interface objects.

31.2.11 Tunnel Templates

Tunnel templates allow users to configure Dynamic LSP, LSP path, and SDP templates to define common characteristics for a tunnel or templatable tunnel object.

To simplify creating tunnel templates, the NFM-P provides examples of common tunnel templates that can be copied and customized. You can also create a tunnel template from an existing templatable object.

After a template is configured, Dynamic LSPs, SR-TE LSPs, LSP paths and SDPs can be configured by choosing a create from template button during configuration. Users can configure

multiple LSP paths for a Dynamic or SR-TE LSP. For example, you can configure one primary, one standby, and many secondary LSP paths.

When the NFM-P auto-tunnel creation is used to create LSP paths, the LSP paths must use a hopless path. If the paths belong to the same LSP, they must use different MPLS paths, even when the MPLS paths are hopless. For a RSVP LSP policy, one LSP template can be specified. When an LSP template with multiple LSP paths is specified, many LSP paths are created.

31.2.12 LSP Path Optimization

The NFM-P allows you to resignal LSP Paths to take advantage of new paths that are less congested, fewer hops, have a lower metric and meet least-fill criteria. NEs periodically check the network to determine whether a more efficient path is available and notify the NFM-P when another path is eligible for re-signaling. The NFM-P maintains a list of LSP Paths that are eligible for re-signaling. When an LSP Path is eligible for optimization, the LSP Path can be routed to a different path. When a LSP Path is not eligible for optimization, the LSP Path cannot use newly available network resources. See 31.24 “To configure an LSP Path optimization policy” (p. 1148) for information about configuring an LSP Path optimization policy.

31.2.13 LDP-over-RSVP tunnels

In networks that contain dozens of NEs spanning multiple routing areas, many RSVP service tunnels may be required. To reduce the number of RSVP tunnels required for service deployment in large networks, you can use the NFM-P to create LSPs using LDP-over-RSVP, sometimes called tunnel-in-tunnel encapsulation. An SDP can ride on multiple LDP-over-RSVP tunnels.

The NFM-P supports automatic, rule-based service-tunnel creation using NEs that are grouped according to their role in the NFM-P-managed network. This functionality greatly reduces the time and effort required to provision a mesh of service tunnels. See Chapter 33, “Service tunnels” for more information about rule-based automatic service-tunnel creation.

The NFM-P can use tunnel rules and groups to create new RSVP LSP bindings between a PE device that is added to the network and the existing ABRs. It can also create a new LDP-over-RSVP tunnel between a new NE and an existing PE NE. See Chapter 33, “Service tunnels” for more information about tunnel creation using rules and groups.

Service traffic that is transported by an LDP-over-RSVP tunnel requires a VC label, an LDP label, and an RSVP label. Unlike T-LDP sessions that use IGP SPF algorithms, RSVP LSPs are not advertised to an IGP instance. When the same FEC applies to the destination of an LDP-over-RSVP LSP tunnel and an IGP-based LDP tunnel, the NFM-P uses the LDP tunnel by default to minimize network overhead.

Using tunnel-in-tunnel encapsulation, a pair of LSP tunnels and a T-LDP session between two NEs are equivalent to two adjacent LDP NEs with a non-tunneled LDP session between them. In other words, the LDP tunnel uses the RSVP LSP as one hop between LSRs in the network.

During the configuration of LDP-over-RSVP, you can specify an explicit list of dynamic and static LSPs, or use the NFM-P to find eligible LSPs. After an LSP is explicitly configured for LDP tunneling, the NFM-P associates the LDP targeted peer with the LSP.

You can configure LDP-over-RSVP to enable an OSPF area router to be a stitching point. You can specify which NEs to use as the stitching points in each area. When there are many NEs in an area, this function helps to reduce the number of LSPs required, because a full mesh is not required.

For an LSP to be eligible for LDP-over-RSVP, the following conditions apply:

- The LSP must be an RSVP-owned LSP (strict or loose)
- The LSP must start and terminate either on:
  - The router ID (system IP address)
  - The router ID (loopback address when used in multi-instance OSPF
- The OSPF system/loopback address must be advertised
- A T-LDP session must exist between the originating and terminating routers using the addresses cited above
- LDP-over-RSVP availability must be enabled in the LSP configuration. This indicates that the LSP is eligible for LDP-over-RSVP, and RSVP signals to the IGP that the LSP should be included in the SPF run.

31.2.14 Shared Risk Link Groups

Shared Risk Link Groups (SRLGs) are constructs which allow you to perform two operations that enhance overall system reliability. Firstly, you can establish a FRR LSP path. In addition, you can also use SRLGs to establish a secondary LSP path which is disjointed from the primary LSP path.

Configured SRLGs are associated with MPLS interfaces. The SRLGs are used by the CSPF when computing a FRR detour/bypass path, or a secondary LSP path. Links which are members of the same SRLG represent resources which are assumed to share the same risk. These links are therefore avoided when computing and setting up an alternate LSP path.

- FRR backup

  The SRLG constraint can be enabled system-wide on a PLR NE, in the computation of a FRR detour or bypass to be associated with a primary LSP path. CSPF includes the SRLG constraint in the computation of a FRR detour or bypass. CSPF then prunes all links with interfaces which belong to the same SRLGs as the egress interface being protected (or the immediate downstream NE, depending on the protection level). If a path is found, the bypass or detour is set up. If not, and you included the Enable SRLG for FRR - Strict option, then the bypass or detour is not set up. If the Enable SRLG for FRR - Strict option is not specified, and a path exists that meets other TE constraints (other than the SRLG constraint), then the bypass or detour is still set up.

- Secondary LSP backup

  The SRLG constraint can also be enabled per LSP path on a head-end LER in the computation of an LSP secondary path that includes the standby path. The SRLG constraint is additional to the admin group constraint on the same secondary LSP path. CSPF includes the SRLG constraint in the computation of the secondary LSP path, which requires that the primary LSP path is configured and active. CSPF prunes links with interfaces that belong to the same SRLGs as the interfaces included in the primary path. If a path is found, the secondary LSP path is set up. Otherwise, CSPF keeps trying to set up the secondary LSP path.

You can enable only the SRLG constraint for the FRR backup operation. However, you can enable both the SRLG constraint and the admin-group include/exclude constraint for the secondary LSP backup path operation. In either case, you can still apply the admin-group constraint for the primary path.

The following conditions also apply to SRLGs:

- An SRLG is modeled as a policy object. It therefore follows the normal policy behavior for creation, listing, updating, deletion, distribution and resynchronization
- SRLGs are defined node-wide
- Configured SRLGs are associated with MPLS interfaces. A specific MPLS interface can belong to multiple SRLGs. Up to 64 SRLGs can be associated with a specific MLPLS interface

31.2.15 Bandwidth-based equal cost RSVP LSP path selection

When multiple equal-cost paths satisfy the constraints of a specific RSVP LSP path, CSPF in the 7x50 head-end NE uses a random number generator to choose a path and return it to MPLS. While this method actually balances the number of LSP paths over the links in the network, it does not necessarily balance the bandwidth utilization across those links.

In order to achieve load balancing of the bandwidth amongst the available LSP paths, CSPF must include the link utilization as a criterion in the path selection. One algorithm that considers this is referred to as the “least-fill” path selection. This algorithm identifies the single link in each of the equal-cost paths that has the least available bandwidth in proportion to its maximum reservable bandwidth. CSPF then selects the path containing the link with the largest such available bandwidth to maximum reservable bandwidth percentage. The net effect of using this algorithm is that over time, LSP paths become spread over the network links in such a way that the percentage of link utilization is balanced.

When comparing the percentages of least available link bandwidth across the sorted paths, if two percentages differ by less than a value you configure as a minimum threshold, CSPF considers them equal. It then applies a random number generator to choose amongst these paths. You can also specify a reoptimization threshold, which allows you include a path cost consideration into the decision of when to alter the paths used by the LSP.

31.2.16 LSP on-demand resynchronization

LSP on-demand resynchronization is supported on the 7750 SR and 7950 XRS. On-demand resynchronization does not affect the existing manual full NE resynchronization functionality. However, when LSP on-demand resynchronization is enabled, any scheduled resynchronization is blocked for the following LSP objects:

- RSVP session
- cross-connect
- in-segment
- out-segment
- actual hop
- CSPF hop

Therefore, these LSP objects are not resynchronized as a result of a relevant trap. The exception is that for CPAM path monitored LSPs, the actual hop trap (a generic trap) is still processed.

By default, the LSP on-demand resynchronization functionality is disabled. See the procedure to enable LSP on-demand resynchronization in the *NSP System Administrator Guide* for more information. After you enable LSP on-demand resynchronization, click on the Resync button on the LSP configure form to start the manual LSP on-demand resynchronization. The following objects are resynchronized, depending on the type of LSP:

- for a Dynamic LSP: the LSP, LSP path, CSPF and actual paths, bypass or detour path, cross-connects, in-segments, out-segments, and RSVP sessions
- for a Manual Bypass LSP: the LSP, LSP path, CSPF and actual paths, cross-connects, in-segments, out-segments, and RSVP sessions
- for a Static LSP: the head-end LSP, the static hops, static label maps, cross-connects, in-segments, and out-segments. You must manually resynchronize all of the planned hop sites in a Static LSP before creating a Static LSP.

31.2.17 Administrative tagging

An administrative tag is an identifier you can apply to a dynamic LSP, a segment routing LSP, or an LSP template. You can use an administrative tag policy to include or exclude tags from automatic tunnel binding, giving you explicit control over which LSPs are used by a given BGP site or VPRN service.

Workflow for configuring administrative LSP tagging

The following workflow describes the steps to tag an LSP and enable administrative tagging filtering on a BGP site or VPRN service.

1. Create administrative LSP tags. See 31.33 “To create an administrative LSP tag” (p. 1162).
2. Create administrative LSP tag policies. See 31.34 “To create an administrative tag policy” (p. 1162).
3. Distribute the created tags and tag policies to NEs where you need to perform administrative tagging filtering. See 49.6 “To release and distribute a policy” (p. 1476).
4. Associate administrative LSP tags with dynamic LSPs, segment routing LSPs, or LSP templates using the Admin Tags tab on the configuration form for each object. See 31.11 “To create a Dynamic LSP” (p. 1126), 31.13 “To configure a Dynamic or segment routing TE LSP” (p. 1132), 31.14 “To create a Dynamic or segment routing LSP from a tunnel template” (p. 1134), and 31.12 “To create a segment routing TE LSP ” (p. 1130) for general information about configuring these LSPs.
5. Enable the Enforce Strict Tunnel Tagging parameter on BGP sites and VPRN services where you need to use administrative LSP tagging.

31.2.18 Forwarding policies

An MPLS forwarding policy allows you to direct traffic flows to certain endpoints or endpoint groups. There are two types of MPLS forwarding policy, endpoint policies and label-binding policies:

- An **endpoint policy** forwards unlabeled packets over a set of direct, indirect, or LSP next-hops. Routes are bound to an endpoint policy when their next-hop matches the endpoint address of the policy.
- A **label-binding policy** forwards labeled packets matching the ILM of the policy label over the next-hops defined in the policy.

The NFM-P supports configuring endpoint policies on MPLS instances. See 31.35 “To create a reserved label block” (p. 1163) for information about creating reserved label blocks for use with forwarding policies. See 31.6 “To configure an MPLS instance” (p. 1116) for information about configuring MPLS, including configuring forwarding policies.

MPLS workflow and procedures

31.3 Workflow to configure MPLS

31.3.1 Stages

1

Enable and configure one or more IGPs such as IS-IS, RIP, and OSPF to include the system interface on all NEs that are to participate in the MPLS network. See the appropriate device documentation for more information.

2

Enable MPLS on the routing instances of all NEs that are to participate in the MPLS network. See 31.5 “To enable MPLS on a routing instance” (p. 1115) .

3

Create or configure an MPLS instance. Assign Layer 3 interfaces, including the system management interface, to the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) or 31.7 “To create an MPLS interface” (p. 1120) .

4

Create a mesh of MPLS paths. See 31.8 “To create an MPLS path” (p. 1122) .

5

As required, create a mesh of LSPs.

a. Create a static LSP. See 31.10 “To create a static LSP” (p. 1124) .

b. Create a Dynamic LSP. See 31.11 “To create a Dynamic LSP” (p. 1126) . See 31.14 “To create a Dynamic or segment routing LSP from a tunnel template” (p. 1134) for more information about how to create a Dynamic LSP from a tunnel template.

c. Create a Segment Routing TE LSP. See 31.12 “To create a segment routing TE LSP ” (p. 1130).

d. Create a Point-to-Multipoint LSP. See 31.18 “To create a Point-to-Multipoint LSP” (p. 1137) .

e. View MVPN Point-to-Multipoint LSP objects as required. See 31.19 “To view an MVPN Point-to-Multipoint LSP object” (p. 1140) .

f. Create or configure a Manual Bypass LSP. See 31.20 “To create a Manual Bypass LSP” (p. 1141) or 31.21 “To configure a Manual Bypass LSP” (p. 1143) .

6

If your mesh includes Dynamic LSPs or Bypass-only LSPs, configure an LSP path. See 31.22 “To configure an LSP path” (p. 1144) or 31.23 “To create an LSP path using a tunnel template” (p. 1147) .

7

Create an LSP Path optimization policy, if required. See 31.24 “To configure an LSP Path optimization policy” (p. 1148) .

8

Create and apply administrative LSP tags, if required. See 31.2.17 “Administrative tagging” (p. 1111).

Note: Only Static LSPs and Static FRR LSPs are supported on the OS 9700E and OS 9800E NEs.
RSVP is not supported on the OS 9700E and OS 9800E NEs.
Label actions cannot be modified on NEs after the labels are created by the NFM-P.

31.4 Sample MPLS configuration

31.4.1 Overview

The following figure shows an example of an MPLS service configuration. The actual configuration depends on the specific network requirements.

Figure 31-1 Sample MPLS configuration

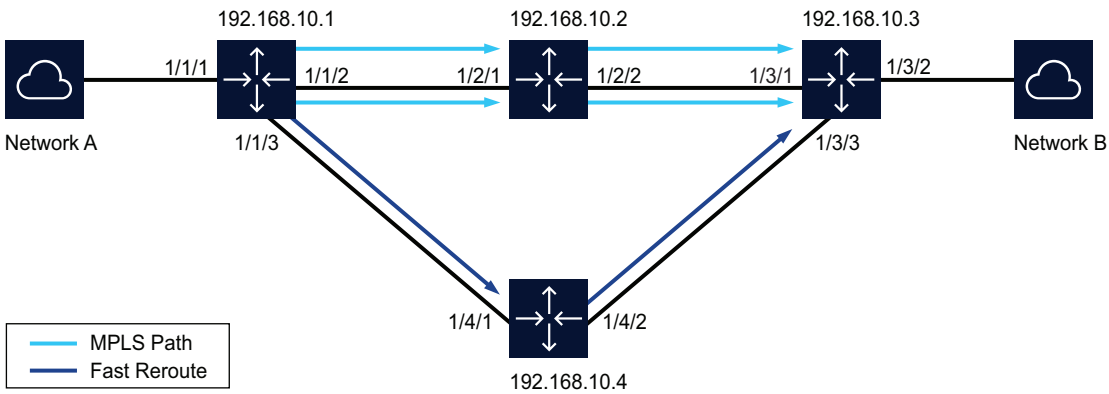

31.4.2 Stages

This process describes the high-level tasks that are required to configure the MPLS service example shown in the figure.

Preconfiguration

1

Before you begin, verify that the following preconfigurations have been made.

- An IGP is enabled on all participating NEs and includes the system interface in the IS-IS or OSPF area.
- Any additional Layer 3 interfaces that you wish to use are configured.

Configuration

2

Create an MPLS path.

Use the following steps:

1. Specify 192.168.10.1 as the source network element.
2. Specify 192.168.10.2 and 192.168.10.4 as hops.
3. Specify interface 1/3/2 on 192.168.10.3 as the destination site. This is the final hop in the MPLS path.

3

Create a Dynamic LSP.

Use the following steps:

1. Specify interface 1/1/1 on 192.168.10.1 as the source IP address.
2. Specify 192.168.10.3 as the destination IP address.

4

Bind the Dynamic LSP created in Stage 3 to the MPLS path created Stage 2 .

31.5 To enable MPLS on a routing instance

31.5.1 Steps

1

Choose Routing from the NFM-P navigation tree view selector.

2

Navigate to Routing→NE→Routing Instance.

3

Right-click on the routing instance and choose Properties. The Routing Instance (Edit) form opens.

4

Click on the Protocols tab.

5

Enable the MPLS Enabled parameter. RSVP is enabled by default and cannot be disabled using the NFM-P.

6

Save the changes and close the forms.

END OF STEPS

31.6 To configure an MPLS instance

31.6.1 Steps

1

Enable MPLS on an NE, as described in 31.5 “To enable MPLS on a routing instance” (p. 1115)
.

2

Choose Routing from the NFM-P navigation tree view selector.

3

Navigate to Routing→NE→Routing Instance→MPLS.

4

Right-click on the MPLS object in the navigation tree and choose Properties. The MPLS (Edit) form opens.

5

Configure the Administrative State parameter in the States panel.

6

Configure the required parameters in the Configuration panel.

The Static Label Range parameter affects the range of the Ingress Label parameter on MPLS static hops, MPLS static label maps, and VPLS, IES, and VLL spoke and mesh bindings.

The minimum value you can enter for the Start Segment Routing Label parameter is the value of the Static Label Range parameter, plus 32.

The PCE Report parameter specifies the PCE report types for the MPLS instance and the router-level configuration of reporting of LSPs to the PCE.

Enable the Logger Event Bundling parameter to generate the vRtrMplsXCBundleChange trap.

7

Configure the required parameters in the LSP Resignal panel.

8

Configure the required parameters in the P2MP LSP Resignal panel.

9

Configure the required parameters in the LSP Bypass Resignal panel.

10

Configure the required parameters in the SR-TE LSP Resignal panel.

11

Configure the Hold Timer (seconds) parameter in the Hold Timer panel.

12

Configure the required parameters in the Shared Risk Link Group panel.

13

Configure the required parameters in the Least-Fill panel.

14

Configure the required parameters in the TTL Propagate panel.

15

Configure the required parameters in the Auto Bandwidth Multipliers panel.

16

Configure the CSPF On Loose Hop parameter in the Inter Area/Level Traffic Engineering panel.

17

Configure LSP Self Ping as needed in the LSP Self Ping panel.

18

Configure the required parameters in the RSVP-TE LSP Resignal panel.

19

Configure the RSVP-TE and SR-TE tunnel types values in the Tunnel Table Preference panel.

20

Configure the reserved label block name as needed in the LSP BSID Block panel.

21

Configure LSP History as needed in the LSP History panel.

22

To configure OAM diagnostics for the MPLS instance, click on the Tests tab. See Chapter 90, “OAM diagnostic tests” for more information about how to configure OAM diagnostics.

23

To configure MPLS interfaces on the MPLS instance, click on the Interfaces tab. See 31.7 “To create an MPLS interface” (p. 1120) for more information about how to configure MPLS interfaces.

24

To configure the collection of ingress accounting statistics for the MPLS instance, click on the Accounting tab.

25

To configure the collection of ingress statistics for LSPs, perform the following.

Use the following steps:

1. Click on the LSP tab.
2. Click Create. The IngStatsPolicy, Routing Instance (Create) form opens.
3. Configure the parameters in the LSP Information panel.
4. Select an accounting policy in the Ingress Accounting Statistics panel. Only accounting policies with the CombinedMplsLspIngressStats statistics type are available for selection.
5. Configure the required parameters in the Ingress Accounting Statistics panel.

   Note:

   The collection of egress accounting statistics is configured in 31.11 “To create a Dynamic LSP” (p. 1126) for Dynamic LSPs and 31.18 “To create a Point-to-Multipoint LSP” (p. 1137) for P2MP LSPs.

26

To configure the collection of ingress statistics for LSP templates:

Use the following steps:

1. Click on the LSP Template tab.
2. Click Create. The LSP Template Ingress Stats, Routing Instance (Create) form opens.
3. Configure the required parameters in the LSP Information panel.
4. Select an accounting policy in the Ingress Accounting Statistics panel. Only accounting policies with the CombinedMplsLspIngressStats statistics type are available for selection.
5. Configure the required parameters in the Ingress Accounting Statistics panel.

   Note:

   The collection of egress accounting statistics is configured in 31.29 “To create an LSP template MVPN policy” (p. 1156) .

27

To configure policies for class-based forwarding:

1. Click on the Class Forwarding Policies tab.
2. Click Create. The Class Forwarding Policy (Create) form opens.
3. Configure the required parameters in the General tab.
4. Click OK. The Class Forwarding policy is created.

28

To configure MPLS forwarding policies:

1. Click on the Forwarding Policies tab.
2. Click Create. The Forwarding Policies form opens.
3. Configure the parameters in the General tab as required.
4. Click on the Forwarding Policy tab and click Create. The Forwarding Policies - Forwarding Policy form opens.
5. Configure the parameters in the General tab as required.
6. Click on the Next Hop Group tab and click Create. The Next Hop Group form opens.
7. Configure the parameters in the General tab as required.
8. Click on the Next Hop tab and click Create to configure a primary next hop and, as required, a backup next hop.
9. Click OK and close the Next Hop Group, Forwarding Policy, and Forwarding Policies forms.

29

To view information about the MPLS instance, click on the following tabs:

- In Segments to view the MPLS ingress segment information, such as label assignments
- Cross Connects to view the LSP cross-connection information
- Out Segments to view the MPLS egress segment information, such as label assignments
- Static Label Maps to view static hop mapping information
- Statistics to view MPLS site and interface statistics
- Faults to view alarm information for the MPLS instance

30

Click View RSVP Site to view the associated RSVP instance on the NE.

31

Save the changes and close the forms.

END OF STEPS

31.7 To create an MPLS interface

31.7.1 Purpose

Create an MPLS interface by assigning a network interface to an MPLS instance.

This procedure requires the existence of a numbered or unnumbered L3 network interface on the NE that hosts the MPLS instance. See 27.1 “NE routing and forwarding” (p. 817) for more information about how to configure network interfaces.

31.7.2 Steps

1

Choose Routing from the NFM-P navigation tree view selector.

2

Navigate to Routing→NE→Routing Instance→MPLS.

3

Right-click on the MPLS object and choose Create Interface. The MPLS Interface (Create) form opens.

4

Assign a numbered or unnumbered L3 network interface as the MPLS interface by selecting an interface in the Interface panel.

5

Configure the required general parameters.

6

Configure the required parameters in the TE Metric panel.

7

To create a TP MEP interface:

Use the following steps:

1. Click Create in the TP MEP panel. The TP MEP Interface (Create) form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

8

Click Apply. The form displays additional tabs, and the NFM-P creates the MPLS interface and displays it below the MPLS and RSVP instances in the navigation tree.

9

To configure a Shared Risk Link Group:

Use the following steps:

1. Click on the Shared Risk Link Groups tab.
2. Click Create. The SharedRiskLinkGroup (Path/Routing Management: MPLS) Select form opens.
3. Search and select the required SRLGs and click OK. Up to 64 SRLGs can be associated with a specific MPLS interface. The SharedRiskLinkGroup (Path/Routing Management: MPLS) Select form closes and the selected SRLG is displayed.

10

A static label map is required when there are intervening unmanaged NEs between the managed NEs in an MPLS path. To configure a static label map:

Use the following steps:

1. Click on the Static Label Maps tab.

   Note:

   Nokia recommends that you use a static label mapping only to specify unmanaged NEs.
2. Click Create. The Static Label Map (Create) form opens. Alternatively, you can choose a Static Label Map and click Properties to edit the Static Label Map properties.
3. Configure the Label Action parameter. When you change the Label Action parameter from unspecified to Pop or Swap, the Static Label Map (Create) form changes to include other parameters. Configure the required parameters.

   The Egress Label, Next Hop and Enable Implicit Null Label parameters are only configurable when the Label Action parameter value is set to Swap.

   The Egress Label and Next Hop parameters are configurable only when the Label Action parameter value is set to Swap/Protect-Swap.

   The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

   The NFM-P sets the Label Action parameter value to Pop when the hop destination matches the destination LSP.
4. Save the changes and close the form.

11

To assign MPLS administrative groups to the interface, click on the Administrative Groups tab. After you assign administrative groups to an MPLS interface, the total value of the groups is displayed in a bit mask format by the Groups Included (bitmask) indicator on the General tab.

12

Save the changes and close the forms.

END OF STEPS

31.8 To create an MPLS path

31.8.1 Steps

1

Choose Manage→MPLS→MPLS Paths from the NFM-P main menu. The Manage MPLS Paths form opens.

2

Click Create. The Create MPLS Path form opens with the Name the MPLS Path step displayed.

3

Configure the parameters:

- Name
- Description

4

Click Next. The Define Source Site step is displayed.

5

Perform one of the following to specify an NFM-P-managed NE for the Starting Network Element parameter value.

a. Choose an NE from a list.

Use the following steps:

1. Click Select. The Select a Network Element - Select Source Site form opens.
2. Choose an NE and click OK. The management IP address of the NE is displayed as the source site of the MPLS path.

b. Enter the management IP address of the port.

6

Click Next. The Define the provisioned Path step is displayed.

7

Click Insert Hop to insert an MPLS path hop. The Hop for New MPLS Path (Create) form opens.

8

Configure the Specify Site parameter. Perform one of the following actions.

a. Specify the IP address of an unmanaged NE.

Use the following steps:

1. Choose Manually.
2. Configure the IP Address parameter.

b. Choose a managed NE.

Use the following steps:

1. Choose By Selection.
2. Click Select. The Select a Network Element - New MPLS Path form opens.
3. Choose an NE and click OK. The management IP address of the NE is displayed as the IP Address value.
4. Specify an alternative or loopback (only for Dynamic P2P RSVP LSPs) interface on the NE to be used as the hop point, if required. Otherwise, go to Step 9 .
5. Click IP Address parameter Select. The Select a Virtual Router Interface form opens with a list of the available interfaces.
6. Choose an interface and click OK. The Select a Virtual Router Interface form closes and the Hop for New MPLS Path (Create) form displays the IP Address of the selected loopback interface.

c. Specify the SID of an NE.

Use the following steps:

1. Choose SID Label.
2. Configure the Value parameter.

9

Configure the Hop Type parameter.

10

Click Apply.

11

Insert an additional hop, if required, by repeating Step 8 to Step 10 .

12

Click OK. The Create MPLS Path form reappears.

13

To change the hop sequence, choose a hop and click Move Up or Move Down. The first hop in the list is the first hop, and the last hop becomes the destination site when the form changes are saved.

14

Click Next. The Set Initial State step is displayed.

15

Configure the Administrative parameter.

16

Click Finish to save the configuration. You are prompted to view the MPLS path.

17

Enable the View the newly created MPLS path parameter to view the MPLS path configuration after closing the form, if required.

18

Close the forms.

END OF STEPS

31.9 To view an MPLS path

31.9.1 Steps

1

Choose Manage→MPLS→MPLS Paths from the NFM-P main menu. The Manage MPLS Paths form opens.

2

View the MPLS path information as required.

END OF STEPS

31.10 To create a static LSP

Note: When LSP on-demand resynchronization is enabled, manually resynchronize all of the planned hops in the static LSP before you create a static LSP. This is to ensure that validations on the static LSP can be run properly. See the procedure to enable LSP on-demand resynchronization in the *NSP System Administrator Guide* for more information.

31.10.1 Steps

1

Choose Manage→MPLS→Static LSPs from the NFM-P main menu. The Manage Static LSPs form opens.

2

Click Create. The Static LSP (Create) form opens.

3

Configure the parameters in the Identity panel.

Note: If a range policy is applied to a service tunnel, a grey text box appears beside the Service ID parameter to indicate that a range policy is in effect.
If a format policy is applied to a service tunnel, a drop-down menu appears beside the object field during object creation, to indicate that a format policy is enforced. When there is only one matching policy, the drop-down menu is dimmed. When there are multiple matching policies, the drop-down menu is used to choose a policy. The items in the drop-down menu are ordered by the policy's Priority parameter.

4

Configure the parameters in the Source panel.

The Administrative parameter applies to the source NE only, not to the entire static LSP

5

Select a destination site in the Destination panel.

The Administrative parameter applies to the source NE only, not to the entire static LSP

6

Click Apply.

7

To configure one or more static hops:

Use the following steps:

1. Click on the Static Hops tab.
2. Click Create. The Static Hop (Create) form opens.
3. Configure the Hop Index parameter.
4. Select a site in the Site panel.
5. Select an interface in the Interface panel.
6. Configure the parameters:

   The Egress Label, Enable Implicit Null Label and Next Hop parameters are configurable when the Label Action parameter value is set to Swap.

   The Egress Label and Next Hop parameters are configurable only when the Label Action parameter value is set to Swap/Protect-Swap.

   The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 "To configure an MPLS instance" (p. 1116) .

Note: Static hop configuration performed in CLI is not synced to the NFM-P.

8

Save the changes and close the forms.

END OF STEPS

31.11 To create a Dynamic LSP

31.11.1 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Click Create. The Create Dynamic LSP step form opens with the Identification step displayed.

3

Configure the required parameters.

Note: When a range policy is applied to a Dynamic LSP, a dimmed text field appears beside the Service ID parameter to indicate that a range policy is in effect.
If a format policy is applied to a Dynamic LSP, a drop-down menu appears beside the object field during object creation, to indicate that a format policy is in effect. When there is only one matching policy, the drop-down menu is dimmed. When there are multiple matching policies, the drop-down menu is used to choose a policy. The sequence of the options in the drop-down menu are specified by the policy Priority parameter.

4

Click Next. The Define Source and Destination Sites step is displayed.

5

Specify the source and destination sites for the Dynamic LSP.

Note: You can also manually specify an IP address for each parameter in the step.

Use the following steps:

1. Click on the Source Site ID parameter Select. The Select a Network Element - Create Dynamic LSP form opens with a list of the available sites.
2. Choose a site and click OK. The Select a Network Element - Create Dynamic LSP form closes and the Create Dynamic LSP form displays the source site information, which

includes the system Source IP Address parameter value. The parameter is automatically populated with the system IP address of the site.

Use the following steps:

1. If you are creating a P2P RSVP LSP, you can alternatively configure a loopback interface other than the system interface. Otherwise, go to 4 .
2. Click on the Source IP Address parameter Select. The Select a Virtual Router Interface form opens with a list of the available interfaces.
3. Choose an interface and click OK. The Select a Virtual Router Interface form closes and the Create Dynamic LSP form displays the Source IP Address and Source Interface Name information of the selected loopback interface.
4. Click on the Destination Site ID parameter Select. The Select a Network Element - Create Dynamic LSP form opens with a list of the available sites.
5. Choose a site and click OK. The Select a Network Element - Create Dynamic LSP form closes and the Create Dynamic LSP form displays the destination site information. The default displayed destination interface is the system interface.
6. If you are creating a P2P RSVP LSP, you can alternatively configure a loopback interface other than the system interface as the destination. Otherwise, go to Step 6 .
7. Click on the Destination IP Address parameter Select. The Select a Virtual Router Interface form opens with a list of the available interfaces.
8. Choose an interface and click OK. The Select a Virtual Router Interface form closes and the Create Dynamic LSP form displays the Destination IP Address and Destination Interface Name information of the selected loopback interface.

6

Click Next. The Auto Select Hop-Less MPLS Path step form opens. Configure the required parameters.

The Reserved Bandwidth (Mbps) parameter is configurable when the Auto Select Hop-less Path parameter is enabled.

7

If the Auto Select Hop-less Path parameter is enabled, go to Step 11 .

8

Click Next. The Add MPLS Paths form opens. This step binds the LSP to the MPLS path to create an MPLS path.

Use the following steps:

1. Configure the Path Destination Matching parameter.
2. Click Create. The LSP-Path Binding step form opens with the Choose Path Type step displayed. Configure the Type parameter.

Use the following steps:

1. Click Next. The Choose MPLS Path step is displayed. Choose an MPLS path to associate

with the LSP. You can also create an MPLS path by clicking Create MPLS Path. See 31.8 “To create an MPLS path” (p. 1122) for more information.

2. Click Next. The Set Traffic Options form opens. Configure the required parameters.

   The Path Preference is only configurable on Standby LSP paths (Type parameter set to standby).

3. Click Next. The Set Initial States form opens. Configure the required parameters.
4. Click Finish. The LSP-Path Binding form closes and a dialog box appears.
5. Click OK. The Create Dynamic LSP step form reappears.

9

Click Next. The Properties - Configuration step is displayed. Configure the required parameters.

10

Click Next. The Properties - PCE step is displayed. Configure the required parameters.

11

Click Next. The Properties - Traffic Engineering and Protection step is displayed. Configure the required parameters.

12

Click Next. The Properties - IGP Shortcut step is displayed. Configure the required parameters.

13

Click Next. The Properties - Fast Reroute step is displayed. Configure the required parameters.

14

Click Next. The Properties - Signalling step is displayed. Configure the required parameters.

15

Click Next. The Properties - Administrative Groups step is displayed.

16

Assign one or more MPLS administrative groups to the Dynamic LSP.

Use the following steps:

1. Select the required MPLS administrative groups in the Unassigned list.
2. Click on the right arrow button. The groups are assigned to the Dynamic LSP and moved to the Assigned list.

   After you assign administrative groups to an MPLS interface, the total value of the groups is displayed in a bit mask format by the Groups Included (bitmask) indicator on the General tab.

17

Click Next. The Properties - Class Forwarding step is displayed. Configure the parameters.

18

Click Next. The Properties - BFD step is displayed. Configure the parameters.

19

Click Next. The Properties - Entropy Label Capability step is displayed. Configure the Entropy Label parameter.

20

Click Next. The Add Path Profiles step is displayed. Specify a path profile ID for the LSP. If no path profiles are available for selection, then create a path profile.

21

Click Next. The Set Initial State step is displayed. Configure the Administrative parameter.

22

Click Finish. The NFM-P prompts you to view the Dynamic LSP.

23

Enable the View the newly created Dynamic LSP parameter to view the Dynamic LSP configuration after closing the form, if required.

Note: If you enable the View the newly created Dynamic Lsp parameter you can configure the collection of egress accounting statistics using the Dynamic LSP configuration form.

24

Click Close. The Create Dynamic LSP form closes.

25

If the View the newly created Dynamic LSP parameter is enabled in Step 23 , the Dynamic LSP (Edit) form opens with the newly created Dynamic LSP configuration displayed.

Use the following steps:

1. View the configuration, if required.
2. If you need to enable the collection of egress accounting statistics, go to 3 . Otherwise go to 5 .
3. Click on the Accounting tab.
4. Configure the parameters in the Egress Accounting Statistics panel.
   - Click Select. The Select Accounting Policy - Dynamic LSP form opens. Choose the required accounting policy and click on the OK button. Only accounting policies with the CombinedMplsLspEgressStats statistics type are available for selection.

- Collect Accounting Statistics
- Administrative State
- Statistics Mode. RSVP-TE supports Aggregate mode and Per Flow Control mode, SR and SR-TE supports only Aggregate mode.

Note:

The Ingress Accounting Statistics panel displays a read-only view. If an ingress accounting statistics record exists, you can click Properties to view the record.

The collection of ingress accounting statistics is configured in 31.6 “To configure an MPLS instance” (p. 1116) .

5. Click on OK to close the Dynamic LSP (Edit) form.

26

Click OK to close the Manage Dynamic LSPs form.

END OF STEPS

31.12 To create a segment routing TE LSP

31.12.1 Purpose

This procedure describes how to create and configure a segment routing TE LSP. The steps are provided in linear sequence. However, each step gives you the option to click Back to return to previous steps and review or change your configuration, as required.

31.12.2 Steps

1

Choose Manage→MPLS→Segment Routing LSPs from the NFM-P main menu. The Manage Segment Routing LSPs form opens.

2

Click Create. The Create Segment Routing LSP Wizard opens. Configure the required parameters in the Identification step.

3

Click Next. The Define Source and Destination Sites step is displayed. Configure the required parameters.

4

Click Next. The Auto Select Hop-less MPLS Path step is displayed. Perform one of the following:

a. If you want the NFM-P to automatically create a hopless MPLS path for the LSP, configure the required parameters and then go to step Step 8.

b. If you want to create or choose the MPLS path manually, then do not to select the Auto

Select Hop-less Path option and proceed to the next step.

5

Click Next. The Add MPLS Paths step is displayed. Create an MPLS path to bind to the LSP:

1. Click Create. The LSP-Path Binding Wizard opens.

   The Choose Path Type step does not require configuration. The Primary path type is selected by default and cannot be modified.
2. Click Next. The Choose MPLS Path step opens.
3. Choose an MPLS path for the LSP. If no eligible MPLS path is available, then click Create MPLS Path. See 31.8 “To create an MPLS path” (p. 1122). for details about creating an MPLS path.
4. Click Next. The Set Traffic Options step is displayed. Configure the required parameters.
5. Click Next. The Set Initial State is displayed. Set the Administrative state to Up.
6. Click Finish. You are returned to the Create Segment Routing LSP Wizard.

6

Click Next. The Properties – Configuration step is displayed. Configure the required parameters.

7

Click Next. The Properties – PCE step is displayed. Configure the required parameters.

8

Click Next. The Properties – Traffic Engineering And Protection step is displayed. Configure the required parameters.

9

Click Next. The Properties – IGP Shortcut step is displayed. Configure the required parameters.

10

Click Next. The Properties – Signalling step is displayed. Configure the required parameters.

11

Click Next. The Properties – Administrative Groups step is displayed. Configure the required parameters.

12

Click Next. The Properties – BFD step is displayed. Configure the required parameters. A BFD Template must be specified if BFD is enabled.

13

Click Next. The Add Path Profiles step is displayed. Specify a path profile ID for the LSP. If no path profiles are available for selection, then create a path profile.

14

Click Next. The Pce Associations step is displayed. Create an association for the LSP by providing an Association Type and selecting the configured Diversity or Policy.

15

Click Next. The Set Initial State step opens.

16

Click Finish. The Summary step is displayed, informing you that the segment routing TE LSP has been created and giving you the opportunity to view the configuration of the segment routing TE LSP that you created.

The LSP is created with the Administrative State parameter set to Down. See Step 5 in 31.13 “To configure a Dynamic or segment routing TE LSP” (p. 1131) to change the administrative state.

END OF STEPS

31.13 To configure a Dynamic or segment routing TE LSP

31.13.1 Steps

1

Choose Manage→MPLS→(Dynamic|Segment Routing) LSPs from the NFM-P main menu. The Manage (Dynamic|Segment Routing) LSPs form opens.

2

Click Search. A list of the selected type of LSPs is displayed.

3

Choose an LSP and click Properties. The (Dynamic|Segment Routing) LSP (Edit) form opens.

4

Configure the Description parameter.

5

Configure the Administrative parameter in the States panel.

6

Click on the Accounting tab and select an accounting policy.

7

Configure the parameters on the Properties tab.

8

Click on the OAM tab to view or configure OAM tests.

9

Click on the Statistics tab to view or configure statistics collection.

10

Click on the Tunnels tab. A list of LSP path bindings is displayed. Perform one of the following:

a. Create an LSP path binding.

   Use the following steps:

   1. Click Create. The LSP-Path Binding step form opens and the Choose Path Type step is displayed. The Type parameter is set to secondary if a primary path exists.
   2. Click Next. The Choose MPLS Path step is displayed. Choose an MPLS path to associate with the LSP. You can also create an MPLS path by clicking Create MPLS Path. See 31.8 “To create an MPLS path” (p. 1122) for more information.
   3. Click Next. The Set Traffic Options step is displayed. Configure the required parameters.
   4. Click Next. The Set Initial States step is displayed. Configure the Administrative State parameter.
   5. Click Finish to save your changes and close the form.

b. Configure an existing LSP path binding.

   Use the following steps:

   1. Choose an LSP path binding and click Properties. The LSP-Path Binding (Edit) form opens.
   2. Configure the required parameters.
   3. Select a BFD template in the BFD panel and configure the required parameters.

      The BFD template can be assigned to either the LSP-path or the LSP, but not both.

      When the LSP path is up and operational, the BFD on LSP Sessions tab contains an entry for the head and tail, respectively. When a BFD session is removed or the path is deleted, the BFD sessions are removed. The BFD session shows information about the current BFD session. BFD session information is automatically updated when a change occurs on the NE.
   4. Configure the required parameters on the Administrative Groups tab.

11

Click on the PCE-Associations tab. A list of LSP PCE associations is displayed. Perform one of the following:

a. Create an LSP PCE Association.

   Use the following steps:

   1. Click Create. The MPLS LSP PCE-Association form is displayed.
   2. Select the Association type (Diversity or Policy).
   3. Select the pre-configured Associations from the list.
   4. Save your changes and close the form.

b. Configure an existing LSP PCE Association.

   Use the following steps:

   1. Choose an LSP PCE Association and click Properties. The MPLS LSP PCE-Association (Edit) form opens.
   2. Configure the required parameters.
   3. Save your changes and close the form.

12

Save your changes and close the form.

END OF STEPS

31.14 To create a Dynamic or segment routing LSP from a tunnel template

31.14.1 Purpose

Before you can create an LSP from a tunnel template, you must create the tunnel template. You can use an existing LSP to create a tunnel template.

31.14.2 Steps

1

Choose Manage→MPLS→(Dynamic|Segment Routing) LSPs from the NFM-P main menu. The Manage (Dynamic|Segment Routing) LSPs form opens.

2

Click Create from Template. A Create (Dynamic|Segment Routing) LSP from Template form opens with a list of tunnel templates.

3

Choose a template and click OK. A Create (Dynamic|Segment Routing) LSP from Template form opens.

4

Configure the required parameters.

5

Specify the source sites for the LSP.

1. Select a Source Site ID and click OK.
2. Select a Source IP Address and click OK.

6

Configure the Destination Site ID and the Destination IP Address parameters.

7

Click on the Accounting tab and select an accounting policy.

8

Click OK.

9

Configure the Collect Accounting Statistics and the Administrative State parameters.

10

Click on the Properties tab and configure the required parameters.

Note: The number of parameters that can be configured depends on how the template was created. All of the parameters are configurable when the template is based on an object. A subset of the parameters can be configured when the template is based on an object class.

11

Click OK. The Create (Dynamic|Segment Routing) LSP from Template form closes.

12

When the Show Created Object parameter is enabled on the template, the (Dynamic|Segment Routing) LSP (Edit) form opens. Close the (Dynamic|Segment Routing) LSP (Edit) form.

13

Close the Manage (Dynamic|Segment Routing) LSPs form.

END OF STEPS

31.15 To list Dynamic or segment routing LSPs

31.15.1 Steps

1

Choose Manage→MPLS→(Dynamic|Segment Routing) LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Choose an LSP and click Properties. The (Dynamic|Segment Routing) LSP (Edit) form opens.

3

To view LSP path information:

Use the following steps:

1. Click on the Tunnels tab.
2. Double-click on an LSP path.
3. Review the state information.

   The LSP path information includes whether the bypass tunnel for the LSP is active, or whether the LSP is in a fast reroute state.

END OF STEPS

31.16 To view ping results on a BFD LSP session for a Dynamic LSP

31.16.1 Steps

The ping results are available only if the LSP BFD session has been established.

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Choose an LSP and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the BFD On LSP Sessions tab, select a session with a Link Type of LSP Head and click Properties. The BFD On LSP Session - LSP (View) form opens.

Note: The session ping results are available only for LSP Head BFD sessions.

4

Review the information on the General tab, as required, and click on the Ping Result tab.

5

Click Resync on the Ping Result tab to display the latest ping results.

END OF STEPS

31.17 To run an OAM validation test for a Dynamic or segment routing LSP

31.17.1 Steps

1

Choose Manage→MPLS→(Dynamic|Segment Routing) LSPs from the NFM-P main menu. The Manage (Dynamic|Segment Routing) LSPs form opens.

2

Choose an LSP and click Properties. The (Dynamic|Segment Routing) LSP (Edit) form opens.

3

Click One Time Validation.

4

See 90.55 “To run a one-time validation test on a service” (p. 3084) for more information on using the One Time Validation function and reviewing the test results.

END OF STEPS

31.18 To create a Point-to-Multipoint LSP

31.18.1 Steps

1

Choose Manage→MPLS→Point-to-Multipoint LSPs from the NFM-P main menu. The Manage Point-to-Multipoint LSPs form opens.

2

Click Create. The P2MP LSP (Create) form opens.

3

Configure the required parameters.

When you configure the System ID (Loopback IP Address) parameter, the Management IP Address and Site Name parameters are automatically populated with the values associated with that NE. The IP Address parameter is also automatically populated, however, you can specify a new value for the parameter, if required.

4

Click on the Properties tab.

5

Configure the parameters in the Traffic Engineering and Protection panel.

6

Configure the parameters in the Fast Reroute panel.

The Node Protect parameter is only configurable when the Reserved Bandwidth (Mbps) parameter is set to a value greater than 0.

7

Configure the Make before Break parameter.

8

Configure the parameters in the CSPF panel.

9

Configure the parameters in the Signalling panel.

10

To configure a P2MP primary instance:

Use the following steps:

1. Click on the P2MP Primary Instance tab and click Create. The P2MP Instance (Create) form opens.
2. Configure the required parameters.
3. Click on the Properties tab.
4. Configure the parameters in the Traffic Engineering Properties panel.

   When you enable the Inherit Value parameter, the Hop Limit parameter for the P2MP Instance is set to the same value as you configured for the parent P2MP LSP.
5. Configure the parameters in the Make Before Break panel.

   When you enable the Inherit Value parameter, the Make before Break parameter for the P2MP instance is set to the same value as you configured for the parent P2MP LSP.
6. Configure the parameters in the Administrative Groups panel.

   When you enable the Inherit Value parameters, the Included and/or Excluded Groups for the P2MP instance are set to the same values you configure for the parent P2MP LSP. If you do not enable the Inherit Value parameters, you can configure the Included Groups and Excluded Groups for the P2MP Instance independently from the parent P2MP LSP. See 8 .
7. To create one or more S2L paths click on the S2L Paths tab and click Create. The S2L Path (Create) form opens. Configure the required parameters and save and close the form.

Each S2L path object represents a root-to-leaf (S2L) sub-LSP path for the primary instance of the P2MP LSP.

8. To assign one or more MPLS administrative groups to the P2MP Instance, click on the Administrative Groups tab.
9. Choose the required Included Groups in the Unassigned list and click on the right arrow button. The groups are assigned to the Point-to-Multipoint LSP and moved to the Assigned list.
10. Select the required Excluded Groups in the Unassigned list and click on the right arrow button. The groups are assigned to the Point-to-Multipoint LSP and moved to the Assigned list.

    After you assign the Included and Excluded Groups, the total value of the groups is displayed in a bit mask format by the Groups Included and Groups Excluded indicators on the P2MP Instance Properties tab.
11. Save your changes and close the form.

11

To assign one or more MPLS administrative groups to the P2MP LSP:

Use the following steps:

1. Select the required Included Groups in the Unassigned list and click on the right arrow button. The groups are assigned to the Point-to-Multipoint LSP and move to the Assigned list.
2. Select the required Excluded Groups in the Unassigned list and click on the right arrow button. The groups are assigned to the Point-to-Multipoint LSP and move to the Assigned list.

   After you assign the Included and Excluded Groups, the total value of the groups is displayed in a bit mask format by the Groups Included (bitmap) and Groups Excluded (bitmap) indicators on the P2MP LSP Properties tab.

12

Save your changes and close the forms.

13

To complete the configuration of the Point-to-Multipoint LSP, you must:

a. Configure the required PIM tunnel interfaces. See 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

b. Configure tunnel interfaces on the required IGMP routing instance. See 28.104 “To configure an IGMP site on a router” (p. 1017) .

c. Configure tunnel interfaces on the required multicast info policy. See 52.13 “To configure an ingress multicast information policy” (p. 1720) .

14

After you complete the configuration of the Point-to-Multipoint LSP, you can reopen the P2MP LSP form to create, configure and run OAM diagnostics, or configure accounting statistics collection, if required. To configure OAM diagnostics, go to Step 15 . To configure accounting statistics collection, go to Step 16 .

15

Perform the following to configure OAM diagnostics.

Use the following steps:

1. Click on the Tests tab. The Test Suite, P2MP Ping, and P2MP Trace tabs are displayed.
2. Create or search for the test suite or particular type of test you want to run from the tab pages. Click on either the Create or Search buttons as required.
3. Click Execute to run a selected test.

   See Chapter 89, “Service Test Manager” and Chapter 90, “OAM diagnostic tests” for more information about how to configure OAM diagnostics.

16

Perform the following to configure accounting statistics.

Use the following steps:

1. Click on the Accounting tab.
2. Select an accounting policy.
3. Configure the required parameters.

END OF STEPS

31.19 To view an MVPN Point-to-Multipoint LSP object

31.19.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service, and click Properties. The VPRN Service (Edit) form opens.

3

Expand VPRN Service→Sites→Routing Instance - NE System ID→Routing Instance→Routing Instance: *VPRNinstance*→Protocols→PIM. The PIM Site (Edit) form opens.

4

Click on the MVPN Inclusive Tunnel tab, then on the RSVP tab.

5

Choose an interface from the Inclusive Provider Tunnel PMSIs panel and click Properties. The Inclusive RSVP PMSI Interface (Edit) form opens.

6

Click Properties next to the P2MP LSP ID parameter. The P2MP LSP (View) form opens.

7

Click on the tabs to view information about the MVPN P2MP LSP object.

8

Close the forms.

END OF STEPS

31.20 To create a Manual Bypass LSP

31.20.1 Steps

1

Choose Manage→MPLS→Manual Bypass LSPs from the NFM-P main menu. The Manage Manual Bypass LSPs form appears.

2

Click Create. The Create Bypass-only LSP form opens with the Identification step displayed.

3

Configure the required parameters.

4

Click Next. The Define Source and Destination Sites step is displayed.

5

Specify the source and destination sites for the manual bypass LSP.

Note: You can also manually specify an IP address for the source and destination sites in this step.

Use the following steps:

1. Click on the Source Site ID parameter Select button.

2. Choose a site and click OK. The Create Bypass-only LSP form displays the source site information, which includes the Source IP Address parameter value.

Use the following steps:

1. Click on the Destination Site ID parameter Select button.
2. Choose a site and click OK. The Create Bypass-only LSP form displays the destination site information. The displayed destination interface is the system interface because an LSP can only terminate on a system interface.

6

Click Next. The Auto Select Hop-Less MPLS Path form opens.

7

Configure the Auto Select Hop-Less Path parameter.

8

If the Auto Select Hop-less Path parameter in Step 7 is enabled, go to Step 10 .

9

Click Next. The LSP-Path Binding opens, with the Choose MPLS Path step displayed.

Use the following steps:

1. Choose an MPLS path to associate with the LSP or create one by clicking on the Create MPLS Path button. See 31.8 “To create an MPLS path” (p. 1122) .
2. Click Next. The Set Traffic Options step is displayed.
3. Configure the required parameters.
4. Click Next. The Set Initial States step is displayed.
5. Configure the Administrative parameter.
6. Click Finish. The LSP-Path Binding form closes and a dialog box appears.
7. Click OK. The Create Bypass-only LSP form reappears.

10

Click Next. The Properties step is displayed.

11

Configure required parameters.

12

Click Next. The Admin Groups step is displayed.

13

Assign one or more included and excluded MPLS administrative groups to the bypass-only LSP.

Use the following steps:

1. Choose the required MPLS administrative groups in the Included Groups - Unassigned and Excluded Groups - Unassigned lists.
2. Click on the right arrow button. The groups are assigned to the bypass-only LSP and move to the Assigned list.

   Note:

   When you assign included and excluded administrative groups to an MPLS interface, the assigned groups are displayed in bit mask form by the Groups Included/Excluded (bitmap) indicator.

14

Click Next. The Set Initial State step is displayed.

15

Configure the Administrative parameter.

16

Click on the Finish button. The NFM-P prompts you to view the bypass-only LSP.

17

Enable the View the newly created Bypass Only Lsp parameter, if required.

18

Close the form(s).

END OF STEPS

31.21 To configure a Manual Bypass LSP

31.21.1 Steps

1

Choose Manage→MPLS→Manual Bypass LSPs from the NFM-P main menu. The Manage Manual Bypass LSPs form opens.

2

Choose a Bypass-only LSP and click Properties. The Bypass-only LSP (Edit) form opens.

3

Configure the required parameters.

4

Click on the Properties tab.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

31.22 To configure an LSP path

31.22.1 Purpose

Use this procedure to configure an LSP path. This procedure applies to Dynamic LSPs, SR-TE LSPs, and Manual Bypass LSPs, unless otherwise noted. Dynamic LSPs are used to demonstrate the steps.

31.22.2 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Choose a Dynamic LSP from the list and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the Tunnels tab. A list of LSP paths is displayed.

4

If you are accessing this form to update the MPLS path for a Dynamic LSP, then go to Step 8 . Otherwise, go to Step 5.

5

Choose an LSP-Path Binding and click Properties. The LSP-Path Binding (Edit) form opens.

6

Configure the required parameters.

7

Replace the existing MPLS path for a primary or secondary LSP with another MPLS path, if required, by proceeding as follows:

Note: You can update the existing MPLS path only for a primary or secondary LSP path that has the Make before Break parameter enabled. The MPLS path you want to use must already exist on the NE, and must not be used by any other LSP paths under the parent LSP.

Use the following steps:

1. Click on the Update MPLS Path button. The Choose MPLS Path form opens and displays the eligible MPLS paths.
2. Choose the required MPLS path and click OK.

   The Choose MPLS Path form closes and the LSP-Path Binding (Edit) form is refreshed with the parameters related to the updated MPLS path selection.
3. Click OK. The LSP-Path Binding (Edit) form closes and the LSP-Path Binding list is refreshed with the updated MPLS path you selected.
4. Go to Step 9 .

8

Choose the MPLS path that you want to use to replace the existing path.

Note: You can update the existing MPLS path only for a primary or secondary LSP path that has the Make before Break parameter enabled. The MPLS path you want to use must already exist on the NE, and must not be used by any other LSP paths under the parent LSP.

Use the following steps:

1. Click on the Update MPLS Path button. The Choose MPLS Path form opens and displays the eligible MPLS paths.
2. Choose the required MPLS path and click OK.

   The Choose MPLS Path form closes and the LSP-Path Binding (Edit) form is refreshed with the parameters related to the updated MPLS path selection.

9

Click the Set as Secondary or Set as Standby button for a selected LSP-Path Binding to change its Type to one of these choices, if required.

10

Click the Switch Standby to Standby button for a selected LSP-Path Binding if you need to switch from the current active standby LSP path to another standby LSP path. This applies only to Dynamic LSPs.

The button choices include:

- Switch: This option allows you to switch manually from an active standby path to another inactive standby path. You can do this as often as required, without having to select Clear first.
- Force Switch: This option is essentially the same as the Switch option, but you cannot switch back unless you first press Clear to remove the Force Switch flag.
- Clear: This option is used to clear the Force Switch flag.

11

Click on the Administrative Groups tab.

12

Assign an MPLS administrative group to the Dynamic LSP.

Use the following steps:

1. Choose an MPLS administrative group in the Unassigned list.
2. Click on the right arrow button. The group is assigned to the Dynamic LSP and moves to the Assigned list.

   After you assign administrative groups to an MPLS interface, the total value of the assigned groups is displayed in a bit mask format by the Groups Included indicator on the General tab.

13

Click on the General Tab and in the Shared Risk Link Group panel, configure the Enable SRLG for FRR or Enable SRLG parameter, if available and required.

14

Click on the Tests tab to configure OAM diagnostics for the LSP path, if required. The supported test type tabs appear. See Chapter 90, “OAM diagnostic tests” for more information about how to configure OAM diagnostics.

15

View path hops information on the following tabs, as available and required:

- Provisioned Path
- Actual Path
- CSPF Path

16

Save your changes and close the forms.

END OF STEPS

31.23 To create an LSP path using a tunnel template

31.23.1 Purpose

Before you can create an LSP path from a tunnel template, you must create the tunnel template. You can use an existing LSP path to create a tunnel template.

31.23.2 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Click on the Create from Template button. The Create Dynamic LSP from Template list form opens displaying previously created templates.

3

Choose an LSP template. The template must be one where you associated a child LSP Path template to the parent LSP template.

4

Click OK. The Create Dynamic LSP from Template form opens. See 31.14 “To create a Dynamic or segment routing LSP from a tunnel template” (p. 1134) for information about creating a Dynamic LSP from a template.

5

Click OK and perform one of the following:

a. If the Show Created Object parameter is enabled on the template, the Dynamic LSP (Edit) form opens. Go to Step 6 .

b. Open the Dynamic LSPs (Edit) form.

Use the following steps:

1. Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.
2. Select the Dynamic LSP that you created in Step 4 and click Properties. The Dynamic LSP (Edit) form opens.

6

Click on the LSP-Path Bindings tab.

7

Click on the Create from Template button. A pop-up menu appears with a selection of LSP path templates.

8

Click on the template you require for the LSP path.

9

Configure the required parameters.

10

Click on the Inheritance tab.

11

Configure the Overridden Properties parameter.

12

Save your changes and close the forms.

END OF STEPS

31.24 To configure an LSP Path optimization policy

31.24.1 Purpose

An LSP Path optimization policy allows you to filter LSP Path candidates and set the execution rules to determine the candidacy and priority of an LSP Path for re-signaling. You can define LSP Paths that are eligible for re-signaling and if a path is selected, the LSP Path is added to a candidate list. Execution rules determine the sequence in which the LSP Paths can re-signal and the number of seconds between re-signaling. Execution rules apply to the entire candidate list. Each time that optimization starts, the candidate list is filtered according to the candidate definition values and then resequenced according to the sequencing target and order values.

To ensure that NEs are not overloaded by LSP re-signaling, an LSP Path optimization schedule can be created to identify when re-signaling can be evaluated and executed. You can also manually execute LSP Path optimization schedules. To avoid network congestion and maintain QoS standards, only one policy should be executed at a time. Repetitive re-signaling requests are ignored, including manual or scheduled re-signaling requests. However, a re-signaling request that

is initiated from the LSP Path binding properties form executes regardless of whether another LSP optimization policy is in progress.

The NEs notify the NFM-P about the current state and outcome of LSP Path re-signaling. You can access the LSP Path property form to view information about the state of LSP Path optimization. The following parameters are displayed:

- The Resignal Eligible parameter indicates whether an LSP Path is eligible for optimization. When the parameter value is true, the LSP Path can be optimized. When the parameter value is false, the LSP Path is not eligible for optimization.
- The Last Performed Type parameter indicates how re-signaling was performed. Typically the value is Manual Resignal. However, if a schedule was set, the value is Timer Based Resignal.
- The Last Performed parameter indicates the last time the LSP Path was chosen to be re-signaled.
- The Last Performed State parameter indicates the re-signal state of a specific LSP Path.

31.24.2 Steps

1

Choose Manage→MPLS→LSP Path Optimization from the NFM-P main menu. The Manage LSP Path Optimization form opens.

2

Choose LSP Path Optimization Policy (Path/Routing Management: MPLS).

3

Click Create and choose Optimization Policy.

The LSP Path Optimization Policy (Create) form opens with the General tab displayed.

4

Configure the required parameters.

5

Click on the Candidate Definition tab to apply filters to define eligible LSP Paths for optimization. If no filters are applied, then all eligible LSP Paths will be listed in the Execution Candidates form. Perform one of the following for LSP Path Filters and/or Dynamic LSP Filters.

a. Select an existing filter:

Use the following steps:

1. Click on Select beside the Filter Name parameter. A list of filters is displayed.
2. Choose a filter and click OK. The list closes and the Filter Name field is populated with your selection.

b. Create a new filter:

Use the following steps:

1. Click on Select beside the Filter Name parameter. A list of filters is displayed.
2. Click Create. The Create LSP Path/Dynamic LSP Filter forms opens.
3. Choose an Attribute and assign an associated Function and Value to it.
4. Click Add. The Attribute and its Function and Value appear in the Filter field. This is the first portion of the string that will be used to filter the execution candidates.
5. Add to the filter string by selecting an operator from the Operators drop-down menu, and then define a subsequent Attribute and its associated Function and Value.
6. Click Add. The Operator and subsequent Attribute (and its associated Function and Value) are added to the filter string in the Filter field.
7. Continue to build the filter string by repeating 5 and 6 , as required.
8. Click Save. The Save Filter dialog box appears.
9. Configure the parameters and click Save.
10. Close the Create Filter form. The newly-created filter appears in the filter selection list.
11. Select the newly-created filter and click OK. The list closes and the Filter Name field is populated with your selection.

6

To further limit the number LSP Paths that will be listed in the Execution Candidate form, define a Rule-Based Group to add only those NEs that contain the LSP Paths that you need include. Perform one of the following:

a. Select an existing group.

Use the following steps:

1. Click on Select beside the Group Name parameter. The Select Rule-Based Group form is displayed.
2. Choose a rule-based group and click OK. The list closes and the Group Name field is populated with your selection.
3. Go to Step 7 .

b. Create a new group.

Use the following steps:

1. Click on Select beside the Group Name parameter. The Select Rule-Based Group form is displayed.
2. Click Create. The Rule-Based NE Group (Create) forms opens with the General tab displayed.
3. Configure the parameters and click Apply.
4. Click the Group Members tab and click Create. The Select Network Elements list is displayed.
5. Choose the required NEs and click OK. The list closes and the selected NEs are listed in the Group Members tab.

Note:

If you configured the Order parameter on the General tab in 3 as Ordered Group, then you can only select one NE at a time. When you click OK, the selected NE will appear in the list. To add the next NE in the order, click Create again, choose another NE, then click OK. Repeat as required.

6. Click OK to save your configuration and close the Rule-Based NE Group (Create) form. The newly-created group appears in the Select Rule-Based Group form.
7. Select the new group and click OK. The list closes and the Group Name field is populated with your selection.

7

Click on the Execution Rules tab.

8

Configure the parameters and click on the Apply button. The form displays with schedule information.

Note: The Schedule Optimization button does not appear during LSP path optimization policy creation. The LSP path optimization policy must be created first before a schedule can be applied. You must click on the Apply button.
After a schedule is applied to a LSP Optimization Policy, you cannot make changes to the schedule from the Optimization Scheduled Task form.

9

Click on the Schedule Optimization button. An Optimization Scheduled Task (Create) form opens.

10

Configure the parameters.

11

Click on the Select button beside the ID parameter. A Select Schedule - Optimization Scheduled Task form opens.

12

Perform one of the following:

a. Choose an existing schedule.

b. Click Create to create a schedule. The NFM-P Schedule (Create) form opens. See 5.6 “To configure an NFM-P-based schedule” (p. 193) for information on creating a schedule.

13

Click on the OK button. The Optimization Scheduled Task form displays schedule and optimization information.

14

Click on the OK button. The Select Schedule - Optimization Scheduled Task form closes and the LSP Path Optimization Policy form reappears.

15

Click on the Execution Candidates tab.

16

Click on the Search button to display only the eligible LSP Paths that were filtered in Step 5 and Step 6 . If a filter was not applied, all eligible LSP Paths are listed.

17

Click on the Optimize button. To view the results of optimization, see 31.26 “To view LSP Path optimization policy results” (p. 1153) .

18

Close the form.

END OF STEPS

31.25 To terminate an LSP Path optimization policy that is in progress

31.25.1 Steps

1

Choose Manage→MPLS→LSP Path Optimization from the NFM-P main menu. The Manage LSP Path Optimization form opens.

2

Choose LSP Path Optimization Policy (Path/Routing Management: MPLS).

3

Create a filter to search for optimization policies that are in progress.

Use the following steps:

1. Click the Filter icon. The Manage LSP Path Optimization - Filter form opens.
2. Select General→Optimization from the Attributes drop-down menu.
3. Set the Function to EQUALS and the Value to In Progress, then click the Add button.
4. Click Apply and then Close. The Manage LSP Path Optimization - Filter form closes and the filter is applied to the Search list.

4

Choose the required LSP Path optimization policy and click Properties. The LSP Path Optimization Policy (Edit) form opens.

5

Click on the Stop Current Optimization button. The Optimization parameter value changes from In Progress to Not In Progress.

6

Close the form.

END OF STEPS

31.26 To view LSP Path optimization policy results

31.26.1 Purpose

Use the following procedure to view the re-signaling information as a result of the LSP Path optimization policy execution. See 31.24 “To configure an LSP Path optimization policy” (p. 1148) for information on configuring an LSP Path optimization policy.

31.26.2 Steps

1

Choose Manage→MLPS→LSP Path Optimization from the NFM-P main menu. The Manage LSP Path Optimization form opens.

2

Select LSP Path (Path/Routing Management MPLS) and click Search. A list of LSP Paths appears.

3

Choose the LSP Path on which you executed an LSP optimization policy and click Properties. The LSP-Path Binding (Edit) form opens.

4

View the following parameters, as described in 31.24 “To configure an LSP Path optimization policy” (p. 1148) :

- Resignal Eligible
- Last Performed Type
- Last Performed
- Last Performed State

5

Close the forms.

END OF STEPS

31.27 To view detour and bypass path information

31.27.1 Purpose

Perform this procedure to view information regarding detour and bypass tunnel paths and their protected LSPs. Detour and bypass paths are typically configured for fast-reroute-enabled LSPs on NEs for service protection against NE or link failure. The RSVP protocol is used to detect an NE or link loss.

Note: This procedure only applies to Dynamic LSPs.

Detours are employed in One-to-One backup configurations. A separate backup LSP is established for each LSP that is backed up.

Bypass Tunnels are employed in Many-to-One backup configurations in which a single backup LSP is used to back up multiple original LSPs.

Whether a path is a detour or bypass path is determined by the Backup Type parameter value specified during LSP creation. The Backup Type parameter can be set only when the Fast Reroute parameter for the LSP is set to true.

31.27.2 Steps

1

Choose Manage→MPLS→Dynamic LSPs. The Manage Dynamic LSPs form opens.

2

Select the required Dynamic LSP and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the RSVP Sessions tab. The detour or bypass originating, transiting, and terminating RSVP sessions associated with the LSP are listed. Type, site, next and previous hop, and other related information is available.

Note: The SRLG Disjoint FRR Properties and SRLG Not Strict FRR Properties indicators apply to primary LSP sessions only, and not to bypass LSP sessions.

4

Select the desired RSVP session of type detour or bypass and click Properties. The Session (View) form opens.

5

To view the Protected LSPs for the specific bypass or detour.

Use the following steps:

1. Click on the Protected LSP Paths tab. All the protected LSPs for a specific bypass or detour tunnel are listed.
2. Use the scroll bar at the bottom of the list frame to view the available information on the Protected LSPs.

6

To view the Detour or Bypass Tunnel Paths for the detour or bypass tunnel.

Use the following steps:

1. Click on the Detour/Bypass Tunnel Path tab. Information regarding the detour or bypass paths is displayed. Sufficient information is provided to permit tracing a path from its originating site, through its transit hops, and to its termination site.
2. Use the scroll bar at the bottom of the list frame to view the available information on the paths.

7

To view additional specific Detour or Bypass information, perform the following steps, as appropriate:

a. For a selected detour session, click on the Detour tab. The Detour Session ID, Detour PLR ID, and Avoided Downstream Site ID are listed.

b. For a selected bypass session, click on the Bypass Tunnel tab. The Bypass Tunnel's current State is displayed.

Note: For the NFM-P to accurately gather the information about the detour and bypass tunnel path (and the LSPs protected by a detour or a bypass tunnel), the NFM-P must be managing all routers participating in the LSP fast reroute. This is required by NFM-P for all SR and 7705 SAR network elements along the primary path of a fast reroute-enabled LSP.

8

Close the forms.

END OF STEPS

31.28 To view exclude route object information

31.28.1 Purpose

Perform this procedure to view the address of the NE to be avoided as signaled in the EXCLUDE_ ROUTE object.

Note: This procedure only applies to Dynamic LSPs.

31.28.2 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Select the required Dynamic LSP and click Properties. The Dynamic LSP (Edit) form opens.

3

Click on the RSVP Sessions tab.

4

Select the desired RSVP session and click Properties. The Session (View) form opens.

5

View the Exclude Node IP Address property. If no EXCLUDE_ROUTE object is signaled, the address is 0.0.0.0.

6

Close the forms.

END OF STEPS

31.29 To create an LSP template MVPN policy

31.29.1 Purpose

This procedure defines the template used to create LSPs across a multicast VPN. The LSP template MVPN policy can be configured locally on a specific NE or globally across the network.

The following can be created:

- P2MP
- one-hop P2P
- mesh P2P
- one-hop P2P SR TE
- mesh P2P SR TE
- OnDemand P2P SR TE

If you need to use entropy labels, the type must be One-Hop P2P LSP or Mesh P2P LSP.

31.29.2 Steps

1

Choose Policies→MPLS→LSP Template MVPN from the NFM-P main menu. The LSP Template MVPN Policies form opens.

2

Perform one of the following:

a. Configure the Policy scope parameter to Global.

b. Configure the Policy scope parameter to Local and click on the Select button to choose a device.

3

Click Create. The LSP Template MVPN Policy (Create) form opens.

4

Configure the parameters.

5

Click on the Properties tab.

Note: The panels and parameters that appear in the Properties tab can vary based on configuration.

6

Configure the required parameters.

7

Click on the PCE-Associations tab. A list of LSP Template PCE associations is displayed. Perform one of the following:

a. Create an LSP Template PCE Association.

Use the following steps:

1. Click Create. The MPLS LSP Template PCE-Association form is displayed.
2. Select the Association type (Diversity or Policy).
3. Select the pre-configured Associations from the list.
4. Save your changes and close the form.

b. Configure an existing LSP Template PCE Association.

Use the following steps:

1. Choose an LSP Template PCE Association and click Properties. The MPLS LSP Template PCE-Association (Edit) form opens.

2. Configure the required parameters.
3. Save your changes and close the form.

8

Click on the Accounting tab.

9

Configure the parameters.

10

Click on the Administrative Groups tab.

11

Assign one or more MPLS administrative groups to the LSP template MVPN policy.

Use the following steps:

1. Choose the required Included Groups in the Unassigned list.
2. Click on the right arrow button. The groups are assigned to the LSP and moved to the Assigned list for inclusion.
3. Choose the required Excluded Groups in the Unassigned list.
4. Click on the right arrow button. The groups are assigned to the LSP and moved to the Assigned list for exclusion.

   After you assign the Included and Excluded Groups, the total value of the groups is displayed in a bit mask format by the Groups Included and Groups Excluded indicators on the LSP Template MVPN Policy Properties tab.

12

Click OK to apply the changes. The LSP Template MVPN Policy, Global Policy (Create) form closes and the LSP Template MVPN Policies form reappears with the LSP MVPN template displayed.

13

Re-select the LSP MVPN template and click Properties. The LSP Template MVPN Policy - Global Policy (Edit) form opens.

14

Click on the Switch Mode button to distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

Note: The MVPN LSP template cannot be used until you choose a default MPLS path and configure the Administrative parameter in a local policy definition after the global policy is distributed. After the changes are applied to the local policy, the policy can be switched back to Sync With Global mode and the local changes are saved.

15

Click on the Local Definitions tab.

16

Select an LSP MVPN template and click Properties. The LSP Template MVPN Policy, Local Policy (Edit) form opens.

17

Click on the Switch Mode button to change the configuration mode to Local Edit Only.

18

Click on the Properties tab.

19

Click on the Select button to choose a default MPLS path. The Select Default MPLS Path - Lsp MVPN Template, Local Policy form opens with a list of default MPLS paths displayed.

20

Choose an MPLS path and click OK. The default MPLS path is displayed.

21

Click on the General tab.

22

Configure the Administrative parameter.

23

Save your changes and close the form.

END OF STEPS

31.30 To view LSP templates for MVPN created using CLI

31.30.1 Purpose

Perform this procedure to view read-only LSP templates created using CLI.

Note: You cannot modify the local or global definitions of these template types using the NFM-P.

You cannot distribute these template types to other NEs using the NFM-P.

See the specific device documentation for information on how to create a One-hop P2P LSP, Mesh P2P LSP, or SR TE template.

31.30.2 Steps

1

Choose Policies→MPLS→LSP Template MVPN from the NFM-P main menu. The LSP Template MVPN Policies form opens.

2

Perform one of the following:

a. Configure the Policy scope parameter to Global.

b. Configure the Policy scope parameter to Local and click on the Select button to choose a device.

3

Choose one of the following from the Type column drop-down menu:

- One-Hop P2P LSP
- Mesh P2P LSP
- Segment Routing

4

Click on the Search button. A list of Lsp MVPN Templates is displayed.

5

Select the LSP MVPN template and click Properties. The LSP Template MVPN Policy (Edit) form opens.

6

Click on the various tabs to view information about the LSP template.

7

Close the forms.

END OF STEPS

31.31 To view LSPs created by One-hop P2P and Mesh P2P templates

31.31.1 Purpose

Dynamic LSPs are automatically created when you use CLI to create and distribute a One-hop P2P, Mesh P2P, or SR TE LSP template, and then associate the LSP template with an auto-LSP. Perform the following procedure to view the LSPs that are automatically created by the templates.

Note: The dynamically-created LSPs cannot be associated with SDPs, and are not listed when you are creating an SDP.

31.31.2 Steps

1

Choose Manage→MPLS→Dynamic LSPs from the NFM-P main menu. The Manage Dynamic LSPs form opens.

2

Choose one of the following from the Type column drop-down menu.

- One-Hop Point-To-Point
- Mesh Point-To-Point
- Segment Routing

3

Click on Search to display a list of Dynamic LSPs.

4

Select an LSP and click Properties. The properties form for the object appears.

5

View the information on the various tabs, as required.

6

Close the forms.

END OF STEPS

31.32 To list and view MPLS objects

31.32.1 Steps

1

Choose Manage→MPLS→MPLS Objects from the NFM-P main menu. The Manage MPLS Objects form opens.

2

Select an MPLS Object Type and click on Search. A list of those MPLS objects is displayed.

3

Select an object and click Properties. The properties form for the object appears.

4

View the information on the various tabs, as required.

5

Close the forms.

End of steps

31.33 To create an administrative LSP tag

31.33.1 Steps

1

Choose Policies→Routing→Admin Tags from the NFM-P main menu. The Routing Policy - Admin Tags management form opens.

2

Click Create. The Admin Tag (Create) form opens.

3

Configure the tag name.

4

Click OK and confirm to save the form. The new administrative tag appears in the list. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

End of steps

31.34 To create an administrative tag policy

31.34.1 Steps

1

Choose Policies→Routing→Admin Tag Policy from the NFM-P main menu. The Routing Policy - Admin Tag Policy management form opens.

2

Click Create. The Admin Tag Policy (Create) form opens.

3

Configure the general parameters.

4

In the Included Admin Tags panel, click on the Create button and select one or more administrative tags from the list that appears, then click OK. LSPs with the selected tags will be used for auto-binding at sites using this policy.

5

In the Excluded Admin Tags panel, click on the Create button and select one or more administrative tags from the list that appears, then click OK. LSPs with the selected tags will be rejected for auto-binding at sites using this policy.

6

Click OK and confirm to save the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

31.35 To create a reserved label block

31.35.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MPLS. Right-click on a routing instance or MPLS object in the navigation tree and choose Properties. The MPLS (Edit) form opens.

2

Click on the Reserved Label Block tab.

3

Click Create. The Reserved Label Block form opens.

4

In the General tab, configure a name for the label block and the reserved label range.

5

Click OK and confirm to close the forms.

END OF STEPS

31.36 Workflow to collect segment routing TE LSP rate PM statistics

31.36.1 Stages

1

Create an accounting policy for the rate statistics. See *NSP NFM-P Statistics Management Guide* for more information.

Note: Configure the Type parameter of the policy to Combined MPLS SRTE Egress and set the Log To File Type to "No File".

2

Configure a segment routing LSP. See 31.13 “To configure a Dynamic or segment routing TE LSP” (p. 1132) for more information.

- Associate the accounting policy created in Stage 1 under Accounting tab.
- In Statistics tab, select MPLS LSP Egress Stats (Path/Routing Management MPLS) as the Object Type.

Note: The rate statistics are available as MIB-based accounting statistics.

To maintain accuracy and to avoid duplicate records, it is recommended to have the collection interval and the polling interval synchronised.

32 MPLS-TP

MPLS-TP overview

32.1 MPLS-TP overview

32.1.1 Overview

MPLS-TP is a set of MPLS protocol functions that enables the use of MPLS in transport networks and applications. MPLS-TP provides in-band proactive OAM and protection mechanisms that do not rely on a control plane, such as RSVP-TE, to operate. MPLS-TP enables MPLS to be deployed in a statically configured transport network without the need for a dynamic control plane.

Supporting NEs have two primary roles in an MPLS-TP network:

- as a gateway PE between IP/MPLS and MPLS-TP domains
- as an LSR or LER for MPLS-TP LSPs

The NFM-P can configure LSR NEs even when LERs are not configured, or when the LERs or LSRs are not supporting NEs. See the specific NE documentation for more information about MPLS-TP.

The NFM-P provides simple provisioning for MPLS-TP management in a transport network. MPLS-TP relies on static configuration, rather than a control plane, to ensure end-to-end path configuration and consistency. The NFM-P support of MPLS-TP includes:

- bidirectional co-routed MPLS-TP LSPs and PWs
- MPLS-TP identifiers for NEs, LSPs, and PWs
- linear protection for MPLS-TP LSPs with working and protection paths for each LSP
- OAM and protection with both IP and non-IP encapsulation
- proactive CC/V for MPLS-TP LSPs using BFD
- on-demand CV for MPLS-TP LSPs and PWs using LSP ping and trace and VCCV ping and trace diagnostics
- static PW status signaling, with support for PW redundancy, MC-LAG, MC-APS, BGP multi-homing and active and standby dual-homing into IES, VPRN, or VPLS
- stitching of static MPLS-TP PWs to T-LDP dynamically signaled PWs (LDP, RSVP, and BGP) on Epipe, Apipe, and Cpipe VLL switching sites

32.1.2 Proactive OAM templates

A proactive continuity check detects a loss of continuity defect between two MEPs in a MEG. Proactive connectivity verification detects an unexpected connectivity defect between two MEPs, and unexpected connectivity within the MEG with an unexpected MEP. The NFM-P implements both tests with a proactive OAM template.

Proactive OAM templates are based on BFD. BFD packets are sent using configurable timers. Continuity check packets contain a BFD control packet, while connectivity verification packets also

include an identifier for the source MEP. The destination MEP can detect a connectivity defect if it is receiving packets from an incorrect peer MEP. When a defect is detected, the NFM-P generates an alarm.

32.1.3 MPLS-TP LSP provisioning

The NFM-P significantly simplifies the provisioning process of bidirectional MPLS-TP LSPs, where using CLI can be cumbersome and prone to error. Using the NFM-P, all of the components of an end-to-end bidirectional MPLS-TP LSP are automatically configured. Using CLI, all of the components of a bidirectional MPLS-TP LSP, including interfaces and labels, must be configured. The NFM-P automatically creates opposite interfaces based on the physical links between NEs, and assigns in and out labels across all of the interfaces. However, if a physical link does not exist between NEs, the NFM-P cannot auto-create the interfaces. The NFM-P prompts you to manually create the interfaces.

Consider the following when you create a bidirectional MPLS-TP LSP using NFM-P simplified provisioning:

- the NEs must be NFM-P-managed
- MPLS-TP tunnels that are created on the NEs using CLI are discovered by the NFM-P

  Note:

  The NFM-P does not create the bidirectional MPLS-TP LSP when all of the MPLS-TP LSP endpoints (LERs) are not managed by the NFM-P. However, cross-connections (on LSRs) appear in the NFM-P database if the LSRs are managed.

In the example in the following figure, you configure the LER A and LER B using the NFM-P. For the bidirectional MPLS-TP LSP path, you identify the hops by specifying the interfaces. In this example, the interfaces are A1 on LER A and C1 on LSR C. Since MPLS-TP is co-routed, the other interfaces (B1 and C2) are determined by the NFM-P based on the link between the two NEs. The NFM-P automatically chooses all of the labels, in both directions, based on the available labels on the two adjacent NEs. The NFM-P also creates the LSR C object. You must configure proactive OAM on LER NEs.

Figure 32-1 Sample bidirectional MPLS-TP LSP

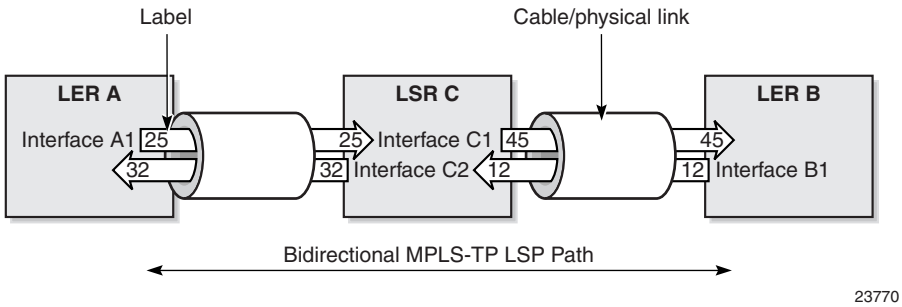

MPLS-TP workflow and procedures

32.2 Workflow to configure MPLS-TP

32.2.1 Stages

1

Enable MPLS-TP on the routing instances of all NEs that are to participate in the MPLS-TP network; see 32.3 “To enable MPLS-TP on a routing instance” (p. 1168) .

2

Create a BFD template policy; see 28.25 “To configure a BFD template policy” (p. 911) .

3

Configure MPLS-TP on the routing instance; see 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) .

4

Perform one of the following, as required:

a. Create a mesh of MPLS-TP LSPs or an MPLS-TP LSR cross-connect path; see 32.5 “To create an MPLS-TP LSP” (p. 1170) or 32.6 “To create an MPLS-TP LSR cross-connect path” (p. 1171) .

b. Create a bidirectional MPLS-TP LSP using simplified provisioning; see 32.7 “To create a bidirectional MPLS-TP LSP” (p. 1173) .

5

Create a IP/MPLS service tunnel, if required, and enable MPLS-TP; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .

6

Create a spoke SDP binding using an MPLS-TP service tunnel on the following, if required:

- VLL service; see 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) and 76.35 “To configure a spoke SDP binding with an L2TPv3 tunnel on a VLL Epipe site” (p. 2165)
- VPLS or MVPLS service; see 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386)
- IES; see 78.10 “To create an L2 SDP spoke termination on an IES service” (p. 2437)
- VPRN service; see 79.69 “To create an L2 SDP spoke termination on a VPRN service” (p. 2638)
- mirror service; see 93.6 “To create a source site on a mirror service” (p. 3169)

7

Configure static MPLS-TP pseudowires on the following, if required:

- VLL service; see 76.15 “To configure an MPLS-TP static pseudowire on a VLL spoke SDP binding” (p. 2136)
- VPLS service; see 77.93 “To configure an MPLS-TP static pseudowire on a VPLS spoke SDP binding” (p. 2392)
- IES; see 78.11 “To configure an MPLS-TP static pseudowire on an IES spoke SDP binding” (p. 2440)
- VPRN service; see 79.70 “To configure an MPLS-TP static pseudowire on a VPRN spoke SDP binding” (p. 2641)
- mirror service; see 93.8 “To configure an MPLS-TP static pseudowire on a mirror SDP binding” (p. 3172)

8

Perform the following OAM diagnostic tests on MPLS-TP objects, if required:

- VCCV ping OAM diagnostic test; see 90.8 “To create and run a VCCV ping OAM diagnostic test from the STM” (p. 3008)
- VCCV trace OAM diagnostic test; see 90.10 “To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM” (p. 3010)
- MPLS-TP LSP ping OAM diagnostic test; see 90.25 “To create and run a MPLS LSP ping OAM diagnostic test from the STM” (p. 3030)
- MPLS-TP LSP trace OAM diagnostic test; see 90.26 “To create and run a MPLS LSP trace OAM diagnostic test from the STM” (p. 3031)

32.3 To enable MPLS-TP on a routing instance

32.3.1 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance.

2

Right-click on a routing instance icon and choose Properties.

3

Click on the Protocols tab and enable the MPLS Enabled and MPLS-TP Enabled parameters.

4

Save your changes and close the form. The MPLS-TP icon appears in the navigation tree below the routing instance.

END OF STEPS

32.4 To configure MPLS-TP on a routing instance

32.4.1 Purpose

Use the following procedure to configure MPLS-TP on a routing instance. Ensure that you enable MPLS-TP on all NEs that are participating in the MPLS-TP network, as described in 32.3 “To enable MPLS-TP on a routing instance” (p. 1168) .

32.4.2 Steps

1

In the navigation tree Routing view, expand Network→*NE*→Routing Instance→MPLS-TP.

2

Right-click on a MPLS-TP icon and choose Properties. The MPLS-TP, Routing Instance (Edit) form opens.

3

Configure the required parameters

The Global ID and Node ID parameters are identifiers of an MPLS-TP site on the NE. The TP-LSP ID Minimum and TP-LSP ID Maximum parameters reserve a range of IDs that you use for MPLS-TP LSPs.

4

Configure the Administrative State parameter in the States panel.

Configure a proactive OAM template

5

Right-click on a MPLS-TP icon and choose Properties. The MPLS-TP, Routing Instance (Edit) form opens.

Perform the following steps:

1. Configure the required parameters.
2. Click Select in the BFD Template panel and click Search. A list of configured BFD template policies appears. See 28.25 “To configure a BFD template policy” (p. 911) for information about how to create a BFD template policy.
3. Choose a BFD template policy and click OK. The Proactive OAM Template (Create) form refreshes with the BFD template information.
4. Save your changes and close the form. The MPLS-TP, Routing Instance (Edit) form reappears.

Configure a protection template

6

Click on the Protection Template tab and click Create.The Protection Template, Routing Instance (Create) form opens.

Perform the following steps:

1. Configure the required parameters.
2. Save your changes and close the form. The MPLS-TP, Routing Instance (Edit) form reappears.

7

Save your changes and close the form.

END OF STEPS

32.5 To create an MPLS-TP LSP

32.5.1 Purpose

Use the procedure to create an MPLS-TP LSP. This procedure requires the existence of a numbered or unnumbered L3 network interface on the NE that hosts the MPLS-TP instance. See 27.1 “NE routing and forwarding” (p. 817) for more information about how to configure network interfaces.

32.5.2 Steps

1

Choose Manage→MPLS→MPLS-TP LSPs from the NFM-P main menu. The Manage MPLS-TP LSPs form opens.

2

Click Create and choose Create LER TP-LSP. The LER TP-LSP (Create) form opens.

3

Configure the required parameters in the Identity and Destination panels.

The TP-LSP ID must be in the range of the TP-LSP ID Minimum and TP-LSP ID Maximum parameters that you specified in Step 3 of 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) .

4

Click Select in the Source Site panel and choose a source site and click OK.

5

Click on the TP-LSP Path tab and click Create. The MPLS-TP LSP Path (Create) form opens.

Note: You can create one working path and one protecting path per LSP. As required, repeat Step 5 to Step 12 to create the alternate path type.

6

Configure the required parameters.

7

Click Select in the Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.

8

Click on the Path MEP tab and click Create. The Path MEP (Create) form opens with the General tab displayed.

9

Configure the required parameters.

10

As required, click Select in the Proactive OAM Template panel and choose a template and click OK. See 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) for information about how to create a Proactive OAM Template.

11

As required, click Select in the Protection Template panel and choose a template and click OK. See 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) for information about how to create a Protection Template.

12

Save your changes and close the forms.

END OF STEPS

32.6 To create an MPLS-TP LSR cross-connect path

32.6.1 Purpose

Use this procedure to create an MPLS-TP LSR cross-connect path. This procedure requires the existence of a numbered or unnumbered L3 network interface on the NE that hosts the MPLS-TP instance. See 27.1 “NE routing and forwarding” (p. 817) for more information about how to configure network interfaces.

32.6.2 Steps

1

Choose Manage→MPLS→MPLS-TP LSPs from the NFM-P main menu. The Manage MPLS-TP LSPs form opens.

2

Click Create and choose Create LSR Path. The LSR Path (Create) form opens.

3

Configure the required parameters in the Identity and States panel. By default, the Forward Path Enabled and Reverse Path Enable parameters are configured to enabled.

4

Click Select in the Site panel and choose a site to be used as the cross-connect and click OK.

5

Configure maintenance intermediate points, if required. The Maintenance Intermediate Points panel displays when you configure a cross-connect site in Step 4 .

Perform the following steps:

1. In the Maintenance Intermediate Points panel, click Create. The Maintenance Intermediate Point (Create) form opens.
2. Configure the required parameters.
3. Click OK to close the form.

6

Perform the following if you enabled the Forward Path Enabled parameter, otherwise, go to Step 8 .

1. Configure the required parameters in the Forward panel.
2. Click Select in the Forward Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.
3. If you choose a numbered outlink interface, configure the Forward Next Hop parameter.

7

Perform the following if you enabled the Reverse Path Enabled parameter, otherwise, go to Step 8 .

1. Configure the parameters in the Reverse panel.
2. Click Select in the Reverse Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.
3. If you choose a numbered outlink interface, configure the Reverse Next Hop parameter.

8

Click on the LSP Path Binding tab and click Create to associate the LSR cross-connect path with LER LSPs.The LSR Path Binding (Create) form opens.

9

Configure the required parameter and click OK.

10

Save your changes and close the form.

END OF STEPS

32.7 To create a bidirectional MPLS-TP LSP

32.7.1 Purpose

A bidirectional LSP consists of MPLS-TP LSPs or LER LSPs that are created on the edge NEs, and the LSR cross-connect—or transit path—that is created on the transit NEs. The NFM-P guides you through the creation of the bidirectional MPLS-TP LSP. When the Apply or OK button is disabled, verify that all necessary parameters are configured. When all possible parameters are configured, the NFM-P displays a Hint button to provide instructions on the necessary actions.

This procedure requires the existence of a numbered or unnumbered L3 network interface on the NE that hosts the MPLS-TP instance. See 27.1 “NE routing and forwarding” (p. 817) for more information about how to configure network interfaces.

Note: NFM-P guided provisioning is available only for the creation of bidirectional MPLS-TP LSPs. The Hint button does not appear if you are configuring an existing bidirectional MPLS-TP LSP.

32.7.2 Steps

1

Choose Manage→MPLS→MPLS-TP LSPs from the NFM-P main menu. The Manage MPLS-TP LSPs form opens.

2

Click Create and choose Create Bidirectional TP-LSP. The Bidirectional TP-LSP (Create) form opens.

3

Configure the required parameters.

The TP-LSP ID must be in the range of the TP-LSP ID Minimum and TP-LSP ID Maximum parameters that you specified in Step 3 of 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) .

4

Click Select in the LER A panel and choose a source site and click OK.

5

Click Select in the LER B panel and choose a source site and click OK.

6

Click on the LSP Path tab and click Create. The Bidirectional TP-LSP Path (Create) form opens.

7

Configure the required parameters

Note: You can create one working path and one protecting path per LSP. As required, repeat Step 6 to Step 17 to create the alternate path type.

The LSP Path Number parameter is a global LSP path number that configures the LSP path number on the LER A working and protecting paths, the LER B working and protecting paths, and on the LSR static hop.The LSP Path Number parameter must be the same on all of the hops.

8

Click on the LER A tab and click Select. The Select Out-Link - Bidirectional TP-LSP Path list form opens.

9

Choose a site and click OK.

10

Click on the Hint button. A window appears displaying an instructional message.

11

Perform one of the following:

a. If the message “You have configured the LER-A Out-Link Interface. As a result, NFM-P has created the next LSR Path Hop, and has auto-filled the data between the LER-A and the newly created LSR Path Hop” appears, perform the following steps:
   1. Click on the LSR Path Hops tab and choose the LSR path hop with a configuration state of Incomplete and click Properties. The LSR Hop (Edit) form opens.
   2. Configure the required parameters and click OK. Go to Step 14 .

b. If the message “You have configured the LER-A Out-Link Interface. NFM-P is unable to determine the next hop (this could be due to an unconnected cable). As a result, NFM-P has enabled the LSR Path Hop creation, and the LER-B Out-Link Interface selection.” appears, perform the following:
   - If the next hop in the path is not the LER B, go to Step 12 to create an LSR path hop.

- If the next hop is the LER B, go to Step 14 .

c. If the message “You have configured the LER-A Out-Link Interface. As a result, NFM-P has auto-completed the data between LER-A and LER-B.” appears, click OK to complete the path configuration.

12

Create an LSR path hop:

Perform the following steps:

1. Click on the LSR Path Hops tab and click Create. The LSR Hop (Create) form opens.
2. Configure the Display Name parameter.
3. Click Select in the Forward Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.
4. Click on the Hint button. The message “You have configured the Forward Out-Link Interface. As a result, NFM-P has auto-completed the data between this LSR Path Hop and the previous hop on the forward path.” appears.
5. Click Select in the Reverse Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.
6. Click OK to close the form.

13

Click on the Hint button. Perform one of the following based on the message returned:

a. If the message “You have configured the LSR Path Hop Reverse Out-Link Interface. NFM-P is unable to determine the next hop (this could be due to an unconnected cable). As a result, NFM-P has enabled the LSR Path Hop creation, and the LER-B Out-Link Interface selection.”, perform the following.

- If the next hop in the path is not the LER B, return to Step 12 to create an LSR path hop.
- If the next hop is the LER B, go to Step 14

b. If the message “You have configured the Reverse Out-Link Interface. As a result, NFM-P has created the next LSR Path Hop, and auto-completed the data between this LSR Path Hop and the newly created LSR Path Hop.” appears, perform the following steps:

1. Click on the LSR Path Hops tab.
2. Choose the LSR path hop with a configuration state of Incomplete and click Properties. The LSR Hop (Edit) form opens.
3. Configure the required parameters and click OK. Go to Step 14 .

c. If the message “You have configured the Reverse Out-Link Interface. As a result, NFM-P has auto-filled the data between this LSR Path Hop and LER-B.” appears, go to Step 17 .

14

Click on the LER B tab.

15

Click Select in the Out-Link panel and choose an interface and click OK. See 31.7 “To create an MPLS interface” (p. 1120) for information about how to configure MPLS interfaces.

16

Click on the Hint button. The message “You have configured the LER-B Out-Link Interface. As a result, NFM-P has auto-completed the data between LER-B and its next hop on the forward path. The path configuration is now complete.” appears, click OK to complete the path configuration.

17

Save your change and close the forms.

18

Create a BFD template on the LER A and LER B. See 28.25 “To configure a BFD template policy” (p. 911) .

19

Create a proactive OAM template on the LER A and LER B. See 32.4 “To configure MPLS-TP on a routing instance” (p. 1169) .

END OF STEPS

33 Service tunnels

Service tunnel overview

33.1 Service tunnel overview

33.1.1 Overview

A service tunnel is an entity used to uni-directionally direct traffic from one device to another device. The service tunnel is provisioned to use a specific encapsulation method, such as GRE or MPLS, and the services are then mapped to the service tunnel. For instance, multi-NE VLL and VPLS traffic is transmitted using uni-directional service tunnels in this way.

The most common type of tunnel used in NFM-P is a Service Distribution Point binding. Service tunnels originate on an SDP on a source NE and terminate at a destination NE. The destination NE directs packets from the service tunnel to the correct service egress interfaces (SAPs) on that device. Services that originate and terminate on the same NE do not require service tunnels, because the same NE is both the source and the destination.

However, because the concept of a tunnel is a logical construct within the NFM-P, a number of different configurable objects can actually be used as service tunnels, including:

- SDP binding:
  Because SDP bindings can be associated to a service of any type, the service tunnel discovery includes services that are associated to SDP bindings of a given service.
- Service to Tunnel Association:
  Discovery adds tunnels that are associated to a given service.
- G.8031 Ethernet tunnel:
  Discovery adds the global Ethernet Tunnel that the Tunnel Endpoint (SAP's terminating port) is a member of.
- G.8032 Ethernet ring:
  Discovery adds the Ethernet Ring that the Ethernet Element is a member of.
- G.8032 Ethernet radio ring:
  Discovery adds the Ethernet radio ring that the Ethernet Element is a member of. See the *NSP Wavence Device Support Guide* for more information.
- PBB:
  Discovery adds the VPLS or MVPLS services which contain the PBB tunnel (B-Sites) of the given Epipe or VPLS.
- P2MP RSVP LSP:
  Discovery adds the P2MP LSPs retrieved from the MVPN LSP Template associated to PIM sites.
- Optical Transport Service:
  Discovery adds optical services that are traversing network interfaces which are being used by an L2/L3 VPN through any static, dynamic or SR-TE LSPs that are bound to SDPs on the service if the following are true:
  − LSPs are bound to service tunnels (SDPs) on the service

  - One or more of the corresponding hops on the LSP active path share the same port with the optical service termination

- Mobile backhaul (applicable to Wavence SM and 7705 SAR on MPLS-based services running over VLAN):

  Discovery adds Wavence VLL services that are associated to an SDP tunnel or are traversing through network interfaces which are being used by an L2/L3 VPN through any static, dynamic or SR-TE LSPs that are bound to SDPs on the service.

The Service Tunnels tab exists on the applicable service configuration forms. This tab lists the objects considered as service tunnels (such as SDPs, Ethernet rings, Ethernet tunnels, other services) that are currently used by the service you are querying. The associated Discover Service Tunnels button removes any previously discovered service tunnels on the service and initiates a manual rediscovery of the tunnels, based on direct usage and current service configurations.

For composite services, the service tunnel discovery action initiates the discovery of service tunnels on the participating services. The listing includes the service tunnels that are currently used by member services.

The Flow-through Services tab is on the configuration forms of the objects that can be used as a service tunnel. The tab lists the other services that are currently using the object you are querying as a service tunnel. The associated Discover Flow-through Services button removes any previously discovered flow-through services and initiates a manual rediscovery of the tunnels, based on direct usage and current service configurations.

For composite services, the concept of using a service as a service tunnel is not supported. The Flow-through Services tab does not appear on the composite service configuration forms.

For alarm management, an association type of alarm relationship exists between a service and the service tunnels that the service uses. This allows the service tunnel alarms to propagate up to the service level. These association alarms are displayed on the Alarms on Related Objects tab, under the Faults tab of the service. Because this is a one-way association, service alarms are not propagated to the service tunnels that the service is using.

33.2 Tunnel selection profiles

33.2.1 Overview

The NFM-P uses a tunnel selection profile to assign transport tunnels for a service when the service is configured for automatic SDP binding creation. Tunnel selection profiles can also be used for services when the SDP binding creation is performed manually. The NFM-P chooses tunnels and return tunnels when the automatic selection of transport tunnels is enabled and a tunnel selection profile is specified.

A tunnel selection profile contains tunnel administrative groups. When the profile is used within a service to create SDP bindings, any tunnels that include the tunnel administrative group in that profile become eligible for consideration in the tunnel selection process. See Chapter 88, “Tunnel administrative groups” for more information about tunnel administrative groups.

33.2.2 Default tunnel selection profiles

When you enable automatic SDP binding creation on a service, the NFM-P uses the Generated Default TSP for tunnel selection unless you specify a different tunnel selection profile. The

Generated Default TSP is configured to use any transport type, and cannot be modified.

33.3 IP/MPLS service tunnels

33.3.1 Overview

Distributed service traffic is transported between PE NEs by circuits aggregated in unidirectional service tunnels.

The operational theory of a service tunnel is that the encapsulation of the data between the two managed edge NEs appears as a Layer 2 path to the service data although it is really traversing an IP or IP/MPLS core.

Tunnel policies from older releases can be converted to the newer template-based policies by using a “Convert” button located on the rule configuration form. You cannot configure CoS forwarding, LDP-over-RSVP, or multiple LSP bindings on SDP features.

A service tunnel uses GRE or MPLS encapsulation. For MPLS, a mesh of MPLS paths and LSPs must be present in the network core. You associate service tunnels with LSPs during service tunnel configuration. The NFM-P supports LSP creation using a basic LDP variant such as T-LDP, or using LDP over RSVP for tunnel-in-tunnel functionality based on traffic classification.

An NFM-P operator binds a service site to a service tunnel during service configuration. A service tunnel has a unique ID on an NE and cannot be deleted when it is currently associated with a service.

ACL IP and ACL MAC policy filters for service tunnels contain options for forwarding packets on a specific service tunnel based on matching criteria, as described in Chapter 51, “Filter policies” .

33.3.2 IP/MPLS service tunnels design considerations

Consider the following before you configure IP/MPLS service tunnels:

- Service tunnels are unidirectional, so they are required in both directions between the source NE and the destination NE.
- The tunnel is not specific to one service or one type of service. After a tunnel exists, multiple SDPs can be aggregated over the tunnel. The SDPs can belong to different services and different customers.
- When a tunnel already exists, the NFM-P does not automatically create a new tunnel, even if the only available tunnel is down. This prevents the creation of multiple inoperable tunnels.
- All services that are mapped to a tunnel use the same transport encapsulation type defined for the tunnel.
- Operations on the tunnel affect all the services that are associated with the tunnel. For example, the operational and administrative states of a tunnel control the state of service circuits that are carried on the tunnel. In the case of LSP-based tunnels, an LSP can be replaced in the tunnel without reconfiguring each service bound to the tunnel.
- A service tunnel is locally unique to an NE. The same tunnel ID can appear on other NEs.
- A service tunnel uses the system IP address to identify the far-end edge NE.
- When a newly created service tunnel terminates on a local loopback network interface of the

base routing instance, the topology map shows the tunnel terminating on a managed NE. If the service tunnel destination address is provisioned with an IP address from an IES or VPRN interface, the tunnel is shown as terminating on unmanaged NE.

- A service tunnel can be configured using a tunnel template. A user who is assigned the create tunnel template scope of command role can create a tunnel template. The template manager can also apply the tunnel template to a service tunnel.
- An SDP tunnel template can be applied to an auto tunnel policy.

See the appropriate device documentation for more detailed information about service tunnels.

33.3.3 Class-based forwarding

Packets of the same CoS and service are forwarded over a specific RSVP LSP (or a static LSP), which is part of an SDP that the service (instance) is bound to. The forwarding decision is based on the forwarding class of the packet, as assigned by the ingress QoS policy defined for the SAP.

When implementing class-based forwarding, consider the following:

- Class-based forwarding is typically implemented using the service tunnel configuration forms, as detailed in 33.9 “To create an IP/MPLS service tunnel” (p. 1190) . Class-based forwarding configurations can also be altered after their initial creation by editing the service tunnel configuration form as detailed in 33.24 “To view and manage service tunnels and tunnel elements” (p. 1224) .
- An LSP can carry packets of a forwarding class for one service and at the same time, carry packets of a different forwarding class for another service. In other words, the same LSP can be used by more than one SDP, even where the forwarding class assignments are different.
- Within an SDP, only one LSP can support a specific CoS.
- A service instance (or service site) can be bound to SDPs of differing CoS configurations, as well as regular SDPs.
- One LSP in the SDP must be designated as the default LSP. The default LSP is used by the SDP if there is no available LSP that matches the packet's forwarding class. When the default LSP is down, the SDP is also brought down.
- Class-based forwarding can be applied to all services supported by the SRs. For VPLS, you can specify an LSP to forward all multicast/broadcast packets (by default, the default LSP is used). For VLL, shared queuing must be enabled on the ingress SAP to support the class-based forwarding.
- Class-based forwarding is configurable on the 7750 SR, 7450 ESS, and 7950 XRS.

33.4 Ethernet G.8031 tunnels

33.4.1 Overview

The IEEE 802.1ah Provider Backbone Bridging specification employs Provider MSTP to ensure loop avoidance in a resilient native Ethernet core. With P-MSTP, failover times depend largely on the size of the network and the connectivity model used in the network. MPLS tunnels provide core scaling and fast failover times using MPLS FRR. A service based on native Ethernet backbone achieves the same fast failover times as MPLS FRR.

Core Ethernet tunnels compliant with the ITU-T G.8031 specification achieve 50 ms resiliency for backbone failures. A configured Ethernet tunnel can be selected when configuring an L2 access interface on a B-VPLS.

The NFM-P uses two different approaches to configure Ethernet tunnels in the network.

- Approach 1 (recommended):
  The NFM-P provisions an end-to-end Ethernet G.8031 tunnel which reduces configuration errors and aids the diagnosis of any problem in the tunnel. The NFM-P automatically configures the Ethernet tunnel endpoint and path endpoint on each of the participating NEs. See 33.15 “To configure an Ethernet tunnel” (p. 1204) for more information.

  The Ethernet tunnel and path are network wide objects that provide the following benefits:
  - the Ethernet tunnel provides the aggregated State of the Tunnel and reports any inconsistency in the configuration of the endpoints or paths
  - faster and easier provisioning that allows you to configure common properties such as holdTime and revertTime at the network level and then apply them to each tunnel endpoint
  - The NFM-P accelerates the creation of global and local MAs, and MEPs by using a continuity check
- Approach 2:
  The NFM-P provides the ability to provision Ethernet tunnel endpoints and path endpoints separately on each NE. See 33.14 “To configure an Ethernet tunnel endpoint” (p. 1203) for more information on creating Ethernet endpoints. If Ethernet tunnel endpoints are created this way (or discovered from the NE) you can associate them to a network-wide Ethernet tunnel and path, but this must be done manually.

Note: Network-wide Ethernet tunnels and paths are not automatically discovered from the NE.

33.4.2 Transit services

You can use the NFM-P to configure and manage transit services on supporting 7210 SAS and 7705 SAR NEs.

Transit services use default QinQ SAPs (which are also called transit SAPs) to forward traffic for data services. Transit SAPs (notation *.*) receive and forward frames with any encapsulation value, and eliminate the need for customer VLAN-specific SAPs on virtual switching instances for service sites with transit SAPs. Data services use transit services for transparent forwarding, which can reduce configuration time and move traffic more efficiently.

Although transit services function as service tunnels, they are managed in the same way as other VPLS and VLL services in the NFM-P: they have service IDs, appear as services in navigation trees, contain sites and interfaces, and appear on service topology maps. Properties forms for transit services share the same tab structure and functionality as forms for other services in the NFM-P. You can use the Transported Services tab to view and access the data services that use the transit service (you must first discover the services). You can also view and access a transit service that is associated with a VLAN Uplink, using the Transit Services tab on the VLAN Uplink form. You can access VLAN Uplink forms from the service topology map. See 4.1.5 “Service topology maps” (p. 171) in 4.1 “Topology map types” (p. 169) for more information.

Any alarms that are generated on a transit service are propagated to the data service that uses the transit service, and appear as alarms on related objects on the Faults tab of the data service properties form.

You can create transit services automatically for VPLS G.8032 Ethernet rings using the Create Transit Services function. See 33.20 “To configure a transit service on an Ethernet ring” (p. 1220) for more information. Transit services created by this method are new services, and contain only transit SAPs. The NFM-P creates linked transit SAPs on all sites on the ring.

For non-G.8032 VPLS and VLL services, manually-configured transit SAPs on the services are used as a transit service by the NFM-P when the transit SAPs are connected by a physical link.

When you use the Connect to Ethernet Ring function to connect a data service to an Ethernet ring configured with a transit service, the NFM-P creates VLAN-specific SAPs for the data service only on the ring site for which Connect to Ethernet Ring is selected. The transit SAPs on the ring transparently forward traffic through other sites on the service.

You can combine transit services using the composite service functionality in the NFM-P. Composite services are created automatically when Auto Discover Composite Services is enabled in System Preferences, or manually using the Rediscover Composite Services function. See Chapter 85, “Composite service management” for more information about composite services.

For more information about transit services and transit SAPs (default QinQ SAPs), including restrictions, see the appropriate NE documentation.

Ethernet G.8032 rings

33.5 Ethernet G.8032 rings

33.5.1 Overview

Ethernet Ring Protection (ERP) as specified in ITU-T G.8032, is a protection mechanism for Ethernet ring topologies that provides a resilient Ethernet network. ERP provides sub-50ms protection and recovery switching for Ethernet traffic in a ring topology, and, at the same time, ensures that loops are not formed at the Ethernet layer. G.8032v1 supports a single ring topology; G.8032v2 supports multiple rings/ladder topology. For more information about Ethernet G.8032 rings, see the ITU website at http://itu.int.

33.5.2 ERP topology

An ERP topology is a collection of Ethernet ring nodes that forms a closed physical loop. Between two and sixteen Ethernet ring nodes are supported per ERP topology. One ERP topology is supported per NE. Up to two ERP instances are supported per ERP topology. Ethernet ring links connect two adjacent Ethernet ring nodes participating in the same ERP topology.

Loop avoidance in the ring is achieved by guaranteeing that at anytime, traffic may flow on all but one of the ring links. This link is called the ring protection link (RPL). One designated node, the RPL Owner (also referred to as the master node), is responsible to block traffic over the RPL.

33.5.3 Supported ERP service types

You can configure the following service types in an ERP G.8032 topology using the NFM-P on NEs that support ERP:

- VLAN service
  - Wavence SM; see the *NSP Wavence Device Support Guide* for more information
- VLL service (Cpipe and Epipe services only)
- VPLS that support ERP SAPs can connect to other rings and Ethernet services using VPLS SAPs. A VPLS can use a single ring (main or sub-ring) or multiple rings.

33.5.4 Ethernet ring discovery

The NFM-P automatically unites discovered Ethernet ring elements in the network into Ethernet rings, and creates path objects to link endpoints within the ring. In order for the NFM-P to create an Ethernet ring, the following must exist:

- at least two Ethernet ring elements with path endpoints
- a physical link between the endpoints

The NFM-P detects the ring elements and the physical link, then creates an Ethernet ring to contain the ring elements, and path objects to link them. If one of the ring elements is already part of an Ethernet ring, then the other element is added to the existing ring. If both elements are part of existing rings, then the rings are merged (see 33.5.7 “Merging Ethernet rings” (p. 1184) below).

See 33.8 “Workflow to configure Ethernet G.8032 rings” (p. 1189) for information about using the NFM-P to automatically create Ethernet rings.

You can manually create Ethernet G.8032 rings using the NFM-P GUI or using the OSS service provisioning client. See 33.18 “To create an Ethernet G.8032 ring” (p. 1212) to create Ethernet G.8032 rings using the NFM-P GUI. See “MPLS, LSP, and service tunnel configuration” in the *NSP NFM-P XML API Developer Guide* for information about configuring Ethernet G.8032 rings using the OSS client application.

33.5.5 Triggering automatic ring creation

Automatic Ethernet ring creation is triggered by the following events:

- creation or detection of a path endpoint on an Ethernet ring element
- creation or detection of a physical link between two path endpoints on Ethernet ring objects
- performance of a global Ethernet ring audit (see 33.5.11 “Auditing Ethernet rings” (p. 1185))

Path endpoints and physical links can be detected through LLDP or when a node is discovered, or created through NFM-P or CLI. In each case, the NFM-P attempts to create an appropriate Ethernet ring for the discovered object.

33.5.6 Ethernet sub-rings

The NFM-P also detects and creates Ethernet sub-rings, and creates interconnects between upper rings and sub-rings when appropriate. You can view the sub-rings associated with a ring in the Sub-Rings tab of the Ethernet ring configuration form, and the parent upper rings on the Upper Rings tab.

33.5.7 Merging Ethernet rings

When the NFM-P detects a physical link between two Ethernet ring element path endpoints on separate Ethernet rings, the rings are automatically merged together. One of the Ethernet rings is deleted, and the elements and paths from that ring are added to the other ring. The following criteria are used to determine which ring is deleted:

- If one of the rings was created manually using the NFM-P, then the other ring is merged into the manually-created ring.
- If both rings were created manually, then the smaller ring is merged into the larger ring, based on the number of Ethernet ring elements.
- If both rings were created automatically, then the smaller ring is merged into the larger ring, based on the number of Ethernet ring elements.
- If the rings have the same number of elements, and were both created the in the same manner, then the ring with the higher Manager ID is merged into the other ring.

33.5.8 Naming of discovered Ethernet rings

Global Ethernet rings inherit the Description parameter of the first element added to the ring. The Name parameter for automatically created rings is always “N/A”.

33.5.9 Deleting Ethernet rings

Automatically created Ethernet rings are deleted when the underlying physical links are removed. When a physical link is deleted - for example, manually, or by disabling LLDP on the related node - the associated path is also deleted, and the existing Ethernet ring is modified or deleted appropriately.

33.5.10 Adding or removing elements in an Ethernet ring

You can cut-in or cut-out an NE in an existing Ethernet ring as part of a network adjustment. The NFM-P alters the existing Ethernet rings and updates the associated endpoints and services to reflect the new configuration.

To add (cut-in) an NE to an Ethernet ring, perform the following:

1. Shut down ports on NEs that will be connected to the cut-in NE.
2. Configure Ethernet ring endpoints on the cut-in NE. See 4.10 “To create a physical link” (p. 183).
3. Update OAM parameters for NEs to be connected to the cut-in NE. See Chapter 77, “VPLS management ”.
4. Re-cable the network to include the cut-in NE in the Ethernet ring.
5. Enable all related ports.
6. (In-band network only) Configure the cut-in NE as an Ethernet ring element (including control SAPs). See 33.16 “To configure an Ethernet Ring Element” (p. 1208).
7. Discover the cut-in node in the NFM-P. See “Discovering devices using the NFM-P” (p. 277).
8. The NFM-P creates paths and merges them appropriately into the Ethernet rings. (Out-of-band network only) Control SAPs for the cut-in NE are created on the associated control VPLS automatically.
9. You can perform an Ethernet ring audit to verify that the rings have been created properly. See 33.27 “To perform an Ethernet G.8032 ring audit” (p. 1226).
10. Trigger an Update Data Service function to create data SAPs on the cut-in NE and to the associated data VPLS. See 33.21 “To manually update data services on an Ethernet ring” (p. 1221).

To remove (cut-out) an NE from an Ethernet ring, perform the following:

1. Shut down ports that are connected to the cut-out NE.
2. Update OAM parameters for NEs connected to the cut-out NE. See Chapter 77, “VPLS management ”.
3. Re-cable the network to exclude the cut-out NE from the Ethernet ring.
4. Enable all related ports.
5. Unmanage the cut-out NE in the NFM-P. See 9.27 “To manage, suspend, or unmanage a device” (p. 316).

33.5.11 Auditing Ethernet rings

You can use the NFM-P to perform a global or local audit of an Ethernet ring. An audit can find missing elements, paths, and interconnects, and potentially merge rings if a link is found. When an audit discovers a problem that cannot be resolved, an alarm is raised against the Ethernet ring, and the problems encountered by the audit are displayed in the alarm details. Performing another audit clears the alarm, if the problems have been resolved.

33.27 “To perform an Ethernet G.8032 ring audit” (p. 1226) describes how to audit Ethernet rings. Performing a global Ethernet ring audit both creates and merges Ethernet rings, even if automatic Ethernet ring creation is disabled.

33.5.12 Service provisioning on Ethernet G.8032 rings

The NFM-P accelerates the provisioning of VPLS instances on Ethernet G.8032 rings by automatically creating the control services and data services on the rings. Services are configured on Ethernet G.8032 rings in the same method as service tunnels.

You can automatically configure transit services on VPLS G.8032 Ethernet rings. See 33.4.2 “Transit services” (p. 1181) in 33.4 “Ethernet G.8031 tunnels” (p. 1180) and 33.20 “To configure a transit service on an Ethernet ring” (p. 1220) for more information.

You can create Ethernet G.8032 rings using the NFM-P GUI or using the OSS service provisioning client. See 33.18 “To create an Ethernet G.8032 ring” (p. 1212) to create Ethernet G.8032 rings using the NFM-P GUI. See “MPLS, LSP, and service tunnel configuration” in the *NSP NFM-P XML API Developer Guide* for information about configuring Ethernet G.8032 rings using the OSS client application.

33.6 L2TPv3 service tunnels

33.6.1 Overview

L2TPv3 transport tunnels enable the encapsulation of L2 Ethernet frames into IPv6 packets transported by L2TPv3. L2TPv3 tunnels provide services similar to Ethernet pseudowire services but without the use of MPLS encapsulation and labels for network and service offerings that are IPv6-only, with no MPLS or IPv4 services.

L2TPv3-encapsulated traffic is transported over the IPv6 address family, without using any session or label identifier or signaling control plane.

Consider the following when you create an L2TPv3 tunnel:

- you must enable the IPv6 Allowed and Loopback Enabled parameters on the network interface of the NE
- you can configure multiple IPv6 addresses on the loopback interface
- you can add the loopback interfaces as ISIS and OSPFv3 routing instances so that the IGP is used to route the traffic between loopback interfaces
- you can configure an address aggregation policy to distribute the loopback interface IP addresses

Only Epipe spoke and mirror SDP bindings can use L2TPv3 tunnels. See 76.35 “To configure a spoke SDP binding with an L2TPv3 tunnel on a VLL Epipe site” (p. 2165) for more information

about how to create an Epipe SDP binding with an L2TPv3 tunnel. See 93.5 “To create a destination site on a mirror service” (p. 3167) for more information about how to add a remote source mirror SDP binding on a mirror destination site with an L2TPv3 tunnel. See 93.6 “To create a source site on a mirror service” (p. 3169) for more information about how to create a mirror SDP binding on a mirror source site with an L2TPv3 tunnel.

Configuring service tunnel workflows and procedures

33.7 Workflow to configure service tunnels

33.7.1 Stages

1

Before you can create an MPLS service tunnel, you must create a static LSP, dynamic LSP or segment routing TE LSP. See the following procedures:

- 31.10 “To create a static LSP” (p. 1124)
- 31.11 “To create a Dynamic LSP” (p. 1126)
- 31.12 “To create a segment routing TE LSP ” (p. 1130)

2

Create an IP/MPLS service tunnel and assign it to an NE as follows:

a. Create an IP/MPLS service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .

b. Create an IP/MPLS service tunnel using a tunnel template. See 33.12 “To create an SDP using a tunnel template” (p. 1200) .
Note: Service tunnels are unidirectional, so they are required in both directions between the source NE and the destination NE.

c. Create a tunnel administrative group and a tunnel selection profile and assign the tunnel selection profile to an SDP binding. See 88.3 “To create a tunnel administrative group” (p. 2897) and 33.9 “To create an IP/MPLS service tunnel” (p. 1190) and 33.13 “To create a tunnel selection profile” (p. 1201) for more information.

3

Create an L2TPv3 tunnel and assign it to an NE as follows:

a. Create an L2TPv3 tunnel; see 33.11 “To configure an L2TPv3 service tunnel” (p. 1199) .

b. Create an L2TPv3 service tunnel using a tunnel template. See 33.12 “To create an SDP using a tunnel template” (p. 1200) .

Note: Service tunnels are unidirectional, so they are required in both directions between the source NE and the destination NE.

4

Perform one of the following to create an Ethernet service tunnel:

a. Create an Ethernet service tunnel endpoint. See 33.14 “To configure an Ethernet tunnel endpoint” (p. 1203) .

b. Create an Ethernet service tunnel and assign endpoints to the tunnel. See 33.15 “To configure an Ethernet tunnel” (p. 1204) .

5

Create transit services on Ethernet rings as required. See 33.20 “To configure a transit service on an Ethernet ring” (p. 1220) .

6

As required, discover and list service tunnel information. See 33.22 “To discover service tunnels” (p. 1222) and 33.23 “To discover flow-through services” (p. 1223) .

7

Perform one of the following to view service tunnel information:

a. View and manage service tunnels and tunnel elements. See 33.24 “To view and manage service tunnels and tunnel elements” (p. 1224) .

b. View the service tunnel topology. See 33.25 “To view the service tunnel topology” (p. 1224) .

8

As required, run an OAM validation test on a service tunnel. See 33.26 “To run an OAM validation test on a service tunnel” (p. 1225) .

33.8 Workflow to configure Ethernet G.8032 rings

33.8.1 Stages

1

Perform one of the following to create Ethernet ring elements:

a. Create an Ethernet ring element for all supported devices except for OmniSwitch and Wavence devices. See 33.16 “To configure an Ethernet Ring Element” (p. 1208) for more information.

b. Create an Ethernet ring element for an OmniSwitch. See 33.17 “To configure an OmniSwitch Ethernet Ring Element” (p. 1209) .

c. Create an Ethernet radio ring element for a Wavence SM. See the *NSP Wavence Device Support Guide* for more information.

2

Perform one of the following to create physical links:

a. Configure LLDP to automatically discover and create physical links. See 12.21 “To enable LLDP on an NE” (p. 358) .

b. Manually create a physical link. See 4.10 “To create a physical link” (p. 183) .

3

If required, perform a global Ethernet ring audit to trigger Ethernet ring creation and mergers. See 33.27 “To perform an Ethernet G.8032 ring audit” (p. 1226) .

33.9 To create an IP/MPLS service tunnel

33.9.1 Purpose

Use this procedure to create an IP/MPLS service tunnel. Before you can create an MPLS service tunnel, you must create a static LSP, dynamic LSP or segment routing TE LSP. See the following procedures:

- 31.10 “To create a static LSP” (p. 1124)
- 31.11 “To create a Dynamic LSP” (p. 1126)
- 31.12 “To create a segment routing TE LSP ” (p. 1130)

The steps that display in the wizard will vary based on parameter selection and source NE and release.

33.9.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Create→IP/MPLS Service Tunnel (SDP). The IP/MPLS Service Tunnel (SDP) wizard opens.

3

Configure the required parameters.

If a range policy is applied to a service tunnel, a gray text box appears beside the ID parameter to indicate that a range policy is in effect.

If a format policy is applied to a service tunnel, a drop-down menu appears beside the object field during object creation to indicate that a format policy is in effect. When there is only one matching policy, the drop-down menu is dimmed. When there are multiple matching policies, use the drop-down menu to choose a policy. The sequence of the options in the drop-down menu are ordered by the policy Priority parameter.

4

Click Next. The Pick Source Node step is displayed.

5

Select a source site or enter its IP address.

6

Click Next. The Pick Destination Node step is displayed.

7

To create an SDP tunnel using MPLS-TP:

Use the following steps:

1. Enable the MPLS-TP Far-End parameter.
2. Select an NE on which MPLS-TP is enabled.

8

To create an SDP tunnel without MPLS-TP, enter the IP address of the destination site or select an NE.

9

Click Next. The Specify Transport step is displayed.

10

Perform one of the following:

a. For GRE service tunnels:

   Use the following steps:

   1. Specify GRE for the Underlying Transport parameter.
   2. Configure the Signaling parameter.
   3. Enable SR-ISIS or SR-OSPF as needed.

b. For MPLS service tunnels:

   Use the following steps:

   1. Specify MPLS for the Underlying Transport parameter.

      Specify MPLS for the Underlying Transport parameter. The following considerations apply:

      - When you implement class-based forwarding or specify a transport destination address, you must choose MPLS as the Underlying Transport parameter.
      - When you plan to bind an SR-TE LSP to the service tunnel, you must choose MPLS as the Underlying Transport parameter and not configure any other parameters in the Specify Transport step.
      - When you specify a transport destination address, you cannot implement class-based forwarding.
      - When you configure a service tunnel with a source site that is a 7705 SAR, you can choose IPv4 as the Underlying Transport parameter if the destination site is also a 7705 SAR, a GNE, or an unmanaged IP address.

   2. Configure the required parameters.

The Revert Time (seconds) parameter appears when the Mixed LSP Mode parameter is set to true.

When you implement class-based forwarding, you must use one of the following parameter configurations:

- Set Mixed LSP Mode to true.
- Set Mixed LSP Mode, Enable LDP, and Enable BGP-Tunnel parameters to false.

When you specify a transport destination address, you must use the following parameter configurations:

- Set Enable LDP to true.
- Set Mixed LSP Mode and Enable BGP-Tunnel to false.

You cannot set Enable BGP-Tunnel to true when Enable LDP or Mixed LSP Mode is set to true.

You cannot create a BGP or an LDP tunnel when an LSP exists between the source and destination NEs that you specified in Step 5 and Step 8 respectively.

3. If you set the Enable LDP parameter to true, go to Step 18 .
4. If you set the Enable BGP-Tunnel parameter to true, go to Step 20 .

11

Click Next. The Specify Class Forwarding step is displayed.

12

Configure the parameters.

When the Class Forwarding Capability parameter is set to Off, the existing class-forwarding configurations are removed.

The Administrative State and Enforce Diff-Serv Lsp-Fc Map parameters are configurable when the Class Forwarding Capability parameter is set to On.

13

Click Next. The Associate LSPs step is displayed.

14

Click Add to bind an LSP to the service tunnel. The Bind LSPs to Service Tunnel form opens with the Bind LSP to Service Tunnel Stage step displayed.

If no eligible LSP is available for binding to the service tunnel, the form provides actions buttons to:

- Create Dynamic LSP—for details, see 31.11 “To create a Dynamic LSP” (p. 1126)
- Create Static LSP—for details, see 31.10 “To create a static LSP” (p. 1124)
- Create SR-TE LSP—for details, see 31.12 “To create a segment routing TE LSP ” (p. 1130)

15

Choose an LSP and click Finish.

16

Click Close. The Bind LSPs to Service Tunnel form closes and the LSP is listed on the Create Service Tunnel (SDP) form.

17

If you are enabling class-based forwarding on the LSP:

Use the following steps:

1. Choose an LSP and click Set As Default LSP. This is mandatory. The Default LSP Name appears in the associated field.
2. For a VPLS, specify an LSP to forward all multicast/broadcast packets. Otherwise, the default LSP is used. If the LSP is specified, the Multicast LSP Name appears in the associated field.
3. For any LSPs used for class-based forwarding (including the default), choose the LSP and click Choose a Forwarding Class. Choose the required forwarding class from the contextual menu, which then appears in the Forwarding Class to LSP Mappings table.

18

Click Next. The Specify Transport Destination Address step is displayed.

Note: This step appears if you set the Underlying Transport parameter to MPLS in Step 10 and set the Enable LDP parameter to true.

19

Configure the Transport Destination Address parameter.

Note: You cannot assign tunnel administrative groups to the SDP if you configure the Transport Destination Address parameter with an IPv6 address.

20

Click Next. The Specify Local End IP Address step is displayed, if applicable.

21

Select a virtual router or enter its IP address.

22

Click Next. The Specify Hello Parameters step is displayed.

23

Configure the required parameters.

24

Click Next. The Specify MTU Values step is displayed.

25

Configure the required parameters.

26

Click Next. The Specify Metric step is displayed, if applicable.

27

Configure the Metric parameter, if applicable.

28

Click Next. The VC Type Related Parameters step is displayed, if applicable.

29

Configure the VLAN VC Ethertype parameters.

30

Click Next. The Specify Initial State step is displayed.

31

Configure the Administrative parameter.

32

Click Next. The Booking Factor step is displayed, if applicable.

Note: This step appears if you set the Underlying Transport parameter to MPLS in Step 10 and set the Ldp Enabled parameter to false.

33

Configure the SDP Bandwidth Booking Factor (%) parameter, if required.

34

Click Next. The Associate Service step is displayed.

35

Select a service.

36

Click Next. The Tunnel Admin Groups step is displayed.

37

Select the tunnel administrative groups in the Unassigned list that you want to assign to the tunnel and click the right arrow to move them into the Assigned list.

Note: You cannot assign tunnel administrative groups to the SDP if you configured the Transport Destination Address parameter with an IPv6 address in Step 19 .

38

Click Next. The Network Domain step is displayed.

39

Select a network domain.

Note: All SDPs are associated with a network domain. This association is done during the creation or when editing an SDP in the Tunnel (Edit) form. You can also modify the network domain in the Tunnel (Edit) form. See Step 58 . The Egress Interfaces Consistency State status specifies if all the interface associations to SDP belong to a particular domain.

40

Perform one of the following:

a. If you specified MPLS for the Underlying Transport parameter and TLDP or None for the Signaling parameter in Step 10 , click Next and go to Step 41 .

b. If you specified MPLS for the Underlying Transport parameter, and BGP for the Signaling parameter in Step 10 , go to Step 43 .

c. If you specified GRE for the Underlying Transport parameter in Step 10 , click Next and go to Step 54 .

41

The Port Binding step is displayed.

42

Select an Ethernet or HSMDA port, or a LAG.

Only ports in hybrid mode or a LAG that contains hybrid ports can be selected.

43

Click Next. The Configure Source B-MAC LSB step is displayed.

44

Configure the required parameters.

45

Click Next. The Weighted ECMP step is displayed, if applicable.

46

Configure the parameter if required.

47

Click Next. The Configure PW Ports step form is displayed, if applicable.

48

Click Create. The PW Port Binding (Create) form opens.

49

Perform one of the following:

a. Select an existing PW port in the PW Port ID panel.

b. Create a PW port:

Use the following steps:

1. Click Select in the PW Port ID panel and click Create. The PW Port (Create) form opens.
2. Configure the parameters.
3. Click on Node Redundancy to configure node redundancy for the PW port. The Manage Node Redundancy form appears, preconfigured for MC peer groups.

   The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add PW port sync tags.
4. Search for an existing MC peer group or click Create and configure a new MC peer group, as described in 40.4 “To configure an MC peer group” (p. 1330) .
5. Save the changes and close the form.
6. Choose the created PW port and click OK. The Select PW Port ID - PW Port Binding form closes.

50

Configure the required general parameters.

51

Configure the required parameters in the Encapsulation panel.

52

Select an egress shaper virtual port in the QoS panel.

53

Save the changes and close the form.

The PW port can be used in the Service Access Points of the following service types:

- VPRN (see 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249))
- IES (see 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249))
- VPLS Capture (see 74.26 “To configure a capture SAP” (p. 2049))

See the specific service chapter and forms when creating the required service access points. Specify the PW port as the Terminating Port when you create these service access points.

The PW port can be specified as the terminating port for L3 access interfaces of the following service types:

- VPRN (see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656))
- IES (see 78.28 “To configure an L3 access interface on an IES site” (p. 2472))

The PW port can be specified as the terminating port for L2 access interfaces of the VLL Epipe services. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) .

The service access point and L2 or L3 access interfaces you create appears in the lists on the Service Access Points and L2 Access Interfaces or L3 Access Interfaces tabs of the PW Port (Edit) form.

54

Configure the Allow Fragmentation parameter if required.

55

Click Finish to save the configuration.

56

To view the service tunnel configuration after closing the form, enable the View the newly created tunnel parameter.

57

Close the form.

58

If the View the newly created tunnel parameter in Step 56 is enabled, the Tunnel (Edit) form opens with the newly-created service tunnel configuration displayed.

Use the following steps:

1. View or modify the configuration, if required.
2. Click on the Create a template button, if required.
3. Close the Tunnel (Edit) form.

59

Close the Manage Service Tunnels form.

END OF STEPS

33.10 To configure a service tunnel

33.10.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Select a service tunnel and click Properties. The Tunnel (Edit) form opens.

Note: The parameters and tabs that appear depend on the underlying transport of the service tunnel and the NEs on which the service tunnel is configured.

3

Configure the required parameters.

You can only configure the Allow Fragmentation parameter on a GRE tunnel.

4

Review the parameters in the Segment Routing panel. The values of the parameters in the Segment Routing panel cannot be modified. Only one check box (SR-ISIS, SR-OSPF or SR-TE-LSP) can be selected, indicating the type of LSP with which the service tunnel is associated.

5

Click on the Maintenance tab and configure the required parameters.

6

To configure an accounting policy, click on the Accounting tab and select an accounting policy.

7

To configure an associated service, click on the Associated Service tab and select an associated service.

8

To configure a port binding, click on the Port Binding tab.

1. Select a port in the Port panel.
2. Click Create in the PW Port Bindings panel. The PW Port Binding (Create) form opens.
3. Select an existing PW port in the PW Port ID panel or click Create. The PW Port (Create|Edit) form opens. Configure the required parameters and save your changes.

   See Step 49 of 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for more information about how to create a PW port and to configure node redundancy for the PW port.

4. Configure the required parameters.
5. To configure an operational group to be monitored on the PW port binding of the service tunnel, select a monitored group in the Operational Group panel.

   Note:

   The operational state of the PW port binding follows the operational state of the operational group. You can view the Operational Group Down state cause indicator on the PW Port Binding (Edit) form.
6. To configure an HS override on the interface, select a secondary shaper in the PW SAP Secondary Shaper panel.
7. Save your changes and close the form.

9

To configure LSPs for an MPLS service tunnel, click on the LSPs tab.

1. Click Create. The Bind LSPs to Service Tunnel (SDP) step form opens.
2. Select an existing LSP.

   If no eligible LSP is available for binding to the service tunnel, the form provides actions buttons to:
   - Create Dynamic LSP—for details, see 31.11 “To create a Dynamic LSP” (p. 1126)
   - Create Static LSP—for details, see 31.10 “To create a static LSP” (p. 1124)
   - Create SR-TE LSP—for details, see 31.12 “To create a segment routing TE LSP ” (p. 1130)
3. Click Finish.

10

Save your changes and close the forms.

END OF STEPS

33.11 To configure an L2TPv3 service tunnel

33.11.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Create→L2TPv3 Tunnel (SDP) or select an existing L2TPv3 tunnel entry and click Properties. The L2TPv3 Tunnel (SDP) (Create) form opens.

3

Configure the parameters in the Identity panel.

If a range policy is applied to a service tunnel, a gray text box appears beside the ID parameter to indicate that a range policy is in effect.

If a format policy is applied to a service tunnel, a drop-down menu appears beside the object field during object creation to indicate that a format policy is in effect. When there is only one matching policy, the drop-down menu is dimmed. When there are multiple matching policies, use the drop-down menu to choose a policy. The sequence of the options in the drop-down menu are ordered by the policy Priority parameter.

4

Configure the parameters in the Source panel.

Use the following steps:

1. Select a source site.
2. Select a local end IP address.

   The local end IP address must be configured under a loopback network interface on the NE. The IPv6 Allowed and Loopback Enabled parameters must be enabled on the network interface of the NE.

5

Configure the parameters in the Destination panel.

Use the following steps:

1. Select a destination site.
2. Select a far end IP address.

   The destination end IP address must be configured under a loopback network interface on the NE. The IPv6 Allowed and Loopback Enabled parameters must be enabled on the network interface of the NE.

6

Configure the remaining parameters, as required.

7

Save your changes and close the form.

END OF STEPS

33.12 To create an SDP using a tunnel template

33.12.1 Purpose

Use this procedure to create an SDP using a tunnel template. Before you can create an SDP from a tunnel template, you must create the SDP tunnel template.

33.12.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels search form opens.

2

Click Create from Template. The Create Service Tunnel From Template form opens.

3

Select a tunnel template and click OK. The Create Tunnel from Template form opens.

4

Configure the required parameters.

5

Click on the SDP Bandwidth tab and configure the SDP Bandwidth Booking Factor parameter.

6

Click on the Maintenance tab and configure the required parameters.

7

Click on the Accounting tab and select an accounting policy for the service tunnel.

8

Save your changes and close the forms.

END OF STEPS

33.13 To create a tunnel selection profile

33.13.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Create→Tunnel Selection Profile or select an existing tunnel selection profile and click Properties. The Tunnel Selection Profile (Create | Edit) form opens.

3

Configure the required parameters.

As part of the configuration of a tunnel selection profile, you need to select a transport preference. The Transport Preference panel lists all the transport types in static priority order, followed by the PBB Only transport preference option. When configuring your transport preferences, take into account the following considerations:

- You can select either transport types or the PBB Only option. If you select a transport type, the PBB Only option is disabled. If you select the PBB Only option, the transport type options are disabled.
- You need to select one transport preference to be able to create a tunnel selection profile.
- The Allow Auto Tunnel Creation option becomes active only if you selected at least one transport type that supports the automatic tunnel creation. The automatic tunnel creation depends on the support of the selected transport type by both NEs.

Note: On selection of Mixed LSP mode transport type, if multiple tunnels with mixed-mode are operationally up and available, and if Auto SDP Binding is enabled in the service, then the SDP binding association takes following precedence: RSVP, LDP, BGP. This is irrespective of LSP type selected in Tunnel Selection Profile.

When determining what tunnel to use for binding, apply the following guidelines:

- If tunnels that meet the constraints are operationally up and available, then select them in static priority order.
- If no tunnels that meet the constraints are operationally up, then select available tunnels that are operationally down.
- If there are no tunnels, then create tunnels following the static priority order.

4

Select an SDP binding template for the profile, if required.

Note: An SDP binding template allows the NFM-P to use the configuration information contained in the template for the automatic creation of SDP bindings.

5

Click on the Tunnel Admin Groups tab to assign included and excluded admin groups to the profile.

6

In either of Included- or Excluded-Unassigned panels, click on the admin group that you want to include in the profile, and click on the Right arrow to move the selected admin group to the Assigned panel.

You can select multiple groups by holding down the Ctrl key and clicking on the groups.

If required, you can move an assigned admin group to the unassigned portion of the form by selecting the group and clicking on the Left arrow.

7

Save your changes and close the forms.

END OF STEPS

33.14 To configure an Ethernet tunnel endpoint

33.14.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Perform one of the following:

a. Click Create→Ethernet Tunnel Endpoint. The Select Network Elements form opens.

   Select a site on which to create an Ethernet tunnel endpoint and click OK. The Ethernet Tunnel Endpoint (Create) form opens.

b. Select an existing tunnel endpoint and click Properties. The Ethernet Tunnel Endpoint (Edit) form opens.

3

Configure the required parameters.

The Tunnel Endpoint ID and Site ID parameters are only configurable when you are creating a new Ethernet Tunnel Endpoint.

The Access Adapt QoS, Enable Per Forwarding Path Ingress Queue, and Operational Path Endpoint Threshold parameters are only displayed when the Protection Type parameter is set to Load Sharing.

4

Click on the Path Endpoints tab to configure path endpoints.

Use the following steps:

1. Click Create or select an existing path endpoint and click Properties. The Ethernet Tunnel Path Endpoint (Create|Edit) form opens.
2. Configure the required parameters.

   The Path ID parameter is only configurable when you are creating a new Ethernet tunnel path endpoint.

   The APS Command parameter is only displayed and configurable when you are configuring an existing Ethernet tunnel path endpoint and when the Protection Type parameter you configured in Step 3 is set to G8031 1:1. The associated Perform APS Command button is used to perform the selected command on the path endpoint.

3. Click on the Properties button in the Ethernet Path panel if you need to view or alter parameters on the Ethernet Path or configure a CFM Test. The Ethernet Path - Properties button is only displayed and selectable when you are configuring an existing Ethernet Tunnel Path Endpoint.
4. Click on the Port tab to configure a terminating port for a new Ethernet tunnel path endpoint.
5. Select a terminating port.
6. Configure the Control Tag (Outer Encapsulation Value) and Control Tag (Inner Encapsulation Value) parameters.
7. Save your changes and close the form.

5

Click on the MEPs tab to configure a MEP for a new Ethernet tunnel path endpoint.

Use the following steps:

1. Click Create. The MEP (Create) form opens.
2. Configure the required parameters.

   The Control MEP parameter is only displayed if the Interface Type parameter is set to Ethernet Tunnel Path Endpoint.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

   When configuring a Down MEP on a Subscriber Group Interface SAP, the Direction parameter cannot be configured.
3. If the Maintenance Domain for this MEP has a Name Type of none and its Maintenance Association has a Name Format of icc-based, then the Y.1731 TEST and the AIS tabs will also be displayed. See 74.25 “To configure a MEP on a SAP” (p. 2047) for more information.
4. Save your changes and close the form.

6

Save your changes and close the form.

END OF STEPS

33.15 To configure an Ethernet tunnel

33.15.1 Purpose

Use this procedure to configure an Ethernet tunnel. Before you can create an Ethernet tunnel, you must configure an Ethernet tunnel endpoint. See 33.14 “To configure an Ethernet tunnel endpoint” (p. 1203) .

33.15.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Create→Ethernet Tunnel or select an existing Ethernet tunnel and click Properties. The Ethernet Tunnel (Create|Edit) form opens.

3

Configure the required parameters.

The Access Adapt QoS, Enable Per Forwarding Path Ingress Queue, and Operational Path Endpoint Threshold parameters are only displayed when the Protection Type parameter is set to Load Sharing.

4

Automatically configure a CFM MEG on an Ethernet tunnel path and MEPs for each Ethernet tunnel path endpoint.

Perform the following steps:

1. Configure the parameters in the Path CFM Defaults panel.
2. Select a maintenance domain.

   Note:

   The configuration performed here when creating an Ethernet tunnel object cascades to each Ethernet tunnel path that is subsequently configured.

5

Click on the Apply button. The Ethernet Tunnel form refreshes with additional tabs displayed.

6

Click on the Components tab.

7

Right-click on the Tunnel Endpoints object and choose one of the following:

a. Create Ethernet Tunnel Endpoint.

   Perform the following steps:

   1. See 33.14 “To configure an Ethernet tunnel endpoint” (p. 1203) to create an Ethernet Tunnel Endpoint. The endpoint will be listed as Endpoint A on the Ethernet Tunnel form in the Tunnel Endpoints Section.

      Note:

The parameter values specified in Step 3 are automatically applied in the Ethernet Endpoint form when creating an endpoint.

2. Repeat 1 to create Endpoint B.

b. Add Existing Endpoint. The Select Endpoints form opens.

Perform the following steps:

1. Select an endpoint and close the form.

   The endpoint is listed as Endpoint A on the Ethernet Tunnel form in the Tunnel Endpoints Section.

2. Repeat 1 to select Endpoint B.

   Note:

   The parameter values specified in Step 3 are not automatically applied to the selected Ethernet endpoint.

8

Right click on the Paths object and choose Create Ethernet Path. The Ethernet Path (Create) form opens.

9

Configure the required parameters.

10

Click on the Endpoints tab. Perform one of the following:

a. Manually configure the required parameters for endpoints A and B.

A path endpoint cannot be administratively enabled if one of the following is true:

- An operationally up same-fate SAP on the Ethernet tunnel endpoint does not have a tag configured for the path endpoint.
- The Control Tag (Outer Encapsulation Value), Control Tag (Inner Encapsulation Value), and Member Port parameters are not configured.

b. Select existing path endpoints for endpoints A and B. The endpoint parameters are automatically populated.

11

Click on the Intermediate Services tab to view and specify the intermediate services for the Ethernet path. The services that you add to the list (or that already appear there) are in the order of proceeding from Endpoint A to Endpoint B. Perform the following to add an intermediate service:

Perform the following steps:

1. Click Create. The Select Service form opens.
2. Click on the Search button to display a list of intermediate services available to include in this path. Only VLL services can be included.

3. Select the required service and click OK.

   The NFM-P does not perform a validation to confirm whether the selected service is actually along the path between the two path endpoints.

   If you have an intermediate service highlighted in the list and want to add another service immediately above it, click on the Insert Component button rather than the Create button. The Create button places an additional service at the bottom of the list, whereas the Insert Component button places it directly above a highlighted service.

   You can change the ordering of intermediate services in the list by selecting a particular service and using the Move Up or Move Down buttons.

12

Click on the CFM Continuity Check tab to configure a CFM Continuity Check test for each path.

You must complete the following steps before selecting a test:

- Create a maintenance domain with the Name Type parameter set to None. See Chapter 91, “Ethernet CFM” for more information.
- Create a CFM continuity check test. See 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122) for information.

13

In the CFM Test panel, select a CFM test.

14

Perform one of the following:

a. Perform one of the MEP auto creation functions:
   - Click Run Continuity Check Protocol to create maintenance associations and local and remote MEPs. Any existing remote MEPs that do not match the local MEPs are deleted. MEPs are turned up and CCM messages are enabled, and the control MEP property is set on the MEPs.
   - Click Create MEPs to create local MEPs. Any existing remote MEPs that do not match the local MEPs are deleted.

b. Continue to Step 15 . Perform the Run Continuity Check Protocol or Create MEPs function at a later date.

15

If required, click on the MEP tab to display the MEP form and select the check box to configure the Control MEP parameter.

16

Save your changes and close the form.

17

If the Protection Type parameter was set to G8031 1:1 in Step 3 , repeat Step 8 to Step 16 for the other path.

If the Protection Type parameter was set to Load Sharing in Step 3 , repeat Step 8 to Step 16 for all paths.

18

Click Apply. The Path Id, Operational State, and CC Protocol State parameters are updated and the path endpoints with this specific configuration are sent to the NEs.

The CC protocol state appears under the global path (in the navigation tree and on the CFM Continuity Check form). The CC protocol state displays the state of the continuity check protocol running between the MEPs of the two endpoints of the path.

19

Save your changes and close the form.

END OF STEPS

33.16 To configure an Ethernet Ring Element

33.16.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Perform one of the following:

a. Click Create→Ethernet Ring Element. The Select Network Elements form opens.

   Select a site for the Ethernet ring element and click OK. The Ethernet Ring Element (Create) form opens.

   You can create multiple ring elements by selecting multiple sites.

b. Select an existing ring element and click Properties. The Ethernet Ring Element (Edit) form opens.

3

Configure the required parameters.

The ID and Site ID parameters are only configurable when you are creating a new Ethernet ring element.

4

Configure the parameters for a sub-ring, if required.

5

Click on the Path Endpoints tab to configure the path endpoints for the Ethernet ring element.

6

Click Create or select an existing path endpoint and click Properties. The Ethernet Ring Path Endpoint (Create|Edit) form opens.

7

Configure the required parameters.

The Path ID parameter is only configurable when you are creating a new Ethernet Ring Path Endpoint.

8

Click on the Port tab.

Perform the following steps:

1. Select a terminating port for the ring path endpoint.
2. Configure the R-APS Tag (Outer Encapsulation Value) and R-APS Tag (Inner Encapsulation Value) parameters.

9

Save your changes and close the forms.

END OF STEPS

33.17 To configure an OmniSwitch Ethernet Ring Element

33.17.1 Important information

Configuration requirements for OmniSwitch Ethernet ring elements vary depending on the chassis type and release, and whether the service is a stacked VLAN or standard VLAN. See the NE documentation for information about the configuration requirements for a specific chassis type and release.

For ERP created for stacked VLAN services, only NNIs (network ports) can be added as ERP path endpoints. You must use a port or LAG in network mode as the endpoint. To configure a port or LAG in network mode, see 16.29 “To change the port mode” (p. 611).

For stacked VLAN services on some chassis types and releases, the network port must be configured as an ERP binding port, with the Bind Type parameter set to ERP, before configuration as an ERP ring element. However, for some chassis types and releases, a network port or LAG does not have to be configured as an ERP binding port before it is added to the ring.

33.17.2 Steps

1

Perform this step if the chassis type you are configuring requires preliminary configuration of an ERP binding on the network port.

Configure the required network ports with an ERP binding. For each port:

1. Open the port properties form; see 16.24 “To configure Ethernet ports” (p. 599).
2. Click on the Network Interfaces tab, then on the L2 Network Interfaces subtab.
3. Select the required interface in the list and click Properties. The L2 Network Interface (Edit) form opens.
4. Click on the VLAN tab.
5. If the VLAN association is the same as the one in which the ERP binding is required, delete it.
6. Click Create. The VLAN Bindings (Create|Edit) form opens.
7. Select a VLAN, and set the Bind Type parameter to ERP.
8. Save your changes and close the forms.

2

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

3

Perform one of the following:

a. Click Create→Ethernet Ring Element. The Select Network Elements form opens.

   Select an OmniSwitch site for the Ethernet ring element and click OK. The Ethernet Ring Element (Create) form opens.

   You can create multiple ring elements by selecting multiple sites.

b. Select an existing ring element on an OmniSwitch site and click Properties. The Ethernet Ring Element (Edit) form opens.

4

Configure the required parameters.

The ID and Site ID parameters are only configurable when you are creating a new Ethernet ring element.

5

Click on the Path Endpoints tab to configure the path endpoints for the Ethernet ring element.

6

Click Create or select an existing path endpoint and click Properties. The Ethernet Ring Path Endpoint (Create|Edit) form opens.

7

Configure the required parameters.

The Path ID parameter is only configurable when you are creating a new Ethernet ring path endpoint.

8

Click on the Port tab.

9

Select a terminating port for the Ethernet ring path endpoint.

If you are configuring a standard VLAN:

- For some chassis types and releases, the port must be added as a tagged access interface to the corresponding VLAN before being selected as a member port for the ERP.
- For some chassis types and releases, when a port is added in a ring, it is automatically added to the corresponding VLAN as a tagged port.

See the NE documentation for information about the configuration requirements for a particular chassis type and release.

You must select a ring element for each site that will be part of the Ethernet ring.

The Ethernet ring can only have one element configured as an RPL owner or None. This is configured using the Ring Protection Link Type parameter for the element.

The Ethernet ring path can be a LAG or port.

10

Configure the Path ID parameter.

11

Save your changes and close the form.

12

Repeat Step 6 to Step 11 to create another Ethernet ring path endpoint, if required.

If only one Ethernet ring path endpoint is created, the Ethernet ring element is created as a sub-ring.

13

Save your changes in the Ethernet Ring Element form.

14

Perform one of the following:

a. If you are configuring an Ethernet Ring Element on an OmniSwitch NE that supports G.8032v1, continue to Step 15 .

b. If you are configuring an Ethernet Ring Element on an OmniSwitch NE that supports G.8032v2, go to Step 16 .

15

Click on the Protected Services tab to view and specify the services for the Ethernet path. The services that you add to the list (or that already appear there) are in the order of proceeding from Endpoint A to Endpoint B. Perform the following to add a standard or stacked VLAN service:

1. Click Add. The Select Service form opens.
2. Select the required service and click OK.

   The NFM-P does not perform a validation to confirm whether the selected service is actually along the path between the two path endpoints.

   If ERP is created with standard VLAN service, then standard VLAN services for Release 6.4.4/6.6.2 NEs and IPM Enterprise VLAN services for Release 6.6.2 NEs can be added as Protected services.

   If ERP is created with stacked VLAN service, then stacked/IPM stacked VLAN services for 6.4.4/6.6.2 NEs can be added as Protected Services.

16

To edit or view protected services, click on the Properties button of a selected standard VLAN service. The Ethernet Ring Element (Edit) form opens.

17

Click on the Protected Services tab. Select a service and click Properties. The Protected Services (Edit) form opens. Click on the Properties tab to edit or view the protected service.

18

Save your changes and close the form.

END OF STEPS

33.18 To create an Ethernet G.8032 ring

33.18.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels search form opens.

2

Click Create→Ethernet Ring. The Ethernet Ring (Create) form opens.

3

Configure the required parameters.

If you enable the Auto Create Global MEG parameter, the ETH-CFM tab displays all associated Global MEGs created for this Ethernet ring.

If path discovery generates Ethernet ring paths, then the Auto Create Global MEG, Maintenance Domain ID, CCM Interval, and Run CCM on Create parameters are not propagated to the per-path CFM Continuity Check tab configuration. Auto-generation of a MEG and/or MEPs can be configured on a per-path basis.

Alternatively, if path discovery is not operational (that is, no Ethernet ring paths are populated into the Ethernet ring), then Path CFM Defaults parameters on the General tab are propagated and a MEG and/or MEPs are generated upon manual creation of the Ethernet ring paths.

4

Click Apply. The Ethernet Ring (Edit) form opens.

5

Click on the Components tab.

6

Right-click on Ring Elements and choose one of the following:

a. Create Ethernet Ring Element.

   Perform the following steps:

   1. Create an Ethernet ring element as described in 33.16 “To configure an Ethernet Ring Element” (p. 1208) . The element displays on the Components tab on the Ethernet Ring form.

      You must create a ring element for each site that will be part of the Ethernet ring.

      The Ethernet ring can only have one element configured as an RPL owner and one element configured as an RPL neighbor. This is configured using the element’s Ring Protection Link Type parameter.

      The parameter values specified in Step 3 are initially populated in the Ethernet Ring Element form when creating a ring element from the Components tab. However, you can change these as required.

   2. Repeat 1 to create the other required ring elements.

      You can create all the ring elements required for the ring in one step; see 33.16 “To configure an Ethernet Ring Element” (p. 1208) .

   3. Go to Step 7 .

b. Add Existing Element.

   Perform the following steps:

   1. The Select Element form opens.
   2. Click Search, select an element, and click OK.

      The element displays on the Components tab on the Ethernet Ring form.

You must select a ring element for each site that will be part of the Ethernet ring.

The Ethernet ring can only have one element configured as an RPL owner and one element configured as an RPL neighbor. This is configured using the element's Ring Protection Link Type parameter.

When you add an existing Ethernet endpoint, the parameter values specified in Step 3 are not populated into the properties of that endpoint.

3. Repeat 1 to 2 to select the other required ring elements.

7

Create an Ethernet path.

Perform the following steps:

1. Right-click on Paths and choose Create Ethernet Ring Path. The Ethernet Path (Create) form opens.

   You must create a path for each element in the ring. Each element (site) in the ring will therefore have two endpoints from two different paths associated with it.
2. Configure the required parameters.
3. Click on the Endpoints tab to configure the endpoints for the path.

   One and only one endpoint within the entire Ethernet ring must have its Ring Protection Link Type parameter configured as Ring Protection Link End. All other endpoints in the Ethernet ring must have this parameter set to a value of Normal.
4. Configure the required parameters for endpoints A and B.

   Path endpoints cannot be administratively enabled if:
   - the R-APS Tag or Member Port parameters are not configured
   - a MEP is not configured on the path endpoint
5. Select an available path endpoint on endpoints A and B.

8

Click on the CFM Continuity Check tab to configure a CFM Test for the path.

9

Modify the parameters in the Global MEG Auto Creation panel, if required.

If path discovery generates Ethernet ring paths, then these parameters are not propagated to individual paths. Auto-generation of a MEG and/or MEPs may be configured on a per-path basis.

10

Click Select. The Select CFM Test form opens. Perform one of the following:

a. Select the required CFM test.

b. Click Create. The Global Maintenance Entity Group (Create) form opens.

The Initial CCM Interval for the CFM test you select or create must be set to 10 ms or 100 ms, otherwise the subsequent MEP creation will fail.

See 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122) for information about creating the test.

11

Click on the Intermediate Services tab to view and specify the intermediate services for the Ethernet path. The services in the list are in an order proceeding from endpoint A to endpoint B.

To add an intermediate service:

1. Click Create. The Select Service form opens.
2. Click Search, select the required service, and click OK.

   The Select Service form closes and the service is displayed in the list on the Intermediate Services tab.

   The NFM-P does not perform a validation to confirm whether the selected service is along the path between the two path endpoints.
3. Repeat 1 to 2 to add another intermediate service, if required.

   If you have an intermediate service highlighted in the list and want to add another service immediately above it, click Insert Component rather than Create. Clicking Create places an additional service to the bottom of the list, whereas clicking Insert Component places it directly above a highlighted service.

   You can change the ordering of intermediate services in the list by selecting a particular service and using the Move Up or Move Down buttons.

12

Perform this step if you need to create interconnects for sub-rings. The interconnects connect a sub-ring to the main ring.

Perform the following steps:

1. Right-click on Interconnects and choose Create Ethernet Ring Interconnect. The Ethernet Ring Interconnect (Create) form opens.
2. Configure the required parameters.
3. Configure the Element A end node for the interconnect.

   In the Element A panel, select a site, and configure the required parameters.
4. Configure the Element B end node for the interconnect.

   In the Element B panel, select a site, and configure the required parameters.

13

To automatically create a control service for the Ethernet G.8032 ring, go to Step 14 . To manually create the control service, go to Step 15 .

14

To automatically create a control service for the Ethernet G.8032 ring:

Perform the following steps:

1. Click Create Control Service. The Create Control Service form opens.
2. Select a customer.
3. Configure the required parameters if your Ethernet ring is an Ethernet Virtual link sub-rings.
4. Select a template.

   Only use the Control Service Template if you need to customize the parameters associated with the sites or SAPs such as the description or name.

   If you need to create a Control Service Template, choose Manage→Templates from the NFM-P main menu. Click Browse Example and select the Ethernet ring example.

   If a template is not specified, all created objects use their default property values.
5. Save your changes and close the form.

15

Manually create a control service for the Ethernet ring if you did not auto-create one in Step 14 .

Perform the following steps:

1. Perform 77.5 “To create a VPLS” (p. 2249) to create a control service on a VPLS.
2. Perform 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) to create an L2 access interface for each path endpoint in the Ethernet ring. Each site in the control VPLS must have two control L2 access interfaces.

   When creating the L2 access interfaces, ensure the following requirements:
   - The Terminating Port and Encap Type parameters must have the same values used to create the particular path endpoint as in Step 7 .
   - The Outer Encapsulation Value and Inner Encapsulation Value parameters for the port in each L2 access interface must be set to the same value as the R-APS Tag (Outer Encapsulation Value) and R-APS Tag (Inner Encapsulation Value) parameters, respectively, that you used for the particular path endpoint in Step 7 . This defines the interface as a control SAP for the Ethernet ring.

16

Perform 77.5 “To create a VPLS” (p. 2249) to create data services on the VPLS required for the Ethernet ring. The data service must be a regular VPLS, I-VPLS, or B-VPLS type.

17

To automatically connect the data services to the Ethernet ring, go to Step 18 . To manually connect the data service, go to Step 19 .

18

Automatically connect the data service to the Ethernet ring.

Perform the following steps:

1. Select one or more Site icon(s) associated with the data service and choose Connect to Ethernet Ring from the contextual menu. The Connect to Ethernet Ring form opens.

   When you use the Connect to Ethernet Ring function to connect a data service to an Ethernet ring configured with a transit service, the NFM-P creates VLAN-specific SAPs for the data service only on the ring site for which Connect to Ethernet Ring is selected. The transit SAPs on the ring transparently forward traffic through other sites on the service.
2. Select an Ethernet ring.
3. Configure the required parameters.
4. Select a data site template.

   Only use the data site template if you need to create an I-VPLS or B-VPLS type data service.

   If you need to create a data site template, choose Manage→Templates from the NFM-P main menu. Click Browse Example and select the Ethernet ring example.

   If a template is not specified, all created objects use their default property values.
5. Save your changes close the form.

19

To manually connect the data service to the Ethernet ring, perform 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) to create an L2 access interface for each path endpoint in the Ethernet ring. Each site in the data service must have at least two L2 access interfaces.

When creating the L2 access interfaces, the Terminating Port and Encap Type parameters must have the same values that were used to create the particular path endpoint in Step 7 .

END OF STEPS

33.19 To create an Ethernet G.8032v2 ring on an OmniSwitch

33.19.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels search form opens.

2

Click Create→Ethernet Ring. The Ethernet Ring (Create) form opens.

3

Configure the required parameters.

If you set the Service Composition parameter to VPLS or VPLS and VLAN, configure the following parameters:

- Hold Time Down (centiseconds)
- Hold Time Up (deciseconds)
- Revert Time (seconds)
- Guard Time (deciseconds)
- Compatible Version

4

If the Service Composition parameter was set to VLAN or VPLS and VLAN, perform the following steps to configure control service defaults:

Perform the following steps:

1. Configure the VLAN Application Type parameter.
2. Select a VLAN control service.

5

Click Apply. The Ethernet Ring (Edit) form opens.

6

Click on the Components tab.

7

Right-click on Ring Elements and choose one of the following items:

You must create a ring element for each site in the Ethernet ring, and a path for for each of these ring elements. Each element will therefore have two endpoints from two different paths associated with it.

The Ethernet ring can only have one element configured as an RPL owner and one element configured as an RPL neighbor. This is configured using the element's Ring Protection Link Type parameter.

The parameter values specified in Step 3 are initially populated in the Ethernet Ring Element form when creating a ring element from the Components tab, however, you can change these as required.

Only one endpoint within the entire Ethernet ring must have its Ring Protection Link Type parameter configured as Ring Protection Link End. All other endpoints in the Ethernet ring must have this parameter set to Normal.

a. Create Ethernet Ring Element with path endpoints.

   Perform the following steps:

   1. See 33.17 "To configure an OmniSwitch Ethernet Ring Element" (p. 1209) for information

about creating an Ethernet ring element with path endpoints on OmniSwitch NEs. The element is displayed on the Components tab in Ethernet Ring (Edit) form.

2. Repeat 1 to create more required ring elements with path endpoints.

b. Add an existing Ethernet Ring Element.

Perform the following steps:

1. The Select Element form opens.
2. Click Search and select an element.
3. Repeat 2 to select the other required ring elements.

8

Right-click on Paths and choose Properties. The Ethernet Ring Paths (Edit) form opens.

9

Click on the CFM Continuity Check tab to configure a CFM test for the path.

10

Configure the Auto Create Global MEG parameter.

11

Select a maintenance domain.

12

Configure the CCM Interval and Run CCM On Create parameters.

13

Click Select in the CFM Test panel. The Select CFM Test form opens. Perform one of the following:

a. Select the required CFM test.

b. Click Create. The Global Maintenance Entity Group (Create) form opens.

See 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122) for information about creating the test.

The Initial CCM Interval for the CFM test that you select or create must be set to 10 ms or 100 ms, otherwise the subsequent MEP creation fails.

14

Save your changes and close the forms.

END OF STEPS

33.20 To configure a transit service on an Ethernet ring

33.20.1 Purpose

Perform this procedure to create a transit service on a VPLS G.8032 Ethernet ring. The Ethernet ring must already exist; see 33.4.2 “Transit services” (p. 1181) in 33.4 “Ethernet G.8031 tunnels” (p. 1180) .

Transit SAPs apply only on QinQ ports.

33.20.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Choose Ethernet Ring (Ethernet Ring) and click Search. A list of Ethernet ring services is displayed.

3

Choose an Ethernet ring and click Properties. The Ethernet Ring (Edit) form opens.

4

Click Create Transit Service.

5

Select a customer from the Select Customer form.

6

Configure the required parameters.

The default value for these parameters is *, or 4095. When you create a transit service, you must leave the parameters at the default value.

7

Select a transit service template.

Only use the transit service template if you need to customize the parameters associated with the sites or SAPs, such as the description or name.

If you need to create a transit service template, choose Manage→Templates from the NFM-P main menu. Click Browse Example→Ethernet Ring.

If a template is not specified, all created objects use their default property values.

8

Click OK. A dialog box appears.

9

Perform one of the following:

a. Click Yes to view the properties of the service. The properties form for the transit service opens. View the properties, and close the form. The Ethernet Ring (Edit) form is displayed.

b. Click No. The Create Transit Service form closes and the Ethernet Ring (Edit) form is displayed.

10

Discover the transit service.

Perform the following steps:

1. Click on the Flow-through Services tab. A list of services is displayed.
2. Click Discover Services. The transit service is added to the list.

11

Save your changes and close the form.

You can combine transit services using the composite service functionality in the NFM-P. Composite services are created automatically when Auto Discover Composite Services is enabled in System Preferences, or manually using the Rediscover Composite Services function. See Chapter 85, “Composite service management” for more information about composite services.

END OF STEPS

33.21 To manually update data services on an Ethernet ring

33.21.1 Purpose

Use this procedure update the data VPLS for an Ethernet ring configuration after adding NEs in the Ethernet ring.

33.21.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Choose Ethernet Ring (Ethernet Ring) and click Search. A list of Ethernet ring services is displayed.

3

Choose an Ethernet ring and click Properties. The Ethernet Ring (Edit) form opens.

4

Click on the Transported Services tab.

5

Click Update Data Services.

An information window opens, indicating the number of VPLSs being checked. Click Ok.

6

Close the form.

END OF STEPS

33.22 To discover service tunnels

33.22.1 Purpose

Use this procedure to list all service tunnels objects (SDPs, Ethernet rings, Ethernet tunnels, other services, and so forth) that are currently used by the service you are querying.

33.22.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search, choose a service, and click Properties. The Service (Edit) form opens.

3

Click on the Service Tunnels tab.

4

Click Discover Service Tunnels to re-populate the list.

Any previously-discovered service tunnels on the service are removed and an on-demand re-discovery of these tunnels is triggered.

A confirmation dialog appears when the discovery is complete.

5

Click OK.

6

Click on an item in the list and click Properties to view the object information.

7

Close the form.

END OF STEPS

33.23 To discover flow-through services

33.23.1 Purpose

Use this procedure to list all other services that are currently using the object you are querying as a service tunnel. These are collectively referred to as flow-through services. See the "Service tunnel overview" (p. 1177) section for a list of applicable service tunnel objects.

33.23.2 Steps

1

Open the configuration form of the object that you want to examine.

2

Click on the Flow-through Services tab.

3

Click Discover Flow-through Services to re-populate the list.

Any previously discovered flow-through services are removed and a manual re-discovery of these services is triggered.

A confirmation dialog appears when the discovery is complete.

4

Click OK.

5

Click on a service in the list and click Properties to view the service information.

6

Close the form.

END OF STEPS

33.24 To view and manage service tunnels and tunnel elements

33.24.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Search, choose a service tunnel, and click Properties. The Tunnel (Edit) form opens.

If you searched on Service (or on a specific service type) used as a service tunnel, be aware that the generated list contents are subject to change, depending on current usage and on when the last on-demand (manual) discovery was conducted. There is no auto-discovery for this listing.

In other words, only those services will be listed that were known to be used as service tunnels when the last discovery was conducted. Because the discovery is only updated on user demand, the list of services generated here may not be up to date.

33.22 “To discover service tunnels” (p. 1222) details the on-demand discovery of service tunnels (including a service used by another service as a tunnel).

3

View or configure the parameters, as required.

View the States and State Cause indicators, if applicable, for troubleshooting information, such as a failed OAM diagnostic.

4

Select a tunnel and click Rediscover Destination Site as needed to rediscover the service tunnel destination. If the tunnel selected has its Terminating Site ID set with an interface IP address and a managed node is found that owns that IP address, the rediscovery will update the Terminating Site ID with the node ID.

5

Save the changes and close the forms.

END OF STEPS

33.25 To view the service tunnel topology

33.25.1 Steps

1

Choose Application→Service Tunnel Topology from the NFM-P main menu. The Service Tunnel Topology map opens.

2

View the topology map as required. See Chapter 4, “Topology map management” for information about using topology maps.

3

Close the topology map.

END OF STEPS

33.26 To run an OAM validation test on a service tunnel

33.26.1 Purpose

Use this procedure to run an OAM validation test on a service tunnel. An OAM validator test suite must be created for the tested entity; see Chapter 89, “Service Test Manager” .

Note: Alternatively, you can also run an OAM validation test on the service tunnel by performing a one-time validation; see 90.55 “To run a one-time validation test on a service” (p. 3084) .

33.26.2 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens

2

Click Search, choose a service tunnel, and click Properties. The Tunnel (Edit) form opens.

3

Click Validate.

If an OAM validator test suite is not associated with the service tunnel, a dialog box appears. Click OK, choose an OAM validator test suite, and click OK.

4

View the State Cause indicators. When the validation test fails, a check mark appears beside the OAM Validation Failure indicator.

5

Click on the Tests tab, then on the Tested Entity Result tab.

6

Choose an entry and click Properties. The Tested Entity Result (Edit) form opens.

7

Click on the Results tab to display the validation test results.

8

If you need to compare two test results from the same type of test, choose the two test results and click Compare. The Difference form opens. Otherwise, go to Step 10 .

9

Compare the test results.

10

Close the forms.

END OF STEPS

33.27 To perform an Ethernet G.8032 ring audit

33.27.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Search and choose one of the following:

a. To perform a global Ethernet ring audit, go to Step 3 ,

b. To perform a local Ethernet ring audit, go to Step 6 .

3

Click Global Ring Audit. The Global Ring Audit dialog box appears.

4

Click Yes. The dialog box closes, and the Ring Audit progress window appears. When the audit completes, a notification box appears.

5

Click Yes. Go to Step 8 .

6

Select an Ethernet ring and click Properties. The Ethernet Ring (edit) form opens.

7

Click Ring Audit. When the audit completes, a warning or information window appears displaying the results of the audit.

8

If the audit encountered any problems that could not be resolved, an alarm is raised against the Ethernet ring listing the problems. To clear the alarm, resolve the problems and then repeat the audit.

END OF STEPS

34 IPsec

34.1 Overview

34.1.1 IPsec overview

You can use the NFM-P to configure IPsec sessions and the Security Associations (SAs) that are required in a bidirectional IPsec tunnel. You can configure multiple IPsec tunnels for a VPRN.

NFM-P IPsec session configuration supports the following:

- encryption methods such as DES, 3DES, AES-128, AES-192, and AES-256
- authentication and hashing methods such as HMAC-MD5, HMAC-SHA1, and HMAC-SHA2
- key distribution methods such as IKE shared secret with PFS, and manual exchange
- key generation algorithms such as Diffie-Hellman
- IPsec modes
- shared secret authentication
- NAT traversal for IKEv1 and IKEv2
- DPD for the IPsec tunnel

An IPsec VPN service includes IPsec tunnels that terminate on IES or VPRN IPsec gateways. These gateways support L3 forwarding through an interface that connects to an IPsec tunnel. You can use the NFM-P to configure VPRN services to which individual hosts can connect over the Internet to an IES or VPRN IPsec gateway. You can configure one or more tunnel interfaces in a VPRN service, and can configure multiple tunnel security profiles for each tunnel interface.

IKE policies are used to negotiate IPsec SAs between IPsec peers. An SA is a relationship between two or more IPsec peers that defines how the peers communicate securely. IKE policies are exchanged between IPsec peers to negotiate a secure communication channel; the policies specify how traffic is encrypted between source and destination sites in an IPsec VPN by establishing a shared security policy using authentication keys.

IPsec transform policies specify the protocol for the IPsec authentication header and the encryption protocol for the Encapsulating Security Payload (ESP) and define the attributes that are used to secure the data.

IPsec client databases provide a mechanism to create secure LAN-to-LAN tunnels between an IPsec gateway and multiple VPN clients. The system checks the client database during tunnel authentication and the database returns client credentials, a private VRF ID, private interface name, and other IPsec parameters.

See Chapter 49, “Policies overview” for general information about policies.

After an IPsec peer initiates an IPsec session, there are two main phases:

- authentication and protection of IPsec peer identities and negotiation of matching IKE SA policies between peers to establish a secure channel for negotiating IPsec SAs in the next phase
- IPsec SA parameter negotiation and establishment of matching peer SAs

After the second phase, the IPsec peers exchange data over the IPsec tunnel according to the IPsec parameters in the IKE and IPsec transform policies.

You can create a tunnel template to configure shared IPsec transforms and IKE policies. Each IPsec peer configuration can include the following:

- one or more configured IPsec transforms
- one IKE policy
- one unique IPsec tunnel
- one tunnel filter defined in the IPsec tunnel configuration

Each IPsec tunnel between IKE peers is identified by a unique remote peer IP address or a unique local IP address.

You can use the IPsec Application Function Manager to create and manage end-to-end IPsec components to form a secure VPN.

The NFM-P XML API supports IPsec VPN configuration.

34.1.2 IPv6 IPsec

OSPFv3 authentication requires IPv6. IPv6 IPsec requires the following:

- IPsec transport mode — required because the NE acts as an OSPFv3 authentication host
- IPsec static security association — defines the SPI values, algorithms, protocol, and keys to be used, and requires the same configuration at each end of the tunnel
- AH and ESP
- MD5, SHA1, and SHA2

34.1.3 BFD

You can use BFD for static LAN-to-LAN IPsec tunnels on supporting NEs.

Consider the following when implementing BFD over static LAN-to-LAN IPsec tunnels:

- You can have only one BFD session between a source/destination address pair.
- Each tunnel can be associated with only one BFD session. However, one or more tunnels, to a maximum of 500, can be associated with the same BFD session.
- If one BFD session is associated with multiple tunnels, the tunnel that carries the BFD traffic must be operationally up before any of the other tunnels can be operationally up.
- When the NFM-P does not receive BFD packets from a peer before the detection time expires or a signal down notification is sent from a remote peer, the BFD session is considered down.

  When the NFM-P sets the associated IPsec tunnels in a down state, the NFM-P performs the following:
  - sends a Delete Payload message to each remote peer from each associated tunnel and SA
  - removes the state and table entries from each associated tunnel and SA

34.1.4 Temporary MTU

You can configure an IPSec tunnel to propagate ICMP messages for use in temporary MTU learning by configuring the parameters in the IP Fragmentation, ICMP Message Generation, and ICMPv6 Generation panels of an IPSec Tunnel, IP/GRE Tunnel, or IPSec Tunnel template.

34.2 IPsec VPNs

34.2.1 Overview

You can create and manage the association between IPsec components, and public and private services, to form a secure VPN. See 34.2.3 “Typical applications for IPsec corporate services” (p. 1232) for the typical applications of an IPsec VPN.

You use the IPsec VPN step forms to perform the following:

- Configure the corporate service type.
- Enable the link between the corporate and secure service.
- Choose an NE service site.
- Configure the secure VPRN service.
- Configure the delivery service.
- Choose the tunnel group.
- Create the policy.
- Deploy the IPsec VPN.

34.2.2 Tunnel types

The following table lists the tunnel types for an IPSec VPN in different products.

Table 34-1 Tunnel types for an IPsec VPN

Tunnel type	7705 SAR	7750 SR	7450 ESS in mixed mode	7210 SAS
Dynamic (site-to-site)		✓	✓	
Dynamic (soft client)		✓	✓	
Static	✓	✓	✓	✓

The NFM-P creates the following after the successful creation of an IPsec VPN, regardless of the tunnel type:

- a secure VPRN service and NE sites
- the IES or VPRN delivery service and NE sites
- if you specify that the secure and corporate services are to be linked, a composite service that contains the corporate and secure services

The NFM-P performs specific configuration actions after the successful creation of an IPsec VPN, depending on the tunnel type; see 34.26 “To assign policies and configurations for a dynamic site-to-site IPsec VPN” (p. 1264) to 34.28 “To assign policies and configurations for a static IPsec VPN” (p. 1266) for information.

34.2.3 Typical applications for IPsec corporate services

The following figure shows a public VPRN service that is associated with a private VPRN service. The private service belongs to a larger private VPRN. The public service can be a VPRN or IES service. The private service can only be a VPRN service.

Figure 34-1 Static IPsec tunnel

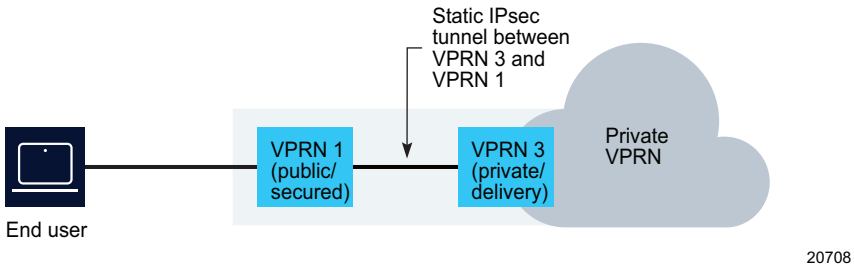

The following figure shows two public L3 VPRNs, VPRN 1 and VPRN 2, which are connected to the private, secure service VPRN 3 through an IPsec gateway. The public services can be VPRN or IES services. The private service can only be a VPRN service.

Figure 34-2 IPsec tunnels for a site

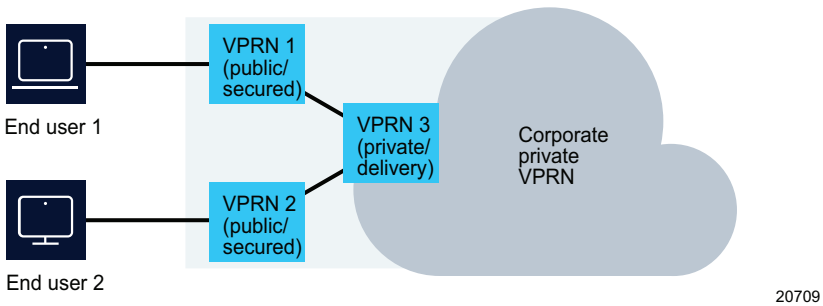

Figure 34-3, “IPsec tunnels for multiple sites” (p. 1233) is the same as Figure 34-1, “Static IPsec tunnel” (p. 1232), but Figure 34-3, “IPsec tunnels for multiple sites” (p. 1233) shows IPsec tunnels across multiple sites. Site A, Site B, and Site C are part of the private service VPRN 4. For the site, there is a secure IPsec tunnel between a public service and a private service. The public services can be L3 VPRN or L3 IES services. The private service can only be a VPRN service.

Figure 34-3 IPsec tunnels for multiple sites

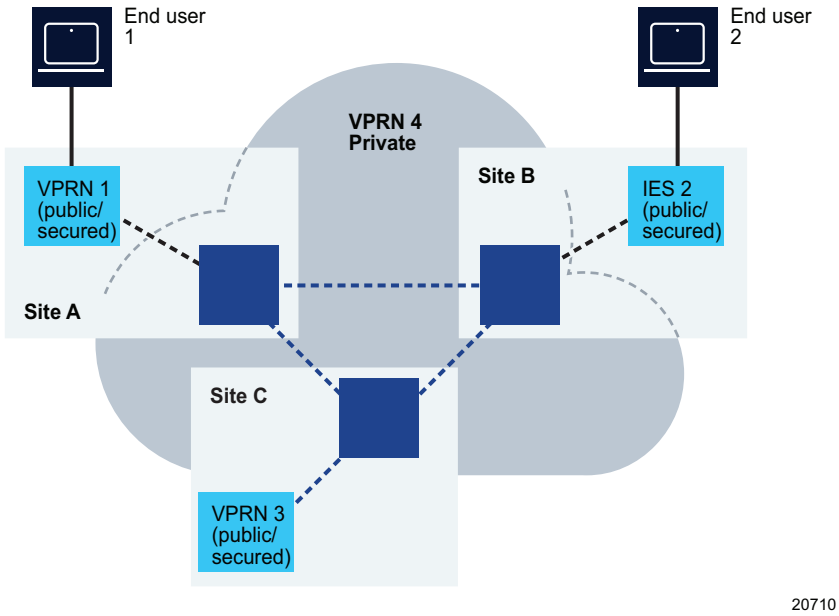

The following figure shows a private VPRN that is connected to a corporate VPLS. The public IES and private VPRN service connect to the VPLS through a CCAG or SCP. A CCAG connects the private VPRN to the VPLS.

Figure 34-4 IPsec VPN in a corporate network

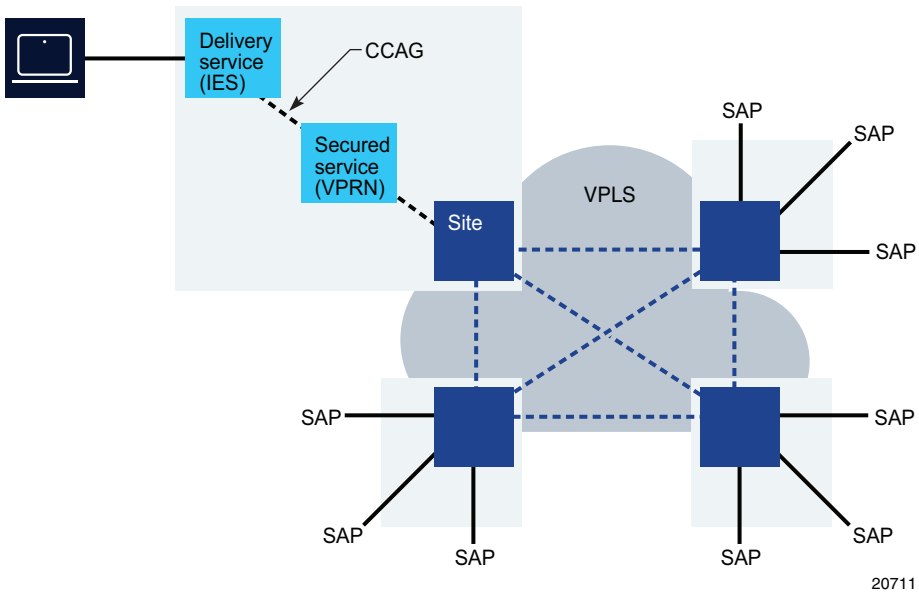

34.3 Multichassis IPsec

34.3.1 Overview

Multichassis (MC) IPsec provides a stateful failover mechanism for IPsec tunnels between two 7750 SRs or 7450 ESSs in mixed mode—an active NE and a standby NE. Stateful failover enables IPsec traffic to continue being forwarded without interruption in the event of a failure. MC IPsec provides protection for NE failure or MS-ISA failure. When an active NE fails, the IPsec tunnels failover to the standby peer without having to re-establish the session. The failover mechanism can occur at the tunnel group level. A tunnel group can failover to the standby NE, independently of other tunnel groups on the active NE; see Chapter 41, “MC IPsec” for more information.

34.4 Sample video wholesale IPsec configuration

34.4.1 Overview

Larger providers or a cooperative of smaller providers often unite to provide a video head end to avoid the costs of investing in satellite head end locations on their own ground station, to provide triple-play features. Each retail subscriber can purchase content from this single station and receive it over IP. However, encryption is required so that the signal cannot be understood if intercepted.

The following figure shows a sample video wholesale configuration with a high-speed encrypted tunnel.

Figure 34-5 Sample video wholesale configuration

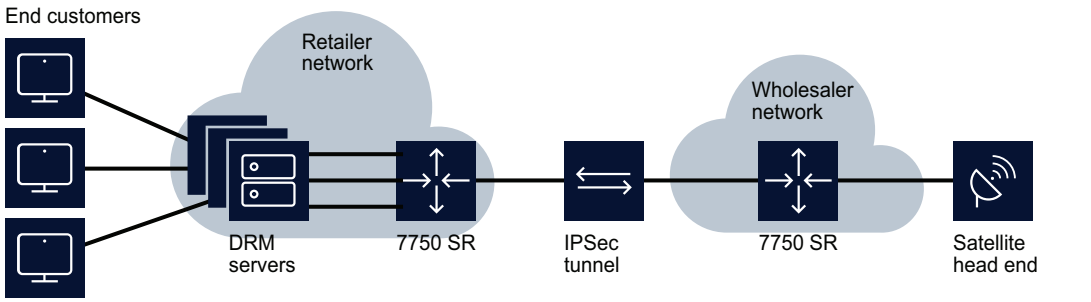

34.5 Workflow to configure IPsec

34.5.1 Stages

1

Provision an ISA tunnel MDA on each participating NE; see Chapter 11, “Working with network objects” and Chapter 13, “Logical group object configuration” for information about IPsec equipment configuration.

Use the following steps:

1. Create or configure ISA-tunnel groups.
2. Assign the active and backup tunnel group members to the ISA-tunnel groups.

2

Configure an IKE policy; see 34.8 “To configure an IPsec IKE policy” (p. 1238).

3

Configure an IPsec transform policy; see 34.10 “To configure an IPsec transform policy” (p. 1240).

4

If your network includes shared IPsec transform and IKE policies, configure an IPsec tunnel template; see 34.12 “To configure an IPsec tunnel template” (p. 1241).

5

Configure an IPsec security policy; see 34.13 “To configure an IPsec security policy” (p. 1242).

6

Configure a RADIUS authentication policy to apply to an IES or VPRN IPsec gateway; see 34.14 “To configure a RADIUS authentication policy” (p. 1243).

7

Configure a RADIUS accounting policy to apply to an IES or VPRN IPsec gateway; see 34.15 “To configure a RADIUS accounting policy” (p. 1244).

8

Configure a trust anchor profile; see 34.17 “To configure a trust anchor profile” (p. 1246).

9

Configure a certificate profile; see 34.18 “To configure a certificate profile” (p. 1247) .

10

If you are configuring IPsec on a VPRN, create the private-facing tunnel interface; see 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249).

Use the following steps:

1. Create a private IPsec SAP.
2. Configure ingress and egress policies.

11

Create the public-facing tunnel interface; see 34.22 “To configure an IES or VPRN IPsec gateway” (p. 1256).

Use the following steps:

1. Create an L3 access interface on an IES or VPRN; see Chapter 78, “IES management” and Chapter 79, “VPRN service management” for information about creating L3 access interfaces.
2. Define the IPsec public SAP for the L3 access interface.
3. Specify the IPsec gateway, if required, on the NE.

12

If you are configuring IPsec on a VPRN, create IPsec tunnels on the VPRN tunnel interface; see 34.21 “To configure an IPsec tunnel on a VPRN tunnel interface” (p. 1253).

13

Configure the static route; see 27.13 “To configure a static route on a routing instance” (p. 852).

34.6 Workflow to configure IPsec VPNs

34.6.1 Stages

1

Create a corporate service; see the appropriate service chapter.

2

Create an IPsec VPN; see 34.25 “To configure an IPsec VPN” (p. 1262).

3

Select the NE service sites for the IPsec VPN; see 34.25 “To configure an IPsec VPN” (p. 1262).

4

Create or select the secure VPRN service for the IPsec VPN; see 34.25 “To configure an IPsec VPN” (p. 1262).

5

Create or select the delivery service for the IPsec VPN; see 34.25 “To configure an IPsec VPN” (p. 1262).

6

Select the tunnel group for the IPsec VPN; see 34.25 “To configure an IPsec VPN” (p. 1262).

7

Assign policies and configurations to the IPsec VPN.

a. If you chose Dynamic (Site-to-Site) for the Tunnel Type parameter in 34.25 “To configure an IPsec VPN” (p. 1262), see 34.26 “To assign policies and configurations for a dynamic site-to-site IPsec VPN” (p. 1264).

b. If you chose Dynamic (Soft Client) for the Tunnel Type parameter in 34.25 “To configure an IPsec VPN” (p. 1262), see 34.27 “To assign policies and configurations for a dynamic soft client IPsec VPN” (p. 1265).

c. If you chose Static for the Tunnel Type parameter in 34.25 “To configure an IPsec VPN” (p. 1262), see 34.28 “To assign policies and configurations for a static IPsec VPN” (p. 1266).

8

Deploy the IPsec VPN.

34.7 Workflow to enable BFD over a static LAN-to-LAN IPsec tunnel

34.7.1 Stages

1

Create a tunnel interface on a VPRN, as described in 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249).

2

Create an IPsec tunnel on the VPRN tunnel interface, as described in 34.21 “To configure an IPsec tunnel on a VPRN tunnel interface” (p. 1253).

3

Enable BFD for the static LAN-to-LAN IPsec tunnel, as described in 34.24 “To enable BFD for a static LAN-to-LAN IPsec tunnel” (p. 1261).

4

Assign a BFD service and interface that can be used for the BFD session on the IPsec tunnel.

34.8 To configure an IPsec IKE policy

34.8.1 Purpose

Use this procedure to set up IPsec IKE policies.

34.8.2 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→IKE Policy, or choose a policy and click Properties. The IPsec IKE Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click on the NAT Traversal tab and configure the required parameters.

5

Click on the DPD tab and configure the required parameters.

The Interval and Max Retries parameters can only be configured when the Dead Peer Detection (DPD) parameter is set to Enable.

6

Click on the Lockout tab and configure the required parameters.

If you select the Enable Lockout check box, you can configure a set of parameters that define the lockout condition.

7

If applicable, click on the IKE Transforms tab and associate IKE Transform policies to the IKE policy. Depending on the NE software release, up to four IKE Transform policies may be applicable.

Perform the following steps:

1. Choose an entry and click Properties. The IPSec IKE Transform Associations form opens.
2. Click Select in the IKE Transform panel and choose or create an IKE transform policy to associate. See 34.9 “To configure an IKE transform policy” (p. 1239).
3. Click OK to close the form.

8

If the IKE version is v2, the Fragment tab becomes available.

If you select the Fragment check box, you can configure a set of parameters that define the fragmentation.

9

Click Apply.

10

Distribute the policy to NEs.

11

Close the IPsec IKE Policy (Create|Edit) form.

12

Close the IPsec Policies form.

END OF STEPS

34.9 To configure an IKE transform policy

34.9.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→IKE Transform, or choose a policy and click Properties. The IPsec IKE Transform (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Close the IPsec IKE Transform (Create|Edit) form.

6

Close the IPsec Policies form.

END OF STEPS

34.10 To configure an IPsec transform policy

34.10.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Transform Policy, or choose a policy and click Properties. The IPsec Transform (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Distribute the policy to NEs.

6

Close the IPsec Transform (Create|Edit) form.

7

Close the IPsec Policies form.

END OF STEPS

34.11 To configure an IPsec static security association

34.11.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Static Security Association Policy, or choose a policy and click Properties. The IPsec Static Security Association (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Distribute the policy to NEs.

6

Close the IPsec Static Security Association (Create|Edit) form.

7

Close the IPsec Policies form.

END OF STEPS

34.12 To configure an IPsec tunnel template

34.12.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Tunnel Template Policy. The IPsec Tunnel Template (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the IPsec Transforms tab.

5

Click Select in the Transform ID 1 to Transform ID 4 panels to choose one or more IPsec transforms.

6

Click Apply.

7

Distribute the policy to NEs.

8

Close the IPsec Tunnel Template (Create|Edit) form.

9

Close the IPsec Policies form.

END OF STEPS

34.13 To configure an IPsec security policy

34.13.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

Click on the Sites tab.

4

Choose a site and click Properties. The VPRN Site (Edit) form opens.

5

Click on the IPsec Security Policies tab.

6

Click Create. The IPsec Security Policy (Create) form opens.

7

Configure the required parameters.

8

Click on the IPsec Security Policy Entries tab.

9

Add one or more security policy entries.

Use the following steps:

1. Click Create. The Security Policy Entry (Create) form opens.
2. Configure the Entry ID parameter.
3. Configure the parameters in the Local IP Address panel.
4. Configure the parameters in the Local IPv6 Address panel.
5. Configure the parameters in the Remote IP Address panel.
6. Configure the parameters in the Remote IPv6 Address panel.
7. Save your changes and close the form.

10

Save your changes and close the forms.

END OF STEPS

34.14 To configure a RADIUS authentication policy

34.14.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Radius Authentication Policy, or choose a policy and click Properties. The IPsec Radius Authentication Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a RADIUS server policy.

5

Click Apply.

6

Distribute the policy to NEs.

7

Close the IPsec Radius Authentication Policy (Create|Edit) form.

8

Close the IPsec Policies form.

END OF STEPS

34.15 To configure a RADIUS accounting policy

34.15.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Radius Accounting Policy, or choose a policy and click Properties. The IPsec Radius Accounting Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a RADIUS server policy.

5

Click Apply.

6

Distribute the policy to NEs.

7

Close the forms.

END OF STEPS

34.16 To configure an IPsec traffic selector list

34.16.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Traffic Selector List, or choose a list and click Properties. The IPsec Traffic Selector List (Create|Edit) form opens.

3

Configure the List Name parameter.

4

Click on the Local Entries tab.

5

Click Create. The IPsec TS Local List Entry (Create) form opens.

6

On the General tab, perform one of the following:

a. Set the Address Format parameter to None.

b. Set the Address Format parameter to Range Address and configure the From IP Address and To IP Address parameters.

c. Set the Address Format parameter to Prefix and configure the IP Prefix and Prefix Length parameters:

7

On the Protocol tab, perform one of the following.

a. Choose Any.

b. Choose a specific protocol option and a port option. If you choose the Manual port option, you need to configure the minimum and maximum port parameters.

c. Choose Manual and configure a Protocol ID.

8

Save your changes and close the IPsec TS Local List Entry (Create) form.

9

Distribute the policy to NEs.

10

Close the IPsec Traffic Selector List (Create|Edit) form.

11

Close the IPsec Policies form.

END OF STEPS

34.17 To configure a trust anchor profile

34.17.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Trust Anchor Profile, or choose a profile and click Properties. The IPsec Trust Anchor Profile (Create|Edit) form opens.

3

Configure the general parameters.

4

Associate one or more certificate authority profiles with the trust anchor profile.

Use the following steps:

1. Click on the Trust Anchor CA Profile tab.
2. Click Create. The Trust Anchor Entry (Create) form opens.
3. Select a certificate authority profile.
4. Save and close the form.

5

Click Apply.

6

Distribute the policy to NEs.

7

Close the IPsec Trust Anchor Profile (Create|Edit) form.

8

Close the IPsec Policies form.

END OF STEPS

34.18 To configure a certificate profile

34.18.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→Certificate Profile, or choose a profile and click Properties. The IPsec Certificate Profile (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Certificate Profile Entries tab.

5

Click Create. The Certificate Profile Entry (Create) form opens.

6

Configure the required parameters.

7

Associate one or more certificate chain CA profiles with the certificate profile entry.

Use the following steps:

1. Click on the Certificate Chain CA Profile tab.
2. Click Create. The Certificate Chain CA Profile (Create) form opens.
3. Select a certificate authority profile.
4. Save your changes and close the form.

8

Click Apply.

9

Distribute the policy to NEs.

10

Close the IPsec Certificate Profile (Create|Edit) form.

11

Close the IPsec Policies form.

END OF STEPS

34.19 To configure an IPsec client database

34.19.1 Steps

1

Choose Policies→ISA Policies→IPsec Policies from the NFM-P main menu. The IPsec Policies form opens.

2

Click Create→IPsec Client DB, or choose a profile and click Properties. The IPsec Client DataBase (Create|Edit) form opens.

3

Configure the required parameters.

Note: The Private Service ID and Private Interface Name parameters can only be configured on local policies.

4

Click on the Client Entries tab.

5

Click Create to create a client entry. Configure the required parameters and close the form.

Up to 8000 entries can be created per database.

6

Save your changes and close the IPsec Client DataBase (Create|Edit) form.

7

To configure the local-only parameters:

1. Choose a global policy and click Properties.

2. Change the distribution mode to Local Edit Only; see 49.9 “To change the distribution mode of a policy” (p. 1482).
3. Configure the Private Service ID and Private Interface Name parameters.
4. Save your changes and close the form.

8

Add the IPsec client database to the IPsec gateway; see 34.22 “To configure an IES or VPRN IPsec gateway” (p. 1256).

The gateway association appears in the IPsec Gateways tab of the global policy.

END OF STEPS

34.20 To configure a tunnel interface on an IES or VPRN

Note: The tabs and parameters that are configurable vary depending on the NE.
The 7705 SAR only supports the configuration of a tunnel interface on a VPRN service.

34.20.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES or VPRN service and click Properties. The service properties form opens.

3

Click on the Interfaces tab. The TMS Interfaces tab is displayed.

4

Click on the Tunnel Interfaces tab.

5

Click Create. The Tunnel Interface (Create) form opens.

6

Configure the required parameters.

7

Configure Cflowd sampling, if applicable.

Use the following steps:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.

2. Configure the parameters.
3. Save your changes and close the form.

8

Configure a port for the interface.

Use the following steps:

1. Click on the Port tab.
2. Select a port.

   Additional tabs may become available based on the type of port selected.
3. Configure the required parameters.

   Note:

   You can configure the NFM-P to automatically assign the lowest unused outer encapsulation value by enabling the Auto-Assign ID parameter.

   You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter in the User Preferences form. Choose Application→ User Preferences from the main menu.

   Private and public tunnel connection points support only dot1q encapsulation.

9

To configure multichassis redundancy, if applicable, see 41.4 “To configure MC IPsec on an MC peer group” (p. 1343).

10

Assign ingress and egress QoS policies to the interface.

Use the following steps:

1. Click on the QoS tab.
2. Configure the required parameters.
3. Click Select in the Ingress Policy and Egress Policy panels to choose ingress and egress QoS policies.

11

Assign an accounting policy to the interface, if applicable.

Use the following steps:

1. Click on the Accounting tab.
2. Select an accounting policy.
3. Configure the Collect Accounting Statistics parameter.

12

Configure scheduling.

Use the following steps:

1. Click on the Schedulers tab.
2. Configure the Aggregate Rate Limit (kbps) parameter.

   Note:

   The Aggregate Rate Limit (kbps) parameter is configurable when the Assign Aggregate Rate Limit check box is enabled and there is no scheduler specified in the Egress Scheduler panel.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.
3. Select ingress and egress schedulers in the Ingress Scheduler and Egress Scheduler panels.
4. Select ingress and egress policer control policies in the Ingress Policer Control Policy and Egress Policer Control Policy panels.

13

Assign ingress and egress ACL filters to the interface.

Use the following steps:

1. Click on the ACL tab.
2. Select ingress and egress ACL filters in the IP Filter and IPv6 Filter panels.

14

Configure BFD.

Use the following steps:

1. Click on the BFD tab.
2. Set the Administration Status parameter to Up. Additional parameters are displayed.
3. Configure the required parameters.
4. To view local and remote session peers, click on the BFD Session tab. The NFM-P retrieves information from the local and remote NEs and lists the current BFD sessions.
5. Select a session and click Properties. The session properties form opens.

   View the following:
   - BFD status
   - protocol used
   - local address
   - remote address
   - operational status and statistics

   The BFD Status field indicates one of the following:
   - no service, when BFD is disabled

- in service, when BFD is running
- out of service, when BFD has failed

6. Close the session properties form.

   Note:

   You cannot enable BFD on an interface if BFD is not configured on the interface, and cannot administratively disable the interface if a routing protocol that uses the interface has BFD enabled.

   See Chapter 28, “Routing protocol configuration” for information about enabling and disabling BFD for routing protocols.

15

Assign a security zone policy to the VPRN tunnel interface, if applicable.

Use the following steps:

1. Click on the Zone tab.
2. Select a security zone policy.
3. Configure the ByPass Zone Config parameter.
4. Save your changes and close the form.

16

Assign an IP address to the interface.

Use the following steps:

1. Click on the Addresses tab and click Create. The IP Address (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

17

Configure ICMPv6 if IPv6 is enabled.

Use the following steps:

1. Click on the ICMPv6 tab.
2. Configure the parameters, as required.
3. Save and close the form.

18

Configure an IP-IP tunnel or GRE tunnel on the interface.

Note: Step 8 must be performed before this step.

Use the following steps:

1. Click on the IP/GRE Tunnels tab and click Create. The IP/GRE Tunnel - IP-GRE|IP-IP Tunnel (Create) form opens.
2. Configure the required parameters.
3. If you are configuring a GRE tunnel, enable the Include GRE Header parameter and configure the GRE tunnel parameters.
4. Select a delivery service in the Delivery Service panel.
5. Click on the States tab and configure the Administrative State parameter.
6. Save and close the form.

19

To specify queue overrides, click on the Override Policy Items tab.

Note: The Override Policy Items tab contains a number of sub-tabs. However, the sub-tabs that are displayed depend on the port type that you have chosen for this interface.

- If the port is not an HSMDA port, the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer sub-tabs are active.
- If the port is an HSMDA port, the Access Ingress Queues, Access Egress HSMDA Queues, and Ingress Policer sub-tabs are active.

20

Save your changes and close the forms.

END OF STEPS

34.21 To configure an IPsec tunnel on a VPRN tunnel interface

Note: Availability of some parameters varies depending on the NE and release; see the NE documentation for more information.

34.21.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

Click on the Interfaces tab. The TMS Interfaces tab is displayed.

4

Click on the Tunnel Interfaces tab.

5

Choose a tunnel interface and click Properties. The Tunnel Interface (Edit) form opens.

6

Click on the IPsec Tunnels tab.

Note: The IPsec Tunnels tab is displayed only when a port is assigned to the tunnel interface.

7

Click Create. The IPsec Tunnel (Create) form opens.

8

Configure the required parameters.

9

Select a security policy in the Security Policy ID panel.

To ensure that no multiple tunnels are created with same security policy, enable the Strict Match parameter.

Note: Strict Match parameter must be set at the same time Security Policy ID is set.

10

Configure the required parameters in the Tunnel Endpoints panel.

11

Configure the required parameters in the IP Fragmentation panel.

12

Configure the required parameters in the ICMPv6 Generation panel.

13

If you set the Keying parameter to Dynamic in Step 8, configure dynamic keying.

Use the following steps:

1. Click on the Dynamic Keying tab.
2. Select an IPsec transform in the Transform ID 1 to Transform ID 4 panels.
3. Select an IKE policy in the IKE Policy panel.
4. Configure the Pre-shared Key and Auto-Establish parameters.

5. Configure the parameters in the Local ID panel.

   Note:

   The parameters in the Local ID panel, and the Certificate File and Key File parameters, are configurable when IKEv2 is specified in the IKE policy associated with the tunnel, and the Authorization Method is set to Certificate Authentication.

6. To specify a single certificate trust anchor, certificate file, and key file where available:
   a. Configure the Certificate File and Key File parameters.
   b. Click Select beside the Certificate Trust Anchor parameter to select a CA profile.

   Note:

   If there is a problem with the Certificate File or the Key File after the tunnel becomes administratively up, the Invalid Certificate File or Invalid Key File operational indicators are enabled on the States tab, and the NFM-P raises an alarm.

7. To specify multiple certificate trust anchors, certificate files, and key files, click Select beside the Trust Anchor Profile and Certificate Profile parameters to select the appropriate profiles.

   Note:

   If there is a problem with a Certificate File or Key File after the tunnel becomes administratively up, the Invalid Certificate File or Invalid Key File operational indicators are enabled on the States tab, and the NFM-P raises an alarm.

14

If you set the Keying parameter to Manual in Step 8, configure manual keying.

Use the following steps:

1. Click on the Manual Keying tab.
2. Click Create. The IPsec Security Association (Create) form opens.
3. Select a security policy entry in the Security Policy Entry panel.
4. Configure the required parameters.
5. Select a transform policy in the Transform panel.
6. Save your changes and close the form.

15

Configure BFD.

Use the following steps:

1. Click on the BFD tab.
2. Configure the required parameters.
3. Select a BFD service in the BFD Service panel.
4. Select an interface in the Interface panel.
5. Configure the Destination Address parameter.

16

Click on the States tab.

17

Configure the Administrative State parameter.

18

Configure static tunnel destination IP addresses.

Use the following steps:

1. Click on the Dest-IP Addresses tab and click Create. The IPsec Tunnel Dest-IP Address (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

19

Save your changes and close the forms.

END OF STEPS

34.22 To configure an IES or VPRN IPsec gateway

Note: Availability of some parameters varies depending on the NE and release; see the NE documentation for more information.

The parameters in the Local ID panel and the Certificate File and Key File parameters, if applicable, are configurable when IKEv2 is specified in the associated gateway IKE policy and the Authorization Method is set to Certificate Authentication.

34.22.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES or VPRN service and click Properties. The service properties form opens.

3

Click on the Interfaces tab, then on the L3 Access Interfaces tab.

4

Choose the L3 access interface on which you want to create the IPsec gateway and click Properties. The L3 Access Interface (Edit) form opens.

Note: The port configured on the L3 access interface must be a Tunnel Group SAP, which is the public-facing interface for an IPsec tunnel.

5

Click on the IPsec Gateway tab.

6

Click Create. The IPsec Gateway (Create) form opens.

7

Configure the required parameters.

8

To associate an IKE policy with the IPsec gateway, select an IKE policy in the IKE Policy panel.

Note: For the IPSec Gateway administrative status to be set to “Up”, the associated IKE Policy must have a IKE Transform policy configured on it.

9

To associate an IPsec tunnel template with the IPsec gateway, select an IPsec tunnel template in the IPsec Tunnel Template panel.

10

To associate an IPsec client database with the IPsec gateway, select an IPsec client database in the IPsec Client DataBase panel.

11

Select the far-end VPRN service site of the tunnel in the Secure Service Id panel.

12

Select the tunnel interface of the far-end site in the Tunnel Interface Name panel.

13

Configure the Local Gateway Address parameter.

14

Configure the parameters in the Local ID panel.

15

Configure the parameters in the Certificate panel as needed.

1. To specify a single CA to use as a trust anchor, click Select beside the Certificate Trust Anchor parameter and select a certificate trust anchor.

2. To specify multiple certificates and CAs, click Select beside the Trust Anchor Profile and Certificate Profile parameters, and select the appropriate profiles.

Note: If there is a problem with a Certificate File or Key File after the tunnel becomes administratively up, the Invalid Certificate File or Invalid Key File operational indicators are enabled on the States tab, and the NFM-P raises an alarm.

16

Select a RADIUS authentication policy in the Radius Authentication Policy panel.

17

Select a RADIUS accounting policy in the Radius Accounting Policy panel.

18

Click on the States tab and configure the Administrative State parameter.

19

Associate a traffic selector list with the IPsec gateway.

Use the following steps:

1. Click on the Traffic Selector Negotiation tab.
2. Click Create. The Traffic Selector Negotiation (Create) form opens.
3. Select a traffic selector policy in the Traffic Selector Policy panel.
4. Save your changes and close the form.

20

On IES and VPRN L3 access interface IPsec gateways, you can lease a local IP address from a pool in a local DHCPv4 or DHCPv6 server defined in the VPRN routing instance or NE base routing instance. To lease an IP address from an external DHCP server, see Step 21.

To configure the local address assignment for the IKEv2 remote access tunnel on an IES or VPRN L3 access interface IPsec gateway:

1. Click on the Local Address Assignment tab.
2. Configure the Administrative State parameter.
3. Select the DHCP server name in the IPv4 and IPv6 panels.

   Note:

   After the tunnel is established and the address is leased, you can view the lease on the Leases tab on the Local DHCP Server (Edit) form. On the VPRN Site (Edit) form, click on the RADIUS/DHCP/Diameter tab and select the required DHCPv4 server to open the Local DHCP Server (Edit) form. You can filter the leases by choosing IPsec from the Client Type column header.

   On the NE base routing instance, navigate to Routing→*NE*→Routing Instance, right click and choose Properties. Click on the Local DHCP Servers tab.

See 79.33 “To configure a local DHCPv4 server on a VPRN site” (p. 2579) and 79.34 “To configure a local DHCPv6 server on a VPRN site” (p. 2581) for information about configuring a local DHCPv4 or DHCPv6 server on a VPRN site. See 27.5 “To configure a local DHCPv4 server on a routing instance” (p. 839) and 27.6 “To configure a local DHCPv6 server on a routing instance” (p. 842) for information about configuring a local DHCPv4 or DHCPv6 server on a base routing instance.

21

Lease an IP address for the IPsec gateway from an external DHCPv4 server.

Use the following steps:

1. Click on the DHCP Address Assignment tab and click Create. The Internal DHCPv4 Address Assignment form opens.
2. Configure the Gi Address and either the Router or Service ID parameters.
3. Specify up to eight DHCPv4 servers in the DHCP Servers panel.
4. Save and close the form.

22

Lease an IP address for the IPsec gateway from an external DHCPv6 server.

Use the following steps:

1. Click on the DHCPv6 Address Assignment tab and click Create. The Internal DHCPv6 Address Assignment form opens.
2. Configure the Link Address and either the Router or Service ID parameters.
3. Specify up to eight DHCPv6 servers in the DHCP Servers panel.
4. Save and close the form.

23

Close the forms.

END OF STEPS

34.23 To view current remote users connected to an IPsec gateway and remote user security associations

34.23.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES or VPRN service and click Properties. The service properties form opens.

3

Expand the Sites object in the service navigation tree.

4

Expand the Routing Instance object in the service navigation tree.

5

Expand the L3 Access Interfaces object in the service navigation tree, click on the appropriate L3 access interface, and click Properties. The L3 Access Interface (Edit) form opens.

6

Click on the IPsec Gateway tab.

7

Choose the IPsec gateway for which you want to view the current remote users and click Properties. The IPsec Gateway (Edit) form opens.

8

Click on the Statistics tab to view and collect statistics as required.

Note: If a remote user has been removed, the previously collected statistics for the remote user will remain visible on the Statistics tab of the IPsec gateway.

9

Click on the Remote Users tab. The current remote users are listed.

10

Select a user and click Properties. The IPsec Remote User (Edit) form opens.

11

Click on the Remote User Security Associations tab. The security associations of the remote user are listed.

12

Select a security association and click Properties. The IPsec Remote User Security Association (Edit) form opens.

13

View general information and parameters that were negotiated with IPsec gateway.

These include but are not limited to:

- Internet IP address
- Private IP address and Second Private IP address

- SA lifetime
- DH group
- Replay window

14

Click on the Algorithms tab to view authentication and encryption algorithm information.

15

Click on the Traffic Selector tab to view the list of IPsec remote user dynamic security association traffic selector entries.

16

Close the forms.

END OF STEPS

IPsec remote user statistics

Starting with Release 20.6, the NFM-P collects and displays IPsec remote user statistics under the Statistics tab of the IPSec gateway instead of the IPsec Remote User. This means that IPsec remote user statistics are retained in the event of a failover.

34.24 To enable BFD for a static LAN-to-LAN IPsec tunnel

34.24.1 Steps

1

On the private side, use the device CLI to specify a tunnel and to enable BFD on the tunnel.

Note: See the device documentation for information about the CLI commands required to perform the action.

2

On the public side, use the device CLI to configure the transmit interval.

Note: See the device documentation for information about the CLI commands required to perform the action.

3

Create a tunnel interface in a VPRN service, as described in 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249).

4

Create an IPsec tunnel on the tunnel interface, as described in 34.21 “To configure an IPsec tunnel on a VPRN tunnel interface” (p. 1253).

5

Close all open forms.

END OF STEPS

34.25 To configure an IPsec VPN

Note: The NFM-P does not support IPsec VPN creation when language localization is enabled.

See 1.2.5 “Localized language support” (p. 79) in Chapter 1, “NFM-P GUI” for information.

34.25.1 Steps

1

Choose Manage→Service→IPsec from the NFM-P main menu. The Manage IPsec VPNs form opens.

2

Click Create. The IPsec VPN (Create) step form opens with the Create IPsec VPN step displayed.

3

Configure the parameters in the IPsec VPN panel.

4

Specify a corporate service.

Use the following steps:

1. Configure the Service Type parameter.
2. Select a corporate service in the Corporate Service panel.

5

Configure the Link Corporate and Secured Service parameter in the Composite Service panel.

6

Click Next. The Select Service Sites step form is displayed.

7

Click Create to select the IPsec VPN service sites.

8

Click Next. The Create/Select Secure VPRN Service step form is displayed.

9

Perform one of the following.

a. Click Select Service to select a service, and configure the parameters on the Service Site tab.

b. Click Create Service to create a service, and configure the parameters on the Service and Service Site tabs.

c. Click Create Service from Template and select a service.

Note: The template must be configured for one site. The template cannot be used if the template is configured for multiple sites.

10

Click Next. The Create/Select Delivery Service(s) step form is displayed.

11

CAUTION

Service Disruption

Service Disruption

When you remove a service, the service is removed without confirmation and cannot be restored.

Perform one of the following.

a. To select a delivery service, click Select Service, select a service, and configure the parameters on the Service Site tab.

b. To create a delivery service, click Create Service, choose IES or VPRN, and configure the parameters on the Service and Service Site tabs.

c. To create a service from a template, click Create Service from Template, choose VPRN Templates or IES Templates, and select a template.

Note: The template must be configured for one site. The template cannot be used if the template is configured for multiple sites.

d. To remove a service, click Remove Service, select a service, and click Remove Service.

12

Click Next. The Tunnel Group Selection step form is displayed.

13

Configure each parameter that has a field with a yellow background by clicking in the field and entering a value or choosing an option.

Note: The parameters are not included in the XML model. However, using the OSS interface, you can create the ServiceSiteStructs and can create the IPsec VPN objects in the ServiceSiteStructs.

14

Click Next. The Create Policy step form opens.

15

Perform one of the following, depending on the Tunnel Type parameter setting in Step 13.

a. If the parameter is set to Dynamic (Site-to-Site), perform 34.26 “To assign policies and configurations for a dynamic site-to-site IPsec VPN” (p. 1263).

b. If the parameter is set to Dynamic (Soft Client), perform 34.27 “To assign policies and configurations for a dynamic soft client IPsec VPN” (p. 1265).

c. If the parameter is set to Static, perform 34.28 “To assign policies and configurations for a static IPsec VPN” (p. 1266).

END OF STEPS

34.26 To assign policies and configurations for a dynamic site-to-site IPsec VPN

34.26.1 Purpose

Perform this procedure if the Tunnel Type parameter in Step 13 of 34.25 “To configure an IPsec VPN” (p. 1262) is set to Dynamic (Site-to-Site).

34.26.2 Steps

1

Select an IPsec tunnel template.

2

Configure the Pre-shared key parameter.

3

Click Finish.

The NFM-P performs the following:

- creates a secure VPRN service and sites
  - creates an access interface
  - creates a tunnel interface
  - configures the service name
- creates the IES or VPRN delivery service and sites

- creates an access interface
- configures the IP address of the SAP
- assigns a tunnel group with an auto-generated outer encapsulation as a public SAP
- creates an IPsec gateway
- assigns the secure service
- assigns the tunnel interface of the secure service
- assigns the tunnel template
- assigns the IKE policy
- configures the local gateway address
- configures the pre-shared key
- configures the service name

- if the Link Corporate and Secure Service parameter is enabled, creates a composite service that contains the corporate and secure services

4

To view the IPsec VPN, click View the newly created IPsec Secured VPN. The IPsec VPN (Create) step form closes, and the IPsec VPN form opens.

5

Close the forms.

END OF STEPS

34.27 To assign policies and configurations for a dynamic soft client IPsec VPN

34.27.1 Purpose

Perform this procedure if the Tunnel Type parameter in Step 13 of 34.25 “To configure an IPsec VPN” (p. 1262) is set to Dynamic (Soft Client).

34.27.2 Steps

1

Select an IPsec tunnel template.

2

Click Select beside the Displayed Name parameter to select a subscriber authentication policy.

3

Click Finish.

The NFM-P performs the following:

- creates a secure VPRN service and sites

  - creates an access interface
  - creates a tunnel interface
  - configures the IP address of the SAP
  - configures the service name
- creates the IES or VPRN delivery service and sites
  - creates an access interface
  - configures the IP address of the SAP
  - assigns the RADIUS authentication policy
  - assigns a tunnel group with an auto-generated outer encapsulation as a public SAP
  - creates an IPsec gateway
  - assigns the tunnel template
  - assigns the IKE policy
  - configures the local gateway address
  - configures the pre-shared key
  - configures the service name
- if the Link Corporate and Secure Service parameter is enabled, creates a composite service that contains the corporate and secure services

4

To view the IPsec VPN, click View the newly created IPsec Secured VPN. The IPsec VPN (Create) step form closes, and the IPsec VPN form opens.

5

Close the forms.

END OF STEPS

34.28 To assign policies and configurations for a static IPsec VPN

34.28.1 Purpose

Perform this procedure if the Tunnel Type parameter in Step 13 of 34.25 “To configure an IPsec VPN” (p. 1262) is set to Static.

34.28.2 Steps

1

Select an IPsec tunnel template.

2

Configure the required parameters.

3

If you set the Keying parameter to Manual, configure the parameters in the Manual Keying - Inbound and Manual Keying - Outbound panels.

4

Select an IPsec transform policy.

5

Configure the Auto-Establish parameter.

Note: The Auto-Establish parameter is configurable only if dynamic keying is enabled.

6

Click Finish.

The NFM-P performs the following:

- creates a secure VPRN service and sites
  - creates the IPsec security policy
  - creates a tunnel interface
  - assigns a tunnel group with an auto-generated outer encapsulation as a private SAP
  - creates the IPsec tunnel
  - assigns the security policy to the tunnel
  - configures the local and remote gateway addresses
  - configures the delivery service
  - configures the replay window
  - configures the keying
  - if set, enables auto-establish
  - configures the static route
  - configures the service name
- creates the IES or VPRN delivery service and sites
  - creates an access interface
  - assigns a tunnel group with an auto-generated outer encapsulation as a public SAP
  - configures the IP address of the SAP
  - configures the service name
- if the Link Corporate and Secure Service parameter is enabled, creates a composite service that contains the corporate and secure services

7

To view the IPsec VPN, click View the newly created IPsec Secured VPN. The IPsec VPN (Create) step form closes, and the IPsec VPN form opens.

8

Close the forms.

END OF STEPS

34.29 To configure an IPsec tunnel on a IES or VPRN service

34.29.1 Before you begin

For VPRN service IPsec tunnel, IPsec security policy and entry should belong to same VPRN.

34.29.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES or VPRN service and click Properties. The service properties form opens.

3

Choose the L3 access interface on which you want to configure an IPsec tunnel and click Properties. The L3 Access Interface (Edit) form opens.

4

Click the IPSec tab.

5

Click Create. The IPSec (Create) form opens. Configure the required parameters.

6

Click IPSec Tunnel tab.

7

Click Create. The IPSec Tunnel (Create) form opens.

8

Configure the parameters on the General tab.

9

Configure the Copy Traffic Class Upon Decapsulation and Keying parameters.

Note: Based on the value selected for the Keying parameter, additional tabs will be displayed.

10

If Dynamic value is selected for the Keying parameter, the Dynamic Keying tab is displayed. Click Dynamic Keying tab.

11

Configure the IPsec transform sets, IKE policies and the Pre-shared key parameter.

12

Save your changes and close the forms.

END OF STEPS

35 ISA-Video

ISA-Video overview

35.1 ISA_Video overview

35.1.1 Overview

The NFM-P supports, on a 7450 ESS or 7750 SR, the equipment and service configuration of the following enhancements to multicast video services based on the ISA-Video MDA:

- Reliable Delivery/Retransmission (RT) Proxy
- Fast Channel Change (FCC)

You can use the NFM-P to configure the required ISA-Video groups and members, as well as the associated multicast info policies. VPRN and IES services can then be configured for the customer video delivery.

These enhancements are provisioned in NFM-P in the following three areas:

- a video group configuration to manage the ISA resources and service grouping.
- multicast information policy configuration parameters and objects to handle the RT and FCC capabilities.

  These include the following:
  - a video policy configuration under the multicast information policy
  - video parameters added to multicast bundles, channels and channel override configurations
- video interface configurations added under VPRN/IES services sites.

The following figure shows the relationship of these elements and how they are associated through configuration.

Figure 35-1 ISA-Video configuration elements

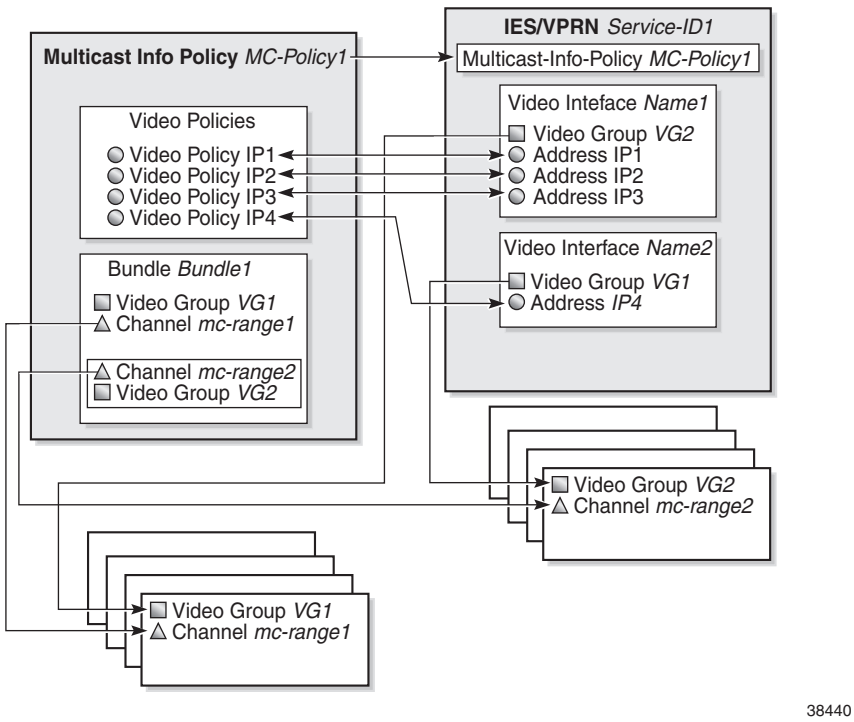

A video interface within a service site can have up to eight IP addresses. The IP addresses assigned to a video interface determine which multicast info policy is applied. The video group you assign to the video interface determines which bundle/channel configuration inside the multicast info policy is applied to the interface. You also configure the ingress and egress QOS policies for the video interface using the video group.

If a request is received on a video interface for a channel not serviced by the video group associated with that specific video interface, then the request is considered invalid and is dropped. For example, in Figure 35-1, “ISA-Video configuration elements” (p. 1272) , a request for mc-range2 received on IP1, IP2, or IP3 is invalid. A request for mc-range2 is only valid on IP4.

A video group manages the ISA-Video MDA resource and the configurations of the enhanced video functionality. You can assign a video group (for example, VG1 in Figure 35-1, “ISA-Video configuration elements” (p. 1272)) for each bundle, and this is the default video group for all of the channels in that bundle. You can then apply specific configuration overrides for each channel within a video group assignment.

Up to four video groups (ID 1-4) can be configured on a single NE. Each video group in turn, can have multiple primary video ISAs assigned to provide load balancing and redundancy protection.

However, any specific video ISA can only be bound to one video group. The binding of a video ISA to the video group is referred as a video group member.

Additional considerations for the Ad Insertion functionality:

- Ingress Ad Insertion channels are configured on the video interface configuration form.
- If Ad Insertion functionality is enabled on a video group, then the group can only have one video group member.
- You can configure zone channels for egress Ad Insertion on each Ad Insertion channel configuration.

Workflow to configure and manage an ISA-Video configuration

35.2 Workflow to configure and manage an ISA-Video configuration

35.2.1 Stages

1

Provision the ISA-Video MDA. See Chapter 11, “Working with network objects” and Chapter 12, “Device object configuration” for more information about ISA-Video MDA equipment configuration.

2

Create or configure ISA-Video groups and add members under the video groups. See 13.13 “To configure an ISA-Video group” (p. 427) for more information.

3

Configure and distribute a global Multicast Information policy. Configure the channel bundle, channel range, channel overrides, and video interface for the NE. See 52.13 “To configure an ingress multicast information policy” (p. 1720) for more information.

4

Create and enable a video interface for the IES or VPRN service. See 35.3 “To add a video interface to an IES or VPRN site” (p. 1275) for more information.

ISA-Video procedures

35.3 To add a video interface to an IES or VPRN site

35.3.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required IES or VPRN service and click Properties. The IES or VPRN (Edit) form opens with the General tab displayed.

3

On the navigation tree, click on the site to which you want to add the video interface and expand the entries for that site.

4

Right-click on Video Interfaces and choose Create Video Interface. The Video Interface (Create) form opens with the General tab displayed.

5

Configure the required parameters.

After you select a Group Number, the following selectable fields are displayed:

- Ingress QoS Policy
- Egress QoS Policy
- Ingress IP Filter
- Egress IP Filter

6

Select the required policies and filters. You can search or choose the entry that appears for each of these items.

7

Click on the Security tab and select an NE DoS Protection policy.

8

Click on the Addresses tab and click Create. The IP Address (Create) form opens.

9

Configure the required parameters.

10

To configure ad insertion, click on the ADI Channels tab and click Create. The Video ADI Channel (Create) form opens with the General tab displayed.

11

Configure the required parameters.

12

Click on the Zone Channels tab and click Create. The Video ADI Zone Channel (Create) form opens.

13

Configure the required parameters and save your changes.

14

Save your changes and close the form.

END OF STEPS

36 Alarm management

Alarm management overview

36.1 Alarm management overview

36.1.1 Overview

The NFM-P converts SNMP traps from NEs and NFM-P events to alarms that are associated with the managed equipment, configured services and policies. The NSP alarm management system provides the following:

- conversion of SNMP traps from NEs to alarms using the X.733 standard
- correlation of alarms with equipment- and service-affecting faults
- updates to the managed-object operational status in near-real-time
- the ability to log the actions performed to resolve a related fault by adding notes to the alarm
- alarm history for performing trend analysis

Viewing and managing alarms is performed using the NSP alarm management. Using the NSP alarm management, you can display and filter alarms, explore root causes, and manage, dismiss, or escalate alarms, among other functions. For more information about using the NSP alarm management, see the *NSP Network and Service Assurance Guide*.

36.2 Correlated alarms

36.2.1 Overview

The NFM-P generates a correlated alarm when a fault condition on one object causes an alarm condition on another object. For example, if a port goes down, an alarm is generated for the port, which is the affected object. Each service that uses the port generates an alarm. The alarm information for the affected object includes the correlated alarms.

The NFM-P automatically promotes the severity of correlating alarms to match the severity of correlated alarms. For example, if a major correlating port alarm has one or more critical correlated service alarms, the NFM-P changes the severity of the correlated port alarm to critical. The correlating alarm cannot have a severity that is lower than the severity of the correlated alarms. However, the severity that is defined by a specific alarm policy takes precedence over the correlated severity. The NFM-P does not lower the severity of correlating alarms when correlated alarms are cleared.

Note: Alarm correlation does not affect the severity of alarms that originate from an external EM system.

When the event that initiates an affected object alarm is resolved and the alarm clears, the correlated alarms also clear. If a correlated alarm does not clear after the affected object alarm clears, the source of the correlated alarm is an event other than the initiating event. In this case, the alarm correlation is removed and the alarm is displayed in the dynamic alarm list.

36.2.2 Service and transport alarm correlation

The NFM-P provides limited service and transport alarm correlation. When an LSP goes out of service, the NFM-P generates an alarm for the LSP and a correlated alarm for each LSP path that is a child object of the LSP. The LSP alarm is listed as an aggregated alarm for each SDP that uses the LSP. An NFM-P operator can use the aggregated LSP alarm for SDP troubleshooting.

There may not be a direct relationship between an LSP and the SDPs that use the LSP. For example, when LDP over RSVP is the transport mechanism, the NFM-P does not correlate SDP alarms with LSP alarms.

37 VRRP

37.1 VRRP

37.1.1 Overview

The NFM-P supports VRRP management. VRRP creates a redundant routing system that takes over packet transmission on a common LAN segment when a router fails. VRRP designates alternative routing paths in the form of virtual routers, or VRs, without changing the IP or MAC address of a protected router.

A VRRP protected router owns the IP address of the VR. The VRRP owner forwards packets using the default gateway. When the owner router fails, packet-forwarding responsibilities are transferred to a designated backup router. This router becomes the master and forwards packets using the VR IP address.

The following figure shows a basic VR that acts in parallel with the real network. Router 2 is the owner with the address from which packets are forwarded. If Router 2 fails, Router 3, which has been configured to route using the master address in a backup role, begins to forward packets using this IP address. Router 1 is also a backup router, but because its priority number is higher, it ranks below Router 3.

Figure 37-1 VR concepts

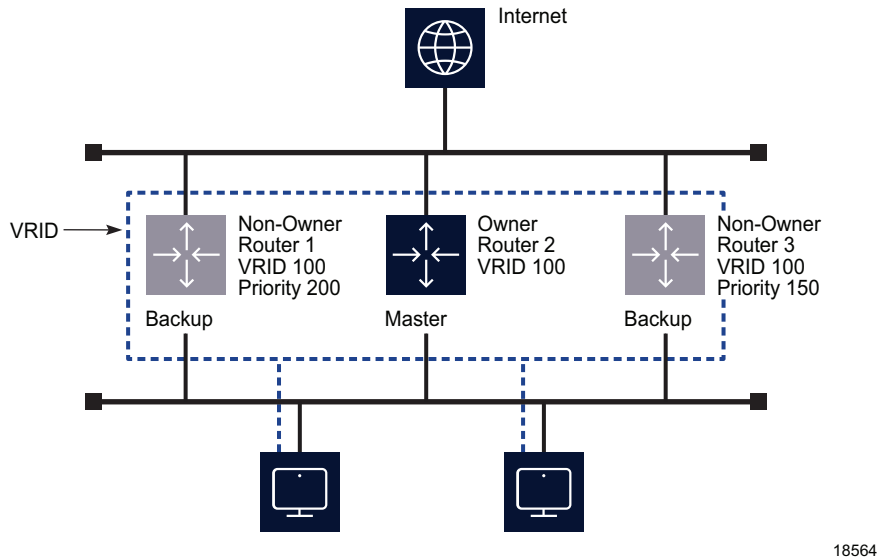

The NFM-P supports the configuration of VRs for network interfaces and for L3 access interfaces.

VRRP in an IES involves interfaces from separate IESs. VRRP in a VPRN requires interfaces that are in the same VPRN service. The NFM-P supports on-demand, but not scheduled, statistics

collection for VRRP in a VPRN. Certain VRRP SNMP traps do not apply to VPRN; see the appropriate NE documentation for information.

You can configure the VR through another set of tabbed forms and a navigation tree, which allows you to add VRRP instances IP owner and non-owner router interfaces, as shown in the following figure.

37.1.2 VR

A VR is a logical entity, managed by VRRP, that acts as a default router for hosts on a shared LAN. The VR consists of a VRID and a subnet (that is, ip_address/mask). A VRRP router can back up one or more VRs. The purpose of supporting multiple IP addresses in a single VR is for multi-netting. This common mechanism allows multiple local subnet attachments on a single routing interface. Up to four VRs are allowed on a single IP interface. The VRs must be in the same subnet.

The following figure shows a common VR setup in which associated routers provide mutual backup using VRRP. Router A forwards packets on IP address 10.1.1.1 to Hosts 1 and 2 on its default gateway. Router B forwards packets on IP address 10.1.1.2 to Host 3 on its default gateway. If Router A fails, VRID 1 uses IP address 10.1.1.1 to forward packets from Router B to Hosts 1 and 2. At the same time, the Router B interface is still configured to deliver packets on IP address 10.1.1.2 to Host 3. If Router B fails, VRID 2 forwards these packets through backup Router A.

Figure 37-2 Sample VRs

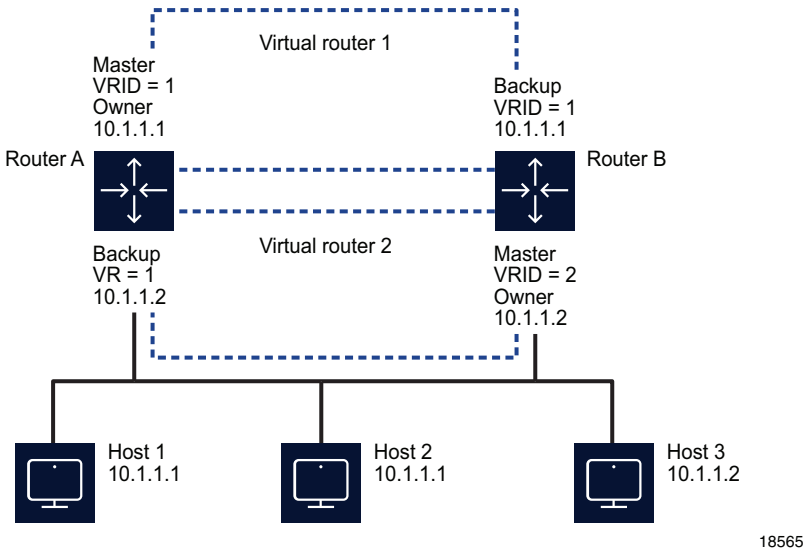

37.1.3 Master router

The VRRP master router, in either a normal or a failover situation, routes all IP packets into the LAN using the physical MAC address for the IP interface as the Layer 2 source MAC address. ARP packets also use the parent IP interface MAC address as the Layer 2 source MAC address.

37.1.4 Owner and non-owner VRRP instances

A VRRP instance is configured in either an owner or non-owner mode.

The owner instance controls the IP address of the VR and is responsible for forwarding packets sent to this IP address. The IP address of the owner VRRP instance is the same as the real interface IP address of the router. The owner assumes the role of the master VR when it is functioning normally in the network. Only one VRRP instance in the domain is configured as the owner. All other instances participating in the domain are non-owners and must have the same VRID.

A backup router becomes the master router after a failover and continues to use the IP address of the original master. As a result, the new master router is the IP address non-owner.

The most important parameter to define for a non-owner VRRP instance is the priority. The priority defines for a VR the selection order. The priority value and the preempt mode combine to determine which VR has the highest priority and becomes the master.

The base priority is used to derive the in-use priority of the VRRP instance as modified by an optional VRRP priority-control policy. VRRP priority-control policies are used to either override or adjust the base priority value depending on events or conditions in the NE. See Chapter 55, “VRRP policies” for more information.

37.1.5 Passive VRRP instances

A VRRP instance can be configured in passive mode. A passive VRRP instance does not transmit or receive keep-alive messages.

The following cannot be configured on a VRRP instance in passive mode:

- Authentication
- BFD Interface
- Init Delay
- Master Inherit Interval
- VRRP Policy association
- Preempt Mode
- Base Priority
- Standby Forwarding

37.1.6 VRRP types

When you create a VR, you specify the VRRP type, for example, network or L3 service, and the VR instance is restricted to the specified VRRP type. Configuring a VR using a mix of network and service interfaces through CLI raises a configuration mismatch alarm.

37.1.7 Primary addresses

A primary IP address is an address that is selected from the set of real interface addresses on the VR. VRRP advertisements between master and backup VRRP instances are sent using the primary IP address as the source of the IP packet.

An IP interface must always have an assigned primary IP address for VRRP to operate on the interface. The primary IP address of the VR and the primary address on the IP interface are always the same.

37.1.8 Backup addresses

A maximum of 16 IP addresses (for either IPv4 or IPv6) in different subnets can be configured for a VRRP instance. One backup address is permitted for a subnet. The number of backup addresses is limited to the number of primary and secondary addresses configured on the IP interface.

The backup IP addresses for the owner VRRP instance must match the primary address or one of the secondary addresses on the IP interface. If the VRRP instance is not the owner, the backup addresses must be in the subnets of the primary and secondary addresses of the IP interface.

The NFM-P includes only eligible IP addresses in the search list.

37.1.9 VRRP message authentication

The type of authentication used by the VR in VRRP advertisement is specified during VRRP instance creation. The current master router uses the configured authentication type when sending VRRP advertisements to backup routers, which authenticate the messages.

37.2 Workflow to configure VRRP

37.2.1 Stages

1

Before you create a VR, ensure that a primary IP address is configured for the network or L3 service interface. Each interface must have a primary IP address. See 27.18 “To configure L3 network interfaces” (p. 863) for more information.

2

Configure an SNMP community for VPRN sites that will contain a VRRP instance. See 79.17 “To configure an SNMP community on a VPRN site” (p. 2551) for more information.

3

Configure the VRRP priority-control policy for a non-owner VRRP instance. See Chapter 55, “VRRP policies” for more information about VRRP priority-control policy events.

4

Create the VR. See 37.3 “To create a VR” (p. 1283) for more information.

5

As required, create VRRP instances in the VR. See 37.4 “To create and configure a VRRP instance” (p. 1283) for more information.

37.3 To create a VR

37.3.1 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Click Create. The Virtual Router (Create) form opens.

3

Configure the required parameters.

Note: If you are creating an IPv6 VR, then you need to specify an IPv6 address as the Backup Address parameter value.

4

Save your changes and close the form.

END OF STEPS

37.4 To create and configure a VRRP instance

37.4.1 Purpose

Before you create a VRRP instance on an interface, ensure that the L3 service or network interface is configured with a primary IP address. The backup IP address that is configured in 37.3 “To create a VR” (p. 1283) must belong to the same subnet as the primary IP address of the owner VRRP instance created in this procedure. The NFM-P automatically configures a backup address using the IP address of the VR. A search of interface IDs in Step 5 of this procedure displays the interface IDs that are eligible for the creation of a VRRP instance for the VR.

Interfaces on VPRN sites must be associated with an SNMP community. See 79.17 “To configure an SNMP community on a VPRN site” (p. 2551) for more information about configuring an SNMP community for a VPRN site.

You can configure up to a total of four VRRP instances (IPv4 plus IPv6) on one IP interface. The maximum number of mixed IPv4 and IPv6 VRRP instances on an NE is 1024.

Note: The 7705 SAR does not support the configuration of a passive VRRP instance or the configuration of VRRP on a network interface.

37.4.2 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Choose a VR and click Properties. The Virtual Router (Edit) form opens.

3

Click on the Components tab.

4

Right-click on VR Instances and choose Create VRRP Instance. The VRRP Instance (Create) form opens.

If the Backup Address parameter value that you specified in Step 3 of 37.3 “To create a VR” (p. 1283) is an IPv6 address, then the contextual menu item is Create VRRP IPv6 Instance, and the VRRP IPv6 Instance (Create) form opens.

5

Assign an interface to the instance.

Perform the following steps:

1. Click Select in the Interface panel. The Select Interface - VRRP Instance or Select Interface - VRRP IPv6 Instance form opens, depending on whether you are creating an IPv4 or IPv6 VRRP instance respectively.
2. Choose an interface, and click OK. The Select Interface - VRRP Instance (or Select Interface - VRRP IPv6 Instance) form closes and the VRRP Instance (Create) or VRRP IPv6 Instance (Create) form reappears with the interface information displayed.

6

Configure the required parameters in the VRRP Instance Information panel.

If your set the Owner parameter to true, then you cannot configure the Administrative State and Base Priority parameters, and cannot associate a policy with this VRRP instance. Go to Step 8 .

7

Assign a VRRP policy to the instance, if required.

The same VRRP policy can be applied to both IPv4 VRRP instances and IPv6 VRRP instances.

1. Click Select in the Policy panel. The Select VRRP Policy form opens.
2. Choose a VRRP policy and click OK. The Select VRRP Policy form closes and the VRRP Instance (Create) form reappears with the VRRP policy information displayed.

8

Click on the Behavior tab and configure the required parameters.

9

If you require bidirectional forwarding detection for the VRRP instance.

Then, perform the following tasks:

1. Select the Enable BFD Interface check box. The BFD Interface tab appears.
2. Click on the BFD Interface tab.
3. Click Select beside the Service Name parameter, and choose a service from the Select Service - BFD Interface form.
4. Click Select beside the Interface Name parameter, and choose an interface from the Select Interface - BFD Interface form.
5. Configure the Destination Address parameter.

10

Click on the Authentication tab, and configure the required parameters.

11

Click on the Backup Addresses tab.

12

Click Create. The Backup Address (Create) form opens.

13

Perform one of the following to configure backup addresses:

a. Enter an IPv4 address (for an IPv4 VR) or IPv6 address (for an IPv6 VR) manually in the IP Address field in the Backup Address section.

b. Select the IP address from a list.

Perform the following steps:

1. Click Choose IP Address adjacent to the IP Address field. The Select IP Address form opens.
2. Choose an IP address and click OK.

   The Select IP Address form closes and the IP Address parameter is populated with your selection.
3. Click Add Link Local Address. The Admin Link Local Address defined for the network L3 access interface appears in the IP Address field.

   If you need to add the Link Local Address as a backup address, then the Admin Link Local Address must be configured and the Admin Link Local Address Preferred parameter must be enabled for the network L3 access interface that you are using. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for more information.

14

Save your changes and close the form.

15

Click OK. The Virtual Router (Edit) form reappears with the VRRP instances displayed in the VR instances tree.

END OF STEPS

37.5 To add a VRRP instance

37.5.1 Purpose

Use the following procedure to add a VRRP instance. A VRRP instance that you add to the VR must have an IP address with the same VRID and subnet as those of the VR.

37.5.2 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Choose a VR and click Properties. The Virtual Router (Edit) form opens.

3

Click on the Components tab.

4

Right-click on VR Instances and choose Add VRRP Instance.

Note: If you chose an IPv6 VR in Step 2 , the contextual menu item is Add VRRP IPv6 Instance, and the Select VRRP IPv6 Instance form opens.

The Select VRRP Instance form opens.

5

Choose a VRRP instance and click OK. The Select VRRP Instance (or Select VRRP IPv6 Instance) form closes and the Virtual Router (Edit) form reappears with the interface information displayed.

END OF STEPS

37.6 To modify a VR or VRRP instance

37.6.1 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Choose a VR and click Properties. The Virtual Router (Edit) form opens.

3

Configure the VR parameters.

To configure parameters for a VRRP instance, click on the Components tab, select and right-click on the VRRP instance, and choose Properties.

You can also choose another configuration option to configure the NE associated with the VRRP instance or to configure the interface associated with the VRRP instance.

4

Click Properties in the Instance panel. The VRRP Instance (Edit) form opens.

5

The following tabs list the VRRP instance properties that you can select and configure:

- General—associations between the VRRP instance and NEs, network interfaces, and policies (non-owner VRRP instances only).
- Behavior—a virtual MAC address that must be the same for all participating VRRP instances
- Authentication—enables authentication of VRRP advertisements among participating VRRP instances in the VR. For more information, see 37.1.9 “VRRP message authentication” (p. 1282) in 37.1 “VRRP” (p. 1279) .
- Backup Addresses—lists backup addresses, whose configuration options you can access by selecting a site and clicking Properties. For more information, see 37.1.8 “Backup addresses” (p. 1282) in 37.1 “VRRP” (p. 1279) .
- VR Instances—lists VRs, whose configuration options you can access by selecting a VR and clicking Properties. For more information, see 37.1.4 “Owner and non-owner VRRP instances” (p. 1281) in 37.1 “VRRP” (p. 1279) .

6

Modify the parameters for the VRRP instance as required.

7

Save your changes. The VRRP Instance (Edit) form closes and the VR Instance (Edit) form reappears.

8

Close the VR Instance (Edit) form.

END OF STEPS

37.7 To view the status of a VR

37.7.1 Purpose

The status of a VR informs you of potential problems, such as problems involving VR preconditions and operational states. You can view the following virtual status information:

- Aggregated Operational State—indicates the collective operational states of the VRRP instances in the VR, such as whether all instances are up, all instances are down, or one or more is down
- Number Of VRRP Instances—the number of owner and non-owner VRRP instances in the current VR
- Multiple Owners configured—one owner for each VR
- VR Instance(s) Down—one or more VRs are not working
- Backup Address Mismatch—the backup address for a VRRP instance does not match the primary IP address of the owner, and so does not match the VR IP address
- Only one instance configured—only one VR instance exists, either a master or backup; both are required
- Subnet Mismatch—the IP addresses of the owner and non-owner routers do not belong to the same subnet
- No Owner configured—an IP address owner is not assigned to the current VR

37.7.2 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Configure the filter criteria. A list of VRs is displayed at the bottom of the Manage VRRP Virtual Routers form.

3

Choose a VR and click Properties. The Virtual Router (Edit) form opens. Status information is in the Status panel.

4

Close the Virtual Router (Edit) form.

END OF STEPS

37.8 To delete a VRRP instance

37.8.1 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Configure the filter criteria. A list of VRs is displayed at the bottom of the Manage VRRP Virtual Routers form.

3

Choose a VR and click Properties. The Virtual Router (Edit) form opens.

4

Click on the Components tab.

5

Right-click on a VR instance and choose Delete. A dialog box appears.

6

Click Yes to confirm your action. The VR is deleted and removed from the VR instances tree.

7

Close the Virtual Router (Edit) form.

END OF STEPS

37.9 To delete a VR

37.9.1 Steps

1

Choose Manage→Networking→VRRP Virtual Routers from the NFM-P main menu. The Manage VRRP Virtual Routers form opens.

2

Configure the filter criteria. A list of VRs is displayed at the bottom of the Manage VRRP Virtual Routers form.

3

Choose a VR and click Delete. A dialog box appears.

4

Confirm your action. The VR is removed from the list.

5

Close the Manage VRRP Virtual Routers form.

END OF STEPS

38 APS

38.1 Overview

38.1.1 Purpose

APS protects SONET/SDH lines from linear bidirectional failures. The NEs in a SONET/SDH network constantly monitor the health of the network. When a failure is detected, the network proceeds to transfer, or switch over, live traffic from the active, or working line, to the standby, or protection line. The transfer occurs quickly to minimize traffic loss. The traffic remains on the protection line until the fault on the working line clears, at which time the traffic can optionally revert to the working line.

In a 1+1 APS architecture, the active OC-*n* signal is transmitted to both the working and protection ports, so the same payloads are transmitted identically to the working and protection ports in the egress direction. The working and protection signals are selected independently in the ingress direction.

38.1.2 Contents

APS overview

38.2 APS overview

38.2.1 Overview

CAUTION

Service Disruption

An NE transmits data only on the working line.

For a single failure, this configuration can cause a delay of up to 100 ms during an APS switch.

On the NFM-P, the main APS elements are the following:

- APS groups

  An APS group is an aggregation of one or two SONET/SDH ports and the associated logical ports, known as APS channels. The first port is known as the working channel and must be configured in the APS group. The second port is the protection channel, which is optionally configurable. In a 1+1 optimized APS group, the protection channel is replaced with a second working channel.

- APS channels

  APS channels are logical and model the physical ports. The set of associated data for an APS channel includes the corresponding physical port identifier, the role, which is working or protection, and the name of the APS group to which the port belongs. If a protection channel is configured, it carries the traffic of the failed working channel. If the protection channel fails, the working channel carries the traffic.

- APS bundles

  APS bundles protect multilink PPP and IMA bundles on channelized ASAP MDAs. All the members of a working or protection APS bundle must belong to the same working or protection line of the APS group.

- APS common port configurations

  All SONET/SDH port parameters on the working and protection channels in an APS group must be identical except for the following:
  - Clock Source
  - Loopback
  - Report Alarms
  - BER Signal Degradation Threshold
  - BER Signal Failure Threshold
  - SONET Section Trace Mode

- APS commands

  Each APS channel has one APS operational command associated with it. The APS operational command is determined by the last command that is issued to both APS channels and the signal condition of the working and protection ports. APS commands affect the APS operational states but do not affect the configuration. The operational state does not persist through an NE restart.

Table 38-1, “Events affecting 1+1 protection” (p. 1293) The following table describes the events that affect 1+1 protection and their relative priority.

Table 38-1 Events affecting 1+1 protection

Priority	Event	Event type
1	Lockout of protection	User initiated
2	Signal failure on protection line	Automatically initiated
3	Forced switch	User initiated
4	Signal failure on working line	Automatically initiated
5	Signal degradation	Automatically initiated
6	Manual switch	User initiated
7	WTR time (revertive switching only)	User initiated (state request)
8	No reversion (non-revertive switching only)	User initiated (state request)

Consider the following when configuring APS:

- A port can belong to only one APS group.
- Two ports that belong to the same APS group must be of the same port type and have the same traffic speed.
- An APS group has one set of APS channels. The set can contain one or two APS channels.
- An APS channel can belong to only one APS group.
- A SONET/SDH port can either belong to only one APS channel.

38.2.2 Bidirectional mode

In 1+1 system bidirectional mode, a signal failure in either direction causes both the near-end and far-end NEs to switch to the protection channels. The highest-priority local request is compared with a remote request; the request that has the greater priority is selected. See Table 38-1, “Events affecting 1+1 protection” (p. 1294) for the list of events that affect 1+1 protection.

38.2.3 Unidirectional mode

In a 1+1 system unidirectional mode, the working interface switches to the protection interface only for the direction in which a signal failure occurs. For example, if there is a signal failure in the transmit direction, the working interface switches to the protection interface for transmission but not for the receipt of data.

38.3 Switching modes

38.3.1 Overview

The following 1+1 system switching modes are available:

- non-revertive (default)

- revertive

In non-revertive switching, a switch to the protection channel is maintained even after the working line has recovered from a failure or a manual switch is cleared.

In revertive switching, the traffic is switched back to the working channel after the working line has recovered from a failure or a manual switch is cleared.

For revertive switching, you can define a period of time that the system must wait before it can restore traffic from the protection line to the working line. This delay, or WTR time, prevents frequent automatic switches from occurring as a result of intermittent failures.

In case of failure on both the working and protection lines, the line that has the less severe error remains active. If there is a signal degradation on both lines, the active line that failed last remains active. If there is a signal failure on both lines, the working line always remains active because a signal failure on the protection line is a higher priority than a signal failure on the working line.

38.4 MLPPP

38.4.1 Overview

MLPPP provides a way to distribute data across multiple links within an MLPPP APS bundle to achieve high bandwidth. MLPPP allows for a single frame to be fragmented and transmitted across multiple links. This reduces latency and allows for a higher maximum received recovery unit, or MRU.

MLPPP is supported in MC APS groups. See 38.7 “MC APS” (p. 1298) for more information about MC APS.

38.4.2 Multiclass MLPPP

Multiclass MLPPP is an extension of the MLPPP standard which allows multiple classes of service to be transmitted over an MLPPP bundle.

Multiclass MLPPP changes the MLPPP header to include either two or four class bits to allow for up to either four or 16 classes of service. This allows multiple classes of services over a single MLPPP connection. The highest priority traffic is transmitted over the MLPPP bundle with minimal delay, regardless of the order in which packets are received.

Multiclass MLPPP is useful in mobile network deployments where multiple types of traffic, each with its own priority level, travel across a single MLPPP link bundle between the base station router and the aggregation router in the point of presence (POP) mobile operator.

Note: Multiclass MLPPP allows for several classes of services to be transmitted over an MLPPP bundle. Link fragmentation and interleaving, however, allows for only two classes of service to be transmitted. Multiclass MLPPP and link fragmentation and interleaving are mutually exclusive.

38.5 APS port configurations

38.5.1 Overview

A 7750 SR network or access port can be connected to another 7750 SR with both the working and protection channels on different IOMs in a single NE, or with the working channel on one NE and the protection channel on another.

A 7705 SAR access port can be connected to another 7705 SAR with both the working and protection channels on different MDAs on a single NE, or with the working channel on one NE and the protection channel on another.

For the 7705 SAR, for clear channels sts3 and sts12, only network mode for APS groups is supported.

Note: Mirroring parameters configured on a specific port or service are maintained during an APS failover.

All SONET/SDH MDAs support APS functionality. Table 38-2, “APS port configurations” (p. 1295)The following table lists the possible port pairings to provide APS protection. Both ports must be of the same type and have the same traffic speed.

Table 38-2 APS port configurations

MDA type—Working channel	MDA type—Protection channel
16 × OC12/OC3 SFP	8 × OC12/OC3 SFP or 16 × OC12/OC3 SFP
8 × OC12/OC3 SFP	8 × OC12/OC3 SFP or 16 × OC12/OC3 SFP
16 × OC3 SFP	8 × OC3 SFP or 16 × OC3 SFP
8 × OC3 SFP	8 × OC3 SFP or 16 × OC3 SFP
4 × OC48 SFP	2 × OC48 SFP or 4 × OC48 SFP
2 × OC48 SFP	2 × OC48 SFP or 4 × OC48 SFP
1 × OC192	1 × OC192
16 × ATM OC3 SFP	16 × ATM OC3 SFP
4 × ATM OC12/OC3 SFP	4 × ATM OC12/OC3 SFP
4 × Channelized OC3 ASAP	4 × Channelized OC3 ASAP
4 × Channelized OC3/OC12 ASAP SFP	4 × Channelized OC3/OC12 ASAP SFP
1 × Channelized OC12 ASAP	1 × Channelized OC12 ASAP
1 × Channelized OC3 CES	4 × Channelized OC3 CES
4 × Channelized OC3 CES	1 × Channelized OC3 CES
1 × Channelized OC12 CES	1 × Channelized OC12 CES

Note: The working and protection channels for the following MDA port pairs are set to the same traffic speed based on the APS group speed configuration:

- 16 × OC12/OC3 SFP and 16 × OC12/OC3 SFP

- 16 × OC12/OC3 SFP and 8 × OC12/OC3 SFP
- 8 × OC12/OC3 SFP and 16 × OC12/OC3 SFP
- 8 × OC12/OC3 SFP and 8 × OC12/OC3 SFP
- 4 × ATM OC12/OC3 SFP and 4 × ATM OC12/OC3 SFP

38.6 SC APS

38.6.1 Overview

1+1 APS can be implemented on a port-by-port basis. If all ports on an MDA or IOM need to be protected, the ports must be individually configured.

The working and protection lines are capable of being connected to:

- two ports on the same MDA (not applicable to the 7705 SAR)
- two ports on different MDAs, and the MDAs are on the same IOM
- two ports on different MDAs on different IOMs

If the working channel and protection channel are on the same MDA, protection is limited to the physical port and the media that connect the two NEs. If different IOMs are used, protection extends to failure of each IOM.

Note: The working and protection channels should be configured on different MDAs but on the same IOM card for the 7705 SAR-8 and 7705 SAR-18.

Working and protection lines can be connected to a 7705 SAR-8, 7705 SAR-8v2, 7705 SAR-18, 7450 ESS, or 7750 SR, and serve as an access port that provides one or more services to the NE. The access port can be a single channel or multiple channels; each channel must support PPP. In the case of the ATM MDA, each channel must support ATM.

The end NE transmits a valid data signal to both the working and protection lines. The signal on the protection line is ignored until the working channel fails or degrades to the degree that requires a switchover to the protection channel. When the switchover occurs, all services, including all service QoS and filter policies, are activated on the protection channel.

The working channel on a 7705 SAR-8, 7705 SAR-8v2, 7705 SAR-18, 7750 SR, or 7450 ESS must connect to the working channel on a peer NE, and the protection channel on a 7705 SAR-8, 7705 SAR-8v2, 7705 SAR-18, 7750 SR, or7450 ESS must connect to the protection channel on a peer NE.

38.6.2 1+1 optimized

You can configure an SC APS group with a 1+1 optimized configuration. A 1+1 optimized APS group is optimized for networks which rely primarily on 1+1 bidirectional switching. A 1+1 optimized APS group is configured with two working channels, as opposed to one working channel and one protection channel. In a 1+1 optimized configuration, one working channel acts as the primary channel, through which the traffic will always default. A 1+1 optimized APS group always operates in bidirectional and nonrevertive modes.

38.7 MC APS

38.7.1 Overview

You can use APS to protect against NE failure by configuring the working channel of an APS group on one 7705 SAR-8, 7705 SAR-8v2, 7705 SAR-18, or 7750 SR, and configuring the protection channel of the same APS group on a different 7705 SAR-8, 7705 SAR-8v2, 7705 SAR-18, or 7750 SR. The two NEs connect using an IP link that is used to establish a signaling path between them.

The working channel on the near-end NE must connect to the working channel on a peer NE, and the protection channel on the far-end NE must connect to the protection channel on a peer NE.

When an MC APS group is configured, the NFM-P automatically creates a container which aggregates the MC APS group configurations of each NE.

Multi-chassis APS configuration is supported for ATM clear channel interfaces.

38.8 APS on channelized ASAP MDAs

38.8.1 Overview

You can protect a channelized SONET/SDH port on a channelized ASAP MDA with a protection port of the same speed on a different channelized ASAP MDA in the same NE. The APS configuration on a channelized ASAP MDA provides protection against a port, MDA, or IOM failure. All SONET/SDH paths and TDM channels in a SONET/SDH port are protected.

Consider the following when you configure APS protection on a deep channel on a channelized ASAP MDA:

- Both SONET and SDH channels are supported.
- Up to three common configuration SONET channels can be created in an APS configuration.

38.9 APS on channelized CES MDAs

38.9.1 Overview

You can protect a channelized SONET/SDH port on a channelized CES MDA with a protection port of the same speed on a different channelized CES MDA in the same NE. The APS configuration on a channelized CES MDA provides protection against a port, MDA, or IOM failure. All SONET/SDH paths and TDM channels in a SONET/SDH port are protected.

Consider the following when you configure APS protection on a deep channel on a channelized CES MDA:

- Both SONET and SDH channels are supported.
- Up to three common configuration SONET channels can be created in an APS configuration.

38.10 APS on multilink bundles

38.10.1 Overview

APS on multilink bundles consists of a working and protection bundle which provide bidirectional APS protection to each other. The members of a working or protection multilink bundle must belong to the same working or protection line of the APS group. User traffic is not sent on the protection line.

The NFM-P supports APS on MLPPP and IMA bundles. In IMA bundles, IMA cells are sent on the protection line as a keep-alive signal during an APS switchover.

APS 1+1 configuration on the channelized MDA provides protection to all the channels on the protected SONET/SDH port and to all of the multilink bundles with member that links reside on the protected SONET/ SDH port.

38.10.2 APS bundles on multiple NEs

You can configure APS bundles to provide bidirectional APS protection across multiple NEs. MC APS bundles are supported on the 7705 SAR and 7750 SR.

You can use APS to protect against failure by configuring the working bundle on one NE, and configuring the protection bundle of the same APS bundle group on a different NE. The two NEs connect to each other with an IP link that is used to establish a signaling path between them.

38.11 1+1 APS configuration example

38.11.1 Overview

Figure 38-1, “Normal operations between two 7750 SRs” (p. 1299)and Figure 38-4, “Normal operations resume on the protection line” (p. 1301) show an example of 1+1 APS for two 7750 SRs that are configured for 1+1 APS in bidirectional and non-revertive modes.

Figure 38-1, “Normal operations between two 7750 SRs” (p. 1299) shows normal operations between two 7750 SRs. There are no faults on the working line, and the protection line is not in use.

Figure 38-1 Normal operations between two 7750 SRs

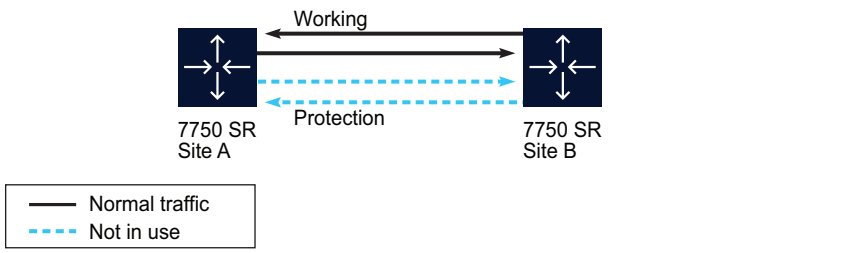

The working line degrades in the direction from site A to site B. Site B detects the fault and notifies site A of the fault using the protection line. The following figure shows the fault on the working line and site B notifying site A of the fault.

Figure 38-2 Site B detects the fault and notifies site A

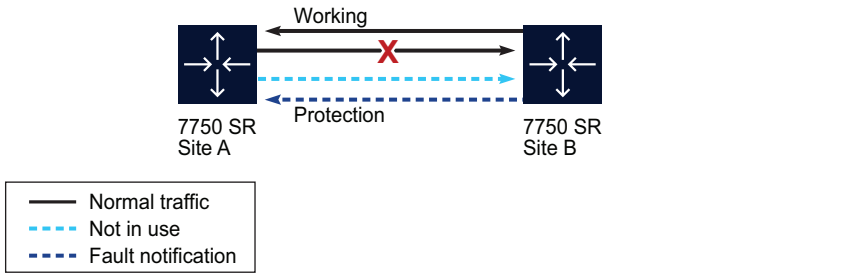

Site B automatically switches to the protection line. Site A receives the fault notification from site B and detects the fault on the working line. Site A acknowledges the fault and notifies site B that it is switching to the protection line. The following figure shows site B switching to the protection line and site A acknowledging the fault.

Figure 38-3 Site B switches to the protection line and site A acknowledges the fault

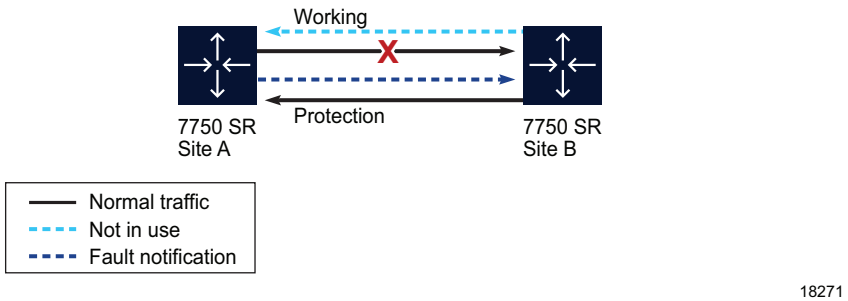

Site B receives the notification from site A and site A automatically switches to the protection line. The following figure shows normal operations resuming on the protection line between site A and site B. The protection line becomes and remains the working line since 1+1 APS is configured for non-revertive switching in this example.

Figure 38-4 Normal operations resume on the protection line

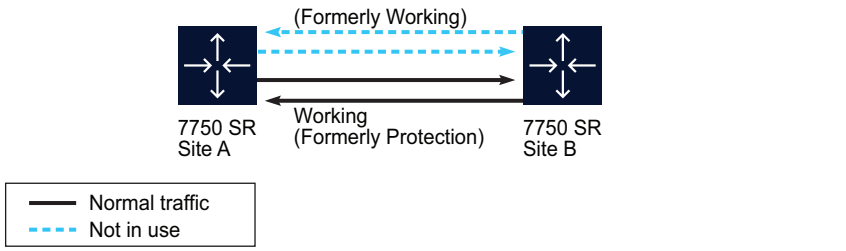

38.12 Configuring SAPs on APS-protected ports

38.12.1 Overview

You can use APS-protected ports to create SAPs for services. The SAP list of a service creation form displays the APS-group SONET channels with all other SONET channels, but uses a different port ID format. An example of an APS group port ID is Channel aps-1.sts3-1. To create a SAP that uses an APS group, you must first configure the SONET channels in the APS group.

APS management procedures

38.13 Workflow to manage APS

38.13.1 Stages

1

As required, create an SC APS group or MC APS group and configure the various channel types supported by APS groups. Configure the operational state for each APS channel. See 38.14 “To create an SC APS group” (p. 1302) and 38.15 “To create an MC APS group” (p. 1305) for more information.

2

As required, create an SC APS IMA or MLPPP bundle or MC APS MLPPP bundle. See 38.16 “To create an SC APS IMA or MLPPP bundle” (p. 1307) and 38.17 “To create an MC APS MLPPP bundle” (p. 1311) for more information.

3

As required, change the operational state of an APS channel; see 38.18 “To change the operational state of an SC APS channel” (p. 1313) .

4

As required, delete the APS group or bundle.

See the following procedures:

- See 38.19 “To delete an SC APS group” (p. 1314) to delete a SC APS group
- See 38.20 “To delete an SC APS bundle” (p. 1314) to delete a SC APS bundle
- See 38.21 “To delete an MC APS group or bundle” (p. 1315) to delete an MC APS group or bundle

38.14 To create an SC APS group

38.14.1 Purpose

Perform this procedure to create an SC APS group and configure the various channel types supported by APS groups including:

- APS working channels
- APS protection channels
- APS SONET channels
- SONET clear channels
- The maximum number of channels on ports that support multiple sub-channels

Note: When a multiservice site is configured on an IOM, the working and protection ports

must be configured on the same IOM.

38.14.2 Steps

1

On the equipment tree, navigate to the Network Element and expand the Shelf icon.

2

Right-click on the APS Groups object and choose Create APS Group. The SC APS Group (Create) form opens with the General tab displayed.

3

Configure the Description parameter.

4

Click on the APS Group tab and configure the required parameters.

5

Click on the States tab and configure the Administrative State parameter.

Note: When a working channel or protection channel is added to an APS group, the channel inherits the current administrative state of the APS group. After the channel is added to the APS group, the administrative state of the physical port can be changed independently.

6

Click Apply. An SC APS Group icon appears under the APS Groups object in the navigation tree.

7

Click on the APS Group tab and configure the required parameters.

8

To configure the working and protection channels for the APS group:

Note: Prior to completing this step, your NE device must be configured with a card that supports APS ports. See Table 38-2, “APS port configurations” (p. 1296) for a list of cards that support APS ports.
To create an APS working channel or APS protection channel, you can also right click on the SC APS Group icon under the APS Groups object in the navigation tree and choose

Create APS Working Channel or Create APS Protection Channel from the contextual menu. The Create APS Protection Channel option is available only when an APS Working Channel already exists.

The working and protection channels should be configured on different MDAs but on the same IOM card for the 7705 SAR-8 and 7705 SAR-18.

1. Click on the APS Channels tab and click Create. The APS Channel (Create) form opens.
2. Configure the Channel Role parameter, if applicable.

   Note:

   If you are creating a 1+1 SC APS group, you must create the working channel before you create the protection channel.
3. Select a port and click OK. The channel is listed on the SC APS Group (Edit) form.
4. Repeat 2 and 3 to add the protection channel or second working channel.
5. Click Apply.

9

To configure the operational state of the working and protection channels:

Use the following steps:

1. Choose a channel and click Properties. The APS Channel - *Type* (Edit) form opens with the APS Channel tab displayed.
2. Configure the Command Switch parameter and click OK.

10

To configure the APS SONET channels for the APS group:

Use the following steps:

1. Right-click on the SC APS Group icon in the navigation tree and choose Create APS SONET Channel. The Sts*n* SONET Channel (Create) form opens with the General tab displayed.
2. Configure the channels as described in 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

11

Perform one of the following for the APS group:

a. To create individual SONET clear channels:

   Use the following steps:

   1. Right-click on the APS Common Config SONET channel icon and choose Create Channel. The DSO Channel Group, (Create) form opens.
   2. Configure the channel as described in 16.60 “To configure SONET clear channels” (p. 653) .

b. To create the maximum number of channels on ports that support multiple sub-channels:

Use the following steps:

1. Right-click on the APS Common Config SONET channel icon and choose Create Maximum # of Channels. The Create Maximum # of Channels form opens.
2. Configure the channels as described in 16.61 “To perform a bulk channel creation on ports that support multiple sub-channels” (p. 654) .

12

Save your changes and close the form.

END OF STEPS

38.15 To create an MC APS group

38.15.1 Purpose

Perform this procedure to create an MC APS group and configure the various channel types supported by APS groups including:

- APS working channels
- APS protection channels
- APS SONET channels
- SONET clear channels
- The maximum number of channels on ports that support multiple sub-channels

38.15.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC APS Group (APS) from the object drop-down menu and click Create. The MC APS Group (Create) form opens.

3

Configure the Group Number parameter.

4

Select a Site Id and Network Interface Address in the First Element and Second Element panels.

Note: Only numbered IPv4 interfaces are listed.
The default interface is the NE system interface.

5

Click Apply. An MC APS Group Member icon appears under the APS Groups object in the navigation tree of the selected First Element and Second Element NEs.

6

Click on the Members tab to configure each MC APS group member.

7

Select a group member in the list and click Properties. The MC APS Group Member (Edit) form opens with the General tab displayed.

8

Configure the Description parameter.

9

Click on the APS Group tab and configure the required parameters.

Note: The Wait To Restore (seconds) parameter is configurable when the Revert Mode parameter is set to revertive.

10

Click on the States tab and configure the Administrative State parameter.

Note: When a working channel or protection channel is added to the APS group, the channel inherits the current administrative state of the APS group. After the channel is added to the APS group, the administrative state of the physical port can be changed independently.

11

Click on the Multi Chassis tab and configure the required parameters.

12

To configure the working and protection channels for the APS group:

Note: Prior to completing this step, your NE device must be configured with a card that supports APS ports. See Table 38-2, “APS port configurations” (p. 1296) for a list of cards that support APS ports.

Use the following steps:

1. Click on the APS Channels tab and click Create. The APS Channel (Create) form opens.
2. Configure the Channel Role parameter.
3. Select a port and click OK. The channel is listed on the MC APS Group (Edit) form.
4. Repeat 2 and 3 to add the protection channel or second working channel.
5. Click Apply.

13

To configure the operational state of the working and protection channels:

Use the following steps:

1. Choose a channel from the list and click Properties. The APS Channel - *Type* (Edit) form opens with the APS Channel tab displayed.
2. Configure the Command Switch parameter and click OK. The APS Channel - *Type* (Edit) form closes.

14

To configure the APS SONET channels for the APS group:

Use the following steps:

1. Right-click on the MC APS Group icon in the navigation tree and choose Create APS SONET Channel. The Channel Type Selection form opens.
2. Choose a channel type and click OK. The Sts*n* SONET channel (Create) form opens.
3. Configure the channels as described in 16.64 “To create VT15 (TU11) or VT2 (TU12) sub-channels” (p. 660) .

15

Perform one of the following for the APS group:

a. To create individual SONET clear channels:

   Use the following steps:

   1. Right-click on the APS Common Config SONET channel icon and choose Create Channel. The DSO Channel Group, (Create) form opens.
   2. Configure the channel as described in 16.60 “To configure SONET clear channels” (p. 653) .

b. To create the maximum number of channels on ports that support multiple sub-channels:

   Use the following steps:

   1. Right-click on the APS Common Config SONET channel icon and choose Create Maximum # of Channels. The Create Maximum # of Channels form opens.
   2. Configure the channels as described in 16.61 “To perform a bulk channel creation on ports that support multiple sub-channels” (p. 654) .

16

Save your changes and close the form.

END OF STEPS

38.16 To create an SC APS IMA or MLPPP bundle

38.16.1 Purpose

Perform this procedure to create an APS bundle using IMA or MLPPP links on one NE. Consider the following when you add a DS0 channel that is a member of an APS group to an APS bundle:

- An APS channel can be a member of an APS bundle only, not of an IMA or multilink PPP bundle.
- All members of a working bundle must belong to a working channel of the same APS group.
- All members of a protection bundle must belong to a protection channel of the same APS group.
- You must create the working bundle before you can add an APS member to the APS bundle. If the member has a protection port, you must create the protection bundle before you can add the member to the APS bundle.
- You cannot delete the working port member of an APS bundle until you remove the member from the APS bundle.

The rules for adding a member to a bundle that is part of an APS bundle also apply to adding a member to a non-APS protected multilink bundle. See 15.90 “To create an MLPPP bundle” (p. 551) and 15.88 “To create an IMA group bundle” (p. 550) for more information.

The following restrictions apply when adding a DS0 channel to an IMA bundle:

- The Clock Source parameter must be set to Node Timed.
- The encapsulation type of the DS0 channel must be ATM.

The following restrictions apply when adding a DS0 channel to a PPP bundle:

- All time slots on the DS0 channel must be selected.
- The encapsulation type of the DS0 channel must be IPCP.

38.16.2 Steps

1

On the equipment tree, right-click on the device where you want to configure an APS bundle and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

3

Right-click on the APS Bundles object and choose Create Bundle. The APS Bundle Display form opens.

4

Configure the Bundle ID and Description parameters and click Next.

5

Set the Bundle Type parameter to IMA Group or PPP and configure the applicable parameters.

6

Click Finish and then click Close. The APS Bundle Display form closes.

7

If the Bundle Type parameter is set to PPP in Step 5 , configure the MLPPP bundle for multiclass service transmission.

Note: Consider the following when you configure MC MLPPP.

- MC MLPPP is supported only on channelized ASAP MDAs on a 7705 SAR or 7750 SR.
- You must configure MC MLPPP before you add bundle members.
- MC MLPPP is configured on the main APS bundle and the parameters are propagated to the working and protection bundles.
- You can only apply QoS profiles to an MC MLPPP bundle when the Class Count parameter is set to 4.

Use the following steps:

1. Right-click on the SC APS Bundle object in the navigation tree and choose Properties. The APS Bundles (Edit) form opens.
2. Click on the MLPPP tab and configure the required parameters.
3. Select a QoS profile in the MLPPP Ingress QoS Profile or MLPPP Egress QoS Profile panel, if required.

8

Click on the States tab and configure the Administrative State parameter.

Note: You must configure the End Point ID parameter on the APS PPP group before you set the Administrative State parameter to Up on the working and protection bundles.

9

Save your changes.

10

Configure the APS working and protection bundles:

Note: You must create the working bundle before you create the protection bundle.

Use the following steps:

1. Right-click on the SC APS Bundle icon in the navigation tree and choose Create APS Working Bundle. The Create APS Working/Protecting Bundle form opens.

2. Configure the required parameters and click Next to view the configured bundle parameters.
3. Click Finish and then click Close. The Create APS Working/Protecting Bundle form closes.
4. Right-click on the SC APS Bundle icon in the navigation tree and choose Create APS Protection Bundle. The Create Working/Protecting Bundle form opens.
5. Repeat 2 and 3 to create the APS protection bundle.

11

To configure the administrative state of the APS working and protection bundles:

Note: You must configure the End Point ID parameter on the APS PPP group before you change the working or protection bundle state.

Use the following steps:

1. Right-click on the working bundle object below the SC APS Bundle object in the navigation tree and choose Properties. The Multilink Bundle (Edit) form opens.
2. Click on the States tab and configure the Administrative State parameter.
3. Click OK.
4. Right-click on the protection bundle object below the SC APS Bundle object in the navigation tree and choose Properties. The Multilink Bundle (Edit) form opens.
5. Repeat 2 and 3 to configure the administrative state of the protection bundle.

12

Add members to the APS bundle.

Use the following steps:

1. Right-click on the SC APS Bundle icon in the navigation tree and choose Create Bundle Members. The Add Bundle Member form opens.
2. Click Next to add DS0 channel groups to the bundle. The Select Channels form opens.
3. Select a compatible channel from the list.

   Note:

   When initially adding bundle members, you must first add one channel group and then repeat the procedure to select up to seven additional channel groups.

   Only compatible channels are listed. If you want to configure and use a channel that is currently incompatible, you can click on the Back button, deselect the Show Only Compatible Channels parameter, click on the Next button, select the channel, click on the Properties button.

   The channel group with the lowest Port ID is the primary member for the bundle. The Encap Type of the primary member is used for all other members. When adding members to a bundle at a later time, the Encap Type must be the same as the primary member Encap Type for the member to be compatible.
4. Click Finish. The Add Bundle Member form closes.

5. Click Apply and then click Add. The Add Bundle Member form opens.
6. Configure the Show Only Compatible Channels parameter and click Next. The Select Channels form opens.
7. Select up to seven additional channels from the list.
8. Click Finish.

13

Save your changes and close the form.

END OF STEPS

38.17 To create an MC APS MLPPP bundle

38.17.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC APS Bundles (Bundle) and click Create. The MC APS Bundles (Create) form opens.

3

Configure the Bundle Number parameter.

4

Select a First Network Element and Second Network Element and click OK. The MC APS Bundle object appears in the navigation tree under the APS Bundles object for each specified NE.

5

Close the Manage Node Redundancy form.

6

To configure the MLPPP bundle for multiclass service transmission:

Note: Consider the following when you configure MC MLPPP.

- You must configure MC MLPPP before you add bundle members.

- MC MLPPP is configured on the main APS bundle and the parameters are propagated to the working and protection bundles.

Use the following steps:

1. Right-click on the new MC APS Bundle object in the navigation tree and choose Properties. The APS Bundles (Edit) form opens.
2. Click on the MLPPP tab and configure the required parameters.
3. Save your changes.

7

Click on the States tab and configure the Administrative State parameter.

Note: You must configure the End Point ID parameter on the APS PPP group before you set the Administrative State parameter to Up on the working and protection bundles.

8

Save your changes and close the form.

9

Configure the MC APS working and protection bundles.

Use the following steps:

1. Right-click on the MC APS Bundle object in the navigation tree and choose Properties. The APS Bundles (Edit) form opens.
2. Click on the Protection and Working Bundles tab and click Create. The Create APS Working/Protecting Bundle form opens.
3. Configure the required parameters.
4. Select a daughter card and click Next.
5. Configure the required parameters.

   Note:

   The working bundle must be on the same NE as the APS working channel and the protection bundle must be on the same NE as the APS protection channel.
6. Click Finish and then click Close. The Create APS Working/Protecting Bundle form closes.
7. Close the APS Bundles (Edit) form.
8. Repeat 1 to 7 to configure the MC APS protection bundle.

10

Configure the administrative state of the MC APS working and protection bundles.

Use the following steps:

1. Right-click on the working bundle object in the navigation tree and choose Properties. The Multilink Bundle (Edit) form opens.

2. Click on the States tab and configure the Administrative State parameter.
3. Save your changes.
4. Right-click on the protection bundle object in the navigation tree and choose Properties. The Multilink Bundle (Edit) form opens.
5. Repeat 2 and 3 .

11

Add multilink bundle members.

Use the following steps:

1. Right-click on the MC APS Bundle object on either NE in the navigation tree and choose Properties. The APS Bundles (Edit) form opens.
2. Click on the Bundle Members tab and click Add. The Add Bundle Member form opens.
3. Configure the Show Only Compatible Channels parameter and click Next. The Select Channels form opens.
4. Select a channel in the list.

   Note:

   When initially adding bundle members, you must first add one channel group and then repeat the procedure to select up to seven additional channel groups.

   Only compatible channels are listed. If you want to configure and use a channel that is currently incompatible, you can click Back, deselect the Show Only Compatible Channels parameter, click Next, select the channel, and click Properties.

   The channel group with the lowest Port ID is the primary member for the bundle. The Encap Type of the primary member is used for all other members. When adding members to a bundle at a later time, the Encap Type must be the same as the primary member Encap Type for the member to be compatible.
5. Click Finish.
6. Click Apply and then click Add. The Add Bundle Member form opens.
7. Configure the Show Only Compatible Channels parameter and click Next. The Select Channels form opens.
8. Select up to seven additional channels from the list and click Finish.

12

Save your changes and close the form.

END OF STEPS

38.18 To change the operational state of an SC APS channel

38.18.1 Steps

1

On the equipment tree, right-click on the device where you want to modify the APS channel and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

3

Right-click on the APS Channel object. The APS Channel (Edit) form opens.

4

Configure the Command Switch parameter.

5

Save your changes and close the form.

END OF STEPS

38.19 To delete an SC APS group

38.19.1 Steps

1

On the equipment tree, right-click on the device where you want to delete an APS group and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

3

CAUTION

Service Disruption

Removing an APS group removes the working and protection channels and associated ports from the APS group, and deletes the APS group ID.

Ensure that you select the correct APS group.

Right-click on the APS Group object and choose Delete.

4

Save your changes.

END OF STEPS

38.20 To delete an SC APS bundle

38.20.1 Steps

1

On the equipment tree, right-click on the device where you want to delete an APS bundle and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

3

Right-click on the APS Bundle object and choose Delete.

4

Save your changes.

END OF STEPS

38.21 To delete an MC APS group or bundle

38.21.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Perform one of the following.

a. Choose MC APS Group (APS) from the object drop-down menu.

b. Choose MC APS Bundles (APS) from the object drop-down menu.

3

Select an entry in the list and click Delete.

4

Save your changes. The NFM-P deletes the MC APS group or bundle and the corresponding configuration on each member site.

END OF STEPS

39 lightRadio Wi-Fi

39.1 lightRadio Wi-Fi

39.1.1 Overview

To provide additional capacity and extended coverage to a network, the NFM-P supports the offloading of traffic to WLANs, also known as Wi-Fi Offload. Users within range of a WLAN access point can establish a secure connection using RADIUS authentication. WPA2 keys are provided by a RADIUS server. Data is then transmitted through a single GRE tunnel to the WLAN Gateway (GW), where a subscriber DHCP session is created. The WLAN GW supports a RADIUS proxy server, which stores all available information from the RADIUS access-accept message.

In cases where the RADIUS server instructs the WLAN GW to forward traffic to a GGSN, a mobile gateway (MG) peer is created on the NE. Read-only information relating to the MG peers can be viewed on a base routing instance configuration form, under the WLAN GW tab.

Wi-Fi access can be offered from:

- hotspots - access points in public places.
- enterprise hotspots - access points in workplaces.
- homespots - access points on wireline CPEs from residential subscribers.

Wi-Fi access can be offered using the following SSID types:

- Private SSID - used by a residential subscriber to access wireline Internet service.
- Community SSID - used by visiting subscribers to access a Wi-Fi community broadband service after IEEE 802.1X authentication.
- Public SSID - used by visiting subscribers to access a free Internet service, or to access a Wi-Fi community broadband service after portal-based authentication.

The following figure shows a typical Wi-Fi Offload scenario. The UE attempts to switch over to Wi-Fi upon detection of a network SSID. Transparent IEEE 802.1X/EAP authentication takes place between the AP and UE. An optional GTP session can be established from the WLAN GW to the GGSN/PGW. This session is established dynamically so that UE traffic can be transferred back to an MMO.

Figure 39-1 Wi-Fi Offload

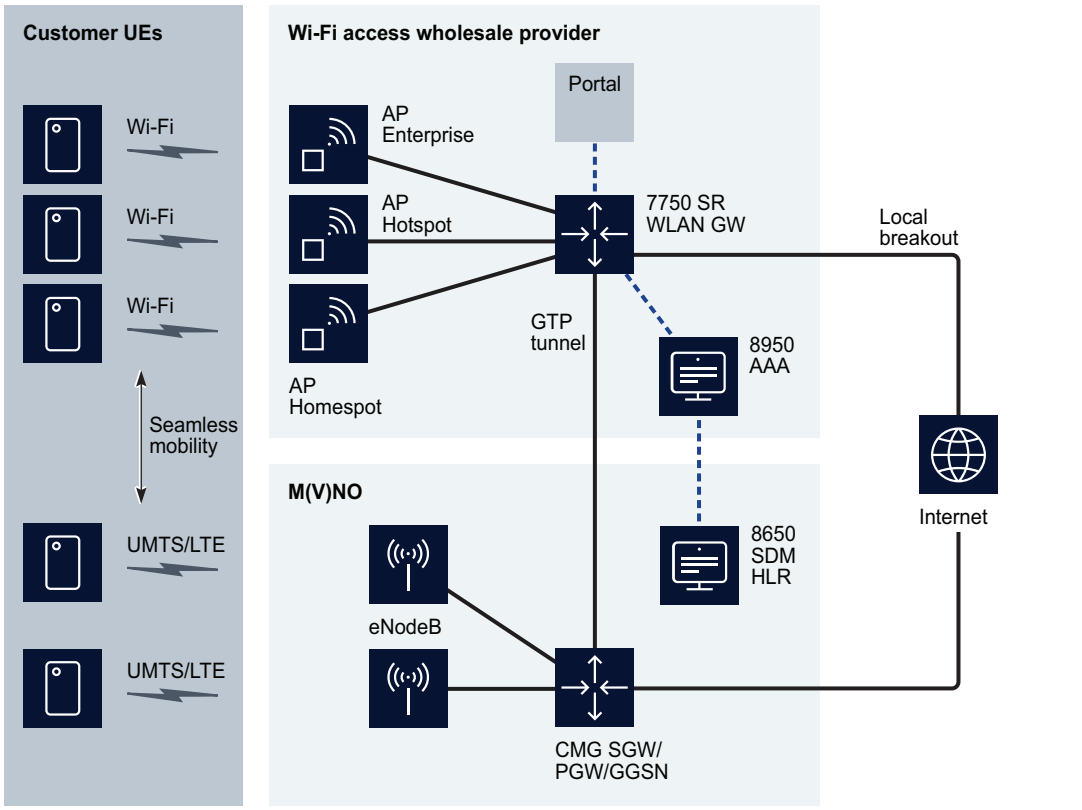

39.1.2 Migrant user handling

The NFM-P supports migrant user handling in a Wi-Fi offload application. Migrant users may have no intention of using the service, or may be invalid users who have failed authentication on the service. Both of these types of subscriber consume IP addresses and ESM resources on a WLAN GW. In a dormant subscriber handling model, the ESM subscriber is created on the IOM only when the following conditions are met:

- The subscriber is authenticated, either by RADIUS authentication or by MAC-based pre-authentication (where the RADIUS server remembers a previously authenticated UE).

- The subscriber is determined to be active by transmitting data beyond the initial DHCP request (data-triggered host creation).

Migrant user support works in conjunction with L2-aware NAT, where each UE is allocated the same inside IP address, and is assigned an outside address based on its subscriber ID (MAC address). In the case of L2-aware NAT and soft GRE, the address is assigned via DHCP, and each user is assigned the same lease. Migrant user handling is configured on an IES or VPRN soft GRE group interface, and includes the following functions:

- Data-triggered authentication and ESM host creation for ISA users.
- ESM host creation based on RADIUS COA request for portal-authenticated users.
- Migrant user DHCP performance.

  IP addresses are assigned on the ISA for L2-aware NAT (without ESM host). CPM-based DHCP is employed on a soft-GRE group interface without L2-aware NAT.

  An HTTP-redirect function is configured on the ISA.

39.1.3 Distributed subscriber management

In a distributed subscriber management (DSM) scenario, the subscriber is instantiated as a distributed subscriber on the ISA, based on a RADIUS VSA during initial authentication. DSM works in conjunction with an L2-aware NAT configuration, under which the same inside IP address is allocated to all subscribers on the same group interface (WLAN GW tunnel endpoint), or on the same VLAN tag (corresponding to SSID). Upstream and downstream forwarding is implemented without the need for an ESM host.

The L2-aware NAT is configured on the anchor ISA, and packets do not require an extra pass through the switch fabric. The NAT configuration must include RADIUS accounting.

In a DSM configuration, traffic forwarding is performed without an IOM-based ESM host. Upstream traffic (AP to network) reaches the anchor ISA as usual via B-VPLS. On the anchor ISA, the subscriber state (and L2-aware NAT state) exist prior to DHCP. Downstream traffic (network to AP) passes through L2-aware NAT on the anchor ISA, resulting in an ISA-based subscriber lookup on the anchor ISA, based on the subscriber MAC address. Traffic does not require an additional switch fabric hop to perform the NAT function. Figure 39-2, “Distributed subscriber management” (p. 1320) The following figure shows a DSM configuration.

Figure 39-2 Distributed subscriber management

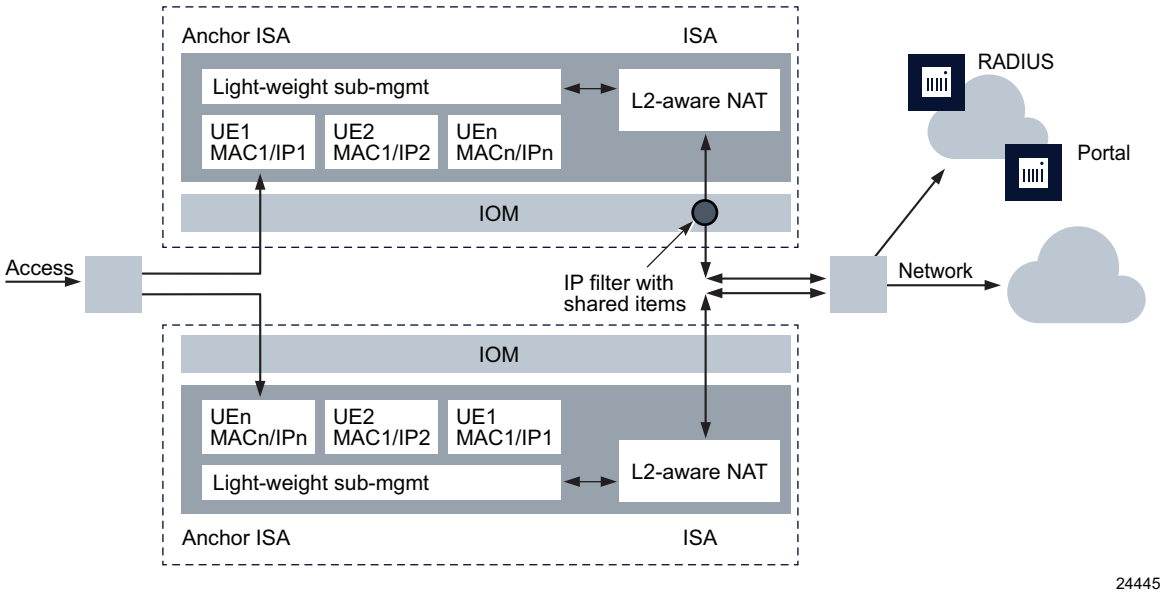

39.1.4 IPv6 pool manager

In order to support IPv6 migrant users and DSM users, an IPv6 pool manager can be configured on an IES or VPRN subscriber interface. The IPv6 pool manager uses an internal DHCPv6 client to delegate IA_NA or SLAAC prefixes to ISAs. Subnets are added to the FIB with respective ISAs as the next hop, allowing downstream forwarding without an ESM host. Up to eight DHCPv6 servers can be specified on the IPv6 pool manager.

The IPv6 pool manager keeps track of the context in which it allocates subnets in order to correctly acquire subnets after a fail-over or reboot. To support fail-over redundancy of this information, an identical virtual chassis identifier string can be configured on a redundant pair of NEs hosting IPv6 pool managers.

39.1.5 Local breakout using NAT

Typically, Wi-Fi traffic on a WLAN GW is considered GTP-tunneled or local breakout. Traffic is treated as local breakout by default. Local breakout traffic is controlled and routed to the Internet/ VPN from the WLAN GW directly.

For certain traffic from a Wi-Fi subscriber, it may be preferable to perform local breakout. This can be performed using the GTP Local Breakout action in the IP filter entries. When this action is applied to traffic marked for GTP-tunneling to the GGSN/PGW, the traffic is routed as local breakout. In this GTP local breakout case, the UE receives its IP address using GTP from the GGSN/PGW, but the matching traffic is routed as local breakout from the WLAN GW directly into the network. This traffic is passed through NAT. NAT is required to ensure the selected traffic is routed back to an IP address on the WLAN GW and not through the GGSN/PGW network. An L2-aware NAT host is created when the address is returned through GTP from the PGW/GGSN. The

returned address is the inside IP address. Internally, a GTP tunnel directs traffic to and from L2-aware NAT. Local and remote tunnel endpoint identifiers are create for the local breakout traffic and associated with the ESM host state. For packets that are not sent in the context of a GTP tunnel, this traffic is forwarded.

39.1.6 Hybrid access

In a hybrid access deployment model, the CPE sets up both an LTE and fixed access uplink to the Internet. Both uplinks share the same IP address and classifiers on the CPE, and the PDN gateway determines which link is to be used for traffic. The double uplinks can be used for both redundancy and bandwidth management. AMBR uplink and downlink bit rates are configurable on the mobile gateway/peer profile, along with radio access type. The PDN gateway is configurable for IPv4, IPv6, or Ipv4v6 address requests on the WLAN GW on a base or VPRN routing instance.

See 39.7 “Workflow to configure hybrid network access” (p. 1326).

39.1.7 VLAN access to anchor ISA

In cases where customers want to employ VLAN access from APs to the WLAN-GW over an L2 network, you can configure L2 APs as part of the WLAN GW configuration on IES and VPRN group interfaces. The NFM-P creates an internal epipe per access SAP or spoke SDP to the WLAN GW tunnel ISA, thereby preserving features requiring soft GRE functionality (for example, migrant user handling and distributed subscriber management). L2 frame processing on WLAN GW ISAs is functionally the same as with a WLAN GW tunnel.

An L2 AP configuration on a WLAN GW specifies an access port, encapsulation information, and an epipe SAP template policy. The epipe SAP template policy allows ingress and egress qos policy configuration and traffic filtering for access SAPs, allowing users to shape traffic based on the Wi-Fi radio technology of the AP.

39.1.8 L2 wholesale

In an L2 wholesale deployment model, subscribers are handled by a wholesale partner (MSO/telco) on behalf of the service provider (retailer), through an L2 network. In this model, the AP is configured with one or more SSIDs per service provider. The wholesaler does not host the subscriber for the service provider. L3 terminations (including address assignment, authentication, and subscriber management information) are performed by individual service providers on their own gateway.

The AP broadcasts an SSID per provider, and is configured with a unique .1q tag per SSID before bridging the L2 frames from the subscriber to the WLAN GW tunnel. VPLS L2 retail service instances are configured statically on the WLAN-GW. Static mappings of .1q tag(s) to L2 retail service instances are configured on the WLAN GW group interface. For each L2 service referenced under a VLAN tag range under the WLAN GW group interface, an internal SAP between the tunnel anchor ISA and IOM is created in that L2 service, on each of the WLAN GW IOMs in the WLAN GW group.

39.2 Workflow to configure lightRadio Wi-Fi

39.2.1 Stages

1

Configure a port policy to be bound to the ISA-WLAN GW group. See 50.65 “To configure a port policy” (p. 1609).

2

Configure a NAT RADIUS server policy to be bound to the ISA-WLAN GW group. See 57.4 “To configure a NAT RADIUS accounting policy” (p. 1795).

3

Configure an ISA-WLAN GW group. See 13.14 “To configure a WLAN GW group” (p. 428).

4

Configure RADIUS proxy functionality:

1. Configure RADIUS servers on a routing instance. See 27.8 “To configure a RADIUS server on a routing instance” (p. 846).
2. Configure RADIUS proxy servers on a routing instance. See 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847).
3. Configure RADIUS server policies. See 57.6 “To configure a RADIUS server policy” (p. 1797).
4. Configure RADIUS proxy caches. See 74.9 “To configure a local user database for subscriber host authentication” (p. 2025).

5

Configure a WLAN GW for an IES or VPRN group interface. See 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462)or 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631).

6

Configure a mobile gateway/peer profile. See 64.20 “To configure a mobile gateway/peer profile” (p. 1862).

7

Configure a WLAN GW for a routing instance or VPRN routing instance. See 27.2 “To configure a routing instance or a VRF instance” (p. 826)or 79.65 “To configure WLAN GW functionality on a VPRN site” (p. 2630).

8

Configure DNS addresses on a VPRN routing instance. See 79.12 “To configure DNS for a VPRN site” (p. 2546).

9

Configure serving network information on an NE. See 12.32 “To configure serving network information on an NE” (p. 367).

39.3 Workflow to configure Wi-Fi local breakout using L2–aware NAT

39.3.1 Stages

1

Configure an ACL IP filter policy. Configure the filter entries on the policy with the required match criteria and a filter action of GTP Local Breakout. See 51.5 “To configure an ACL IP filter policy” (p. 1671).

2

Bind the ACL IP filter policy configured in Stage 1 to the subscriber SLA profile. See 64.5 “To configure an SLA profile” (p. 1845).

3

Under the appropriate routing Instance, configure outside NAT pools with address ranges. See 30.9 “To configure NAT on a routing instance” (p. 1086).

4

Configure a NAT policy and bind it to the outside NAT pool configured in Stage 3 . See 30.5 “To configure a NAT policy” (p. 1082).

5

Bind the NAT policy configured in Stage 4 to the subscriber profile for the subscriber. See 64.4 “To configure a subscriber profile” (p. 1840).

39.4 Workflow to configure distributed subscriber management

39.4.1 Stages

1

Configure distributed subscriber management traffic policers. See 64.27 “To configure a distributed subscriber management traffic policer” (p. 1869).

2

Configure distributed subscriber management IPv4 and IPv6 filter policies. See 64.28 “To configure a distributed subscriber management IP filter policy” (p. 1869).

3

On an L2-aware NAT configuration, specify upstream and downstream IPv4 and IPv6 filter policies. See 30.9 “To configure NAT on a routing instance” (p. 1086).

4

On an IES soft GRE group interface, configure a VLAN tag range with the following (see 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462)):

- Authenticate on DHCP parameter
- distributed subscriber management
- distributed subscriber management traffic policer
- distributed subscriber management IP filter policy
- Idle timeout action, SLAAC host lifetime, and DHCPv6 host lifetime

5

On a VPRN soft GRE group interface, configure a VLAN tag range with the following (see 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631)):

- Authenticate on DHCP parameter
- distributed subscriber management
- distributed subscriber management traffic policer
- distributed subscriber management IP filter policy
- Idle timeout action, SLAAC host lifetime, and DHCPv6 host lifetime

6

To accommodate IPv6 hosts, configure the TCP Maximum Segment Size Adjustment parameter for the WLAN GW configuration on the VPRN site. See 79.65 “To configure WLAN GW functionality on a VPRN site” (p. 2630).

7

To accommodate IPv6 hosts, configure an IPv6 pool manager on IES or VPRN subscriber interfaces. See 78.16 “To configure a subscriber interface on an IES” (p. 2444) and 79.47 “To configure a subscriber interface on a VPRN” (p. 2610).

If required, configure WLAN GW redundancy on sites configured with IPv6 pool managers. See 12.34 “To configure WLAN GW redundancy on an NE” (p. 368).

8

Configure an ISA RADIUS policy with NAS IP Address, NAS Port, and Class accounting attributes, and with NAS IP Address and NAS Port authentication attributes. See 57.12 “To configure an ISA RADIUS policy” (p. 1803).

To accommodate IPv6 hosts, configure the ISA RADIUS poilcy with the IPv6 Address, Framed IPv6 Prefix, Framed Interface ID, and DHCPv6 Options accounting attributes, and the IPv6 Address and DHCPv6 Options authentication RADIUS attributes.

9

Configure an ISA-WLAN GW group with an ISA-AA group. See 13.14 “To configure a WLAN GW group” (p. 428).

10

Configure an HTTP redirect policy with the Prefix Length parameter. See 64.24 “To configure an HTTP redirect policy” (p. 1865).

11

Configure a RADIUS proxy server with RADIUS proxy users, default RADIUS policies, and a WLAN GW group. See 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847).

12

Configure an LI source WLAN distributed subscriber as a mirror source on a mirror service, if required. See 94.29 “To specify an LI WLAN distributed subscriber as an LI source” (p. 3224).

39.5 Workflow to configure VLAN to anchor ISA functionality

39.5.1 Stages

1

Configure an epipe SAP template policy. See 84.2 “To configure an epipe SAP template policy” (p. 2749).

2

Configure L2 access points on a WLAN GW on an IES or VPRN group interface. See 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462) or 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631). Associate the L2 access points with the SAP template policy configured in Stage 1 .

39.6 Workflow to configure L2 wholesale

39.6.1 Stages

1

Configure a VPLS SAP template policy. See 84.3 “To configure a VPLS SAP template policy” (p. 2750).

2

Configure the VPLS SAP template policy from Stage 1 on a VPLS service site. See 77.55 “To configure WLAN GW L2 wholesale forwarding on a VPLS site” (p. 2315).

3

Configure the VPLS service from Stage 2 on a WLAN GW on an IES or VPRN group interface. See 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462) or 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631).

39.7 Workflow to configure hybrid network access

39.7.1 Stages

1

Configure a mobile gateway/peer profile with Radio Access Type on the General tab, and AMBR downlink and uplink limits on the PGW and GGSN tabs. See 64.20 “To configure a mobile gateway/peer profile” (p. 1862).

2

Configure the PDN gateway address request type for the WLAN GW configuration on a base routing instance or VPRN routing instance on which hybrid access is implemented. Apply the mobile gateway/peer profile configured in Stage 1 to the WLAN GW. See 27.2 “To configure a routing instance or a VRF instance” (p. 826) or 79.65 “To configure WLAN GW functionality on a VPRN site” (p. 2630).

40 MC peer groups

40.1 Overview

40.1.1 Purpose

This chapter describes how to configure and manage MC peer groups.

40.1.2 Contents

MC peer groups overview

40.2 MC peer groups overview

40.2.1 Overview

An MC peer group is an NFM-P object that defines the relationship between two peer NEs to provide system redundancy in an Ethernet network. An MC peer group configuration includes a list of protocols and objects with state information that is to be synchronized between the peers.

On supporting NEs, you can create up to 20 MC peer groups using Ethernet ports in access mode. An MC peer group can contain one or more of the following child group objects:

- MC IPsec group—see Chapter 41, “MC IPsec”
- MC endpoint group—see Chapter 42, “MC endpoint groups”
- MC LAG group—see Chapter 43, “MC LAG groups”
- MC synchronization group—see Chapter 44, “MC synchronization groups”
- MC ring group—see Chapter 45, “MC ring groups”

Note: You can create MC endpoints only on a 7450 ESS, 7750 SR, or 7950 XRS.
You can create MC IPsec tunnel groups only on a 7450 ESS-4, 7450 ESS-6, 7450 ESS-6v, 7450 ESS-7, 7450 ESS-12 in mixed mode and in chassis mode D, or on a 7750 SR-7, 7750 SR-12, or 7750 SR-12E in chassis mode D.

When you create a child group object, the NFM-P automatically creates the child group members using the peer objects in the MC peer group.

The NFM-P automatically discovers MC peers that are configured on managed NEs and creates an MC peer group if the source address of each peer matches the peer address of the other. If an MC peer address on an NE is changed after discovery, for example, using a CLI, the NFM-Pdeletes the MC peer group but does not delete the peer configuration on either NE. When the mismatch is corrected, the NFM-P recreates the MC peer group.

When the MC peer addresses match but the NFM-P detects another MC peer group configuration mismatch, the NFM-P raises an alarm and displays a check mark beside the Asymmetrical Configuration Detected indicator on the General tab of the MC Peer Group properties form. The alarm information includes the type of configuration mismatch. When the mismatch is corrected, the alarm and check mark clear.

You can configure NE redundancy using the Manage Node Redundancy form or an NE properties form. When the NFM-P manages both NEs in an MC peer group, you can configure the MC peer and peer group parameters using the Manage Node Redundancy form. When the NFM-P manages only one NE in the MC peer group, you can configure the only the managed peer parameters, and only by using the NE properties form.

You can configure MC IPsec as a failover mechanism for IPsec tunnels between two 7750 SRs. MC IPsec provides protection for NE or MS-ISA failure. When the active peer fails, the IPsec tunnels failover to the standby peer without re-establishing the session. See Chapter 41, “MC IPsec” for more information.

40.2.2 OMCR on MC peer groups

You can configure the warm standby option on MC peer groups to provide oversubscribed multi-chassis redundancy, or OMCR, throughout the peer group.

In the OMCR model, a central standby NE backs up multiple subscriber hosts on multiple NEs; the subscribers are synchronized among the chassis CPMs. By avoiding subscriber host instantiation in the forwarding plane during normal operation, resources are used only as required.

MC peer group management procedures

40.3 MC peer group management workflow and procedures

40.3.1 Stages

1

Create an MC peer group; see 40.4 “To configure an MC peer group” (p. 1330) .

2

As required, configure an MC peer in an MC peer group to correct a configuration mismatch; see 40.5 “To configure an MC peer” (p. 1332) .

3

As required, perform an on-demand protocol synchronization between MC peer group members; see 40.6 “To perform an on-demand protocol synchronization between MC peer group members” (p. 1334) .

40.4 To configure an MC peer group

40.4.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu.

3

Click Create or select a peer group and click Properties. The MC Peer Group (Create|Edit) form opens.

Use the following steps:

1. Choose a site for the peer group in the First Element panel.

   At least one MC peer must be an NFM-P-managed NE.

   The NFM-P raises an alarm when an MC peer group contains an unmanaged NE.
2. To specify an interface other than the NE system interface for MC peer communication, configure the Source Address parameter by typing a valid interface IP address, or by clicking Select and choosing an interface; only numbered, non-multicast IPv4 interfaces are listed.
3. Choose a site for the peer group in the Second Element panel.

4. To specify an interface other than the NE system interface for MC peer communication, configure the Source Address parameter by typing a valid interface IP address, or by clicking Select and choosing an interface; only numbered, non-multicast IPv4 interfaces are listed.
5. If you plan to configure an MC IPsec Group, configure the Authentication Key parameters and select the Enable MC IPsec parameter.
6. Click Apply.

4

Configure the required parameters in the First Element and Second Element panels.

You must configure the Warm Standby parameter only in the First Element or Second Element panel, and not in both panels.

If you do not configure the Peer Name parameter, the parameter is automatically configured using the peer IP address.

5

Click on the Associated Groups tab.

6

Configure an MC LAG group.

Use the following steps:

1. Right-click on the MC LAG Group object and choose Create MC LAG Group, or expand the object, right-click on a group and choose Properties. The MC LAG Group (Create | Edit) form opens.
2. Perform Step 6 to Step 10 of 43.5 “To create an MC LAG group” (p. 1365) .

7

Configure an MC synchronization group.

Use the following steps:

1. Right-click on the MC Sync Group object and choose Create MC Sync Group, or expand the object, right-click on a group and choose Properties. The MC Sync Group (Create | Edit) form opens.
2. Perform Step 6 to Step 11 of 44.4 “To create an MC synchronization group” (p. 1376) .

8

Configure an MC ring group.

Use the following steps:

1. Right-click on the MC Ring Group object and choose Create MC Ring Group, or expand the object, right-click on a group and choose Properties. The MC Ring Group (Create | Edit) form opens.

2. Configure the Name parameter.
3. Click Select to configure the Synchronization Tag parameter by choosing an MC synchronization group.
4. Click OK to save your changes and close the form.

9

Configure an MC endpoint group.

Use the following steps:

1. Right-click on the MC Endpoint Group object and choose Create MC Endpoint Group, or expand the object, right-click on a group and choose Properties. The MC Endpoint Group (Create | Edit) form opens.
2. Perform Step 7 to Step 10 of 42.4 “To configure an MC endpoint group” (p. 1357) .

10

Configure an MC IPsec group.

Use the following steps:

1. Right-click on the MC IPsec Group object and choose Create MC IPsec Group, or expand the object, right-click on a group and choose Properties. The MC IPsec Group (Create | Edit) form opens.
2. Perform Step 7 to Step 10 of 42.4 “To configure an MC endpoint group” (p. 1357) .

11

Click on the Peer Synchronization tab.

Use the following steps:

1. Configure the Sync Administrative State parameter and select peer group synchronization options, as required.

   NE synchronization status is displayed in the bottom half of the form.
2. Click Re-Synchronize to manually update the NE synchronization status.

12

Save your changes and close the forms.

END OF STEPS

40.5 To configure an MC peer

40.5.1 Purpose

Perform this procedure to correct a configuration mismatch between the peers in an MC peer group by modifying an MC peer configuration.

Note: The MC peer parameters are configurable only when there is a configuration mismatch, which is indicated when the Neighbor Match check box on the MC Peer (Edit) form is disabled.

40.5.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search.

3

Click Create or select an MC peer group and click Properties. The MC Peer Group (Create|Edit) form opens.

4

Click on the Members tab.

5

Select a peer and click Properties. The MC Peer (Edit) form opens.

6

Configure the required parameters.

7

Click on the MC LAG tab and configure the parameters, as required.

8

Click on the Synchronization Protocol tab to configure MC synchronization group parameters.

Use the following steps:

1. Click Create or select an entry and click Properties. The MC Sync (Create|Edit) form opens.
2. Configure the required parameters on the General tab.
3. Click on the States tab to configure the Administrative State parameter for the synchronization protocol.
4. Click on the Track SRRP tab to configure a track SRRP instance for the synchronization protocol.
5. Click Create or select an existing track SRRP instance and click Properties. The Track SRRP (Create|Edit) form opens.
6. Configure the required parameters.

   The MCS State Disabled check box must be un-checked in order to configure the L2TP

Tunnel ID Range Start and End parameters.

7. Save your changes and close the form.

9

Save your changes and close the remaining forms.

END OF STEPS

40.6 To perform an on-demand protocol synchronization between MC peer group members

40.6.1 Purpose

Perform this procedure to distribute the most recent protocol synchronization configuration to each member of an MC peer group. This is required when the configurations do not match, for example, after the configuration is changed locally on only one MC peer using a CLI. When the peer configurations do not match, the NFM-P raises an alarm and displays a check mark beside the Asymmetrical Configuration Detected indicator on the General tab of the MC Peer Group properties form.

40.6.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC peer groups is displayed.

3

Select an entry and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Peer Synchronization tab.

5

Click Resync. The NFM-P resynchronizes the information.

6

Close the MC Peer Group (Edit) form.

7

Close the Manage Node Redundancy form.

END OF STEPS

40.7 To view the unmanaged MC peer of an NE

40.7.1 Purpose

Perform this procedure to view an MC peer when the NFM-P manages only one NE in an MC peer group.

40.7.2 Steps

1

In the navigation tree equipment view, right-click on the NE where you need to view an MC peer and select Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab.

3

Click on the MC Peer tab.

4

Click Search. A list of MC peers is displayed.

5

Select a peer and click Properties. The MC Peer (Edit) form opens.

6

View the information, as required.

7

Close the MC Peer (Edit) form.

8

Close the Network Element (Edit) form.

END OF STEPS

40.8 To delete an MC peer group

CAUTION

Service Disruption

Service disruption

Deleting an MC peer group removes all of the MC configurations that are associated with the MC peer group, such as the following:

- MC LAG groups
- MC synchronization groups
- MC ring groups

Ensure that you specify the correct MC peer group for deletion in this procedure.

40.8.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the object drop-down menu.

3

Click Search. A list of MC peer groups is displayed.

4

Select one or more MC peer groups in the list and click Delete.

5

Click View Dependencies. A Warning form opens.

6

View the dependency information.

7

Select the I understand the implications of this action check box.

8

Click Yes. The NFM-P deletes the MC peer group and the corresponding configuration on each member NE.

9

Close the Manage Node Redundancy form.

END OF STEPS

41 MC IPsec

41.1 Overview

41.1.1 Purpose

This chapter describes MC IPsec and how to configure and manage it.

41.1.2 Contents

MC IPsec overview

41.2 MC IPsec overview

41.2.1 Overview

MC IPsec provides a stateful failover mechanism for IPsec tunnels between NEs in an active and standby configuration. Stateful failover allows IPsec traffic to continue to be forwarded without interruption if a failure occurs. MC IPsec provides protection for NE or MS-ISA failure. If an active NE fails, the IPsec tunnels failover to the standby peer without needing to re-establish the session. The failover mechanism can occur at the tunnel group level. A tunnel group can failover to the standby NE, independently of other tunnel groups on the active NE.

A mastership protocol is used to elect the active NE peer. The NFM-P synchronizes IPsec configuration states between the active and standby peers so that existing tunnels do not need to be re-established when a switchover occurs. The IPsec traffic is sent to the active NE peer using an IPsec route policy that exports IPsec routes to the routing protocol. The route metric is then changed according to the changes in active and standby roles.

You can view the role of a peer in an MC IPsec tunnel group by verifying the Master State indicator on the MC IPsec Group configuration form. The states are:

- Not Eligible—no election and not eligible for mastership
- Eligible—no election but eligible for mastership
- Discovery—no election during initial peer discovery interval
- Master—communicating with MC peer, elected active
- Standby—communicating with MC peer, elected standby

You can view the details of the last switchover on the MC IPsec Tunnel Group (Edit) form for a peer that is part of the group. Review the values of the following parameters provided on the States panel:

- Mastership Election State—indicates the result of the mastership election between the tunnel group entity and its peer
- Old Master State parameter—indicates the older mastership election state of the tunnel group after the last state change
- State Change Reason parameter—indicates the reason for the last mastership election state change of the tunnel group

MC IPsec only supports IKEv2 static LAN-to-LAN tunnel. MC IPsec is supported only on the 7450 ESS-4, 7450 ESS-6, 7450 ESS-6v, 7450 ESS-7, and 7450 ESS-12 in mixed mode and in chassis mode D, and on the 7750 SR-7, 7750 SR-12, 7750 SR-c12 and 7750 SR-12E in chassis mode D.

Note: The MC IPsec peers must be of the same device type. For example, the NEs in an MC peer group must each be 7750 SR or 7450 ESS.

You can use the NFM-P to create MC IPsec non-forwarding events on VRRP policies to track the MEP state of MC IPsec-based static routes. See 55.2 “To configure a VRRP priority-control policy” (p. 1774) for more information.

IPsec VPN does not support redundancy.

MC IPsec management procedures

41.3 Workflow to configure and manage MC IPsec

41.3.1 Stages

1

Configure an ISA tunnel MDA and ISA tunnel group on each MC IPsec peer NE; see 13.10 “To configure an ISA-tunnel group” (p. 424) .

You must do the following:

- Enable the Multiple Active ISA Support parameter.
- Enable the IPsec Responder Only parameter to configure the NE to act as an IKE responder when an MC IPsec switchover occurs, and not initiate a new SA policy or rekey an existing SA.

Ensure that the ISA tunnel groups have the same configuration on each MC IPsec peer NE; alternatively, you can create ISA tunnel groups when you configure the MC IPsec group in Stage 3

2

Enable MC IPsec on the MC peer group. See 41.4 “To configure MC IPsec on an MC peer group” (p. 1343) .

3

Create an MC IPsec tunnel group between the MC peer NEs. See 41.5 “To create an MC IPsec group” (p. 1344) .

4

Configure MC IPsec parameters.

The parameters must be configured on the following interfaces:

- VPRN tunnel interfaces; see 41.7 “To configure MC IPsec on a VPRN tunnel interface” (p. 1347) .
- IES or VPRN L3 access interfaces; see 41.8 “To configure MC IPsec on an IES or VPRN L3 access interface” (p. 1348) .

5

Note: The IPsec tunnels must have the same tunnel name, properties, and associated IKE, transform, and security policies.

Tunnels that have a configuration mismatch between peers are lost when a switchover occurs.

Configure two identical static IPsec tunnels using the redundant tunnel group on the two MC peer NEs; see 34.21 “To configure an IPsec tunnel on a VPRN tunnel interface” (p. 1253) .

6

Configure IPsec protocol and state parameters in the From Criteria of a routing policy statement entry; see 54.5 “To configure a routing policy statement” (p. 1745) .

You must do the following:

- Set the Protocol parameter to IPsec.
- Set the State parameter to one of the following:
  - IPsec Master with Peer—The corresponding tunnel group is the master with a reachable peer.
  - IPsec Master No Peer—The corresponding tunnel group is the master with an unreachable peer.
  - IPsec Non Master—The corresponding tunnel group is not the master.

7

If required, perform a manual switchover from the MC IPsec tunnel group and from all of the associated IPsec tunnels that are configured above the MC IPsec tunnel group. See 41.9 “To perform an MC IPsec switchover” (p. 1349) .

41.4 To configure MC IPsec on an MC peer group

41.4.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu.

3

Click Search. A list of MC peer groups is displayed.

4

Choose an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

5

Select the Enable MC IPsec parameter.

6

Click Apply.

7

Configure the parameters for the peer site.

8

Click on the Peer Synchronization tab.

9

Select the MC IPsec parameter.

10

Click Re-Synchronize.

11

Close the MC Peer Group (Edit) form.

12

Close the Manage Node Redundancy form.

END OF STEPS

41.5 To create an MC IPsec group

41.5.1 Purpose

Perform this procedure to create an MC IPsec group of ISA tunnels. You can create up to 16 MC IPsec groups on each MC peer site.

Note: You must first assign an ISA tunnel group to each MC peer site. You must enable the Multiple Active ISA Support parameter on each ISA tunnel group for MC IPsec. You can create ISA tunnel groups when you configure the MC IPsec peer groups.

Alternatively, you can create ISA tunnel groups on the peer NEs before you perform this procedure. See 13.10 “To configure an ISA-tunnel group” (p. 424) .

You must first ensure that MC IPsec is enabled on the MC peer group and ensure that the peers are configured for MC IPsec. See 41.4 “To configure MC IPsec on an MC peer group” (p. 1343) for more information.

An MC IPsec group can exist only when there is a member pair of MC IPsec tunnel groups between MC peer NEs. If one member of a tunnel group is deleted or changed from one peer NE, the NFM-P global MC IPsec group is also deleted or changed.

41.5.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC Peer Groups is displayed.

3

Choose an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on the MC IPsec Group object in the components tree and choose Create MC IPsec Group. The MC IPsec Group (Create) form opens.

6

Configure the Synchronization Tag parameter.

7

Configure an MC IPsec tunnel group.

You must configure the tunnel group in each of the following panels:

- MC IPsec Tunnel Group on First Site
- MC IPsec Tunnel Group on Second Site

Perform the following:

1. Click Select. A list form opens.
2. To use an existing tunnel group, go to 8 .
3. Click Create. The ISA-Tunnel Group (Create) form opens.
4. Enable the Multiple Active ISA Support and IPsec Responder Only parameters.
5. Configure the remaining parameters.
6. Click OK. The MC IPsec Group (Create) form closes.
7. Click Search.
8. Select the tunnel group and click OK.
9. Configure the remaining parameters.

8

Click OK. The MC IPsec Group (Create) form closes.

9

Close the MC Peer Group (Edit) form.

10

Close the Manage Node Redundancy form.

END OF STEPS

41.6 To force the synchronization tag deployment to MC IPsec peers

41.6.1 Purpose

Each member of an MC IPsec group must have the same Synchronization Tag value in order for the group to be operationally up. If the value on a member NE is modified, for example, using a CLI, a mismatch exists and the NFM-P raises an alarm against each member NE.

Perform this procedure when the member NEs of an MC IPsec group have different Synchronization Tag values.

41.6.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC IPsec Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC IPsec groups is displayed.

3

Choose an MC IPsec group and click Properties. The MC IPsec Group (Edit) form opens.

4

Click Force Sync.

Note: If you modify the global Synchronization Tag parameter, you cannot force the synchronization using this procedure. You must click Apply or OK to deploy the value to the NEs.

5

Click Yes. The NFM-P deploys the Synchronization Tag to the member NEs of the MC IPsec group.

6

Close the MC IPsec Group (Edit) form.

7

Close the Manage Node Redundancy form.

END OF STEPS

41.7 To configure MC IPsec on a VPRN tunnel interface

41.7.1 Purpose

Perform this procedure to configure MC IPsec to move traffic to the master NE in the event that the standby NE receives traffic.

Note: The port configured on the tunnel must be created using a tunnel SAP. For example, tunnel-x.private/public:y. See 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249) for more information.

41.7.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search. A list of services is displayed.

3

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

4

Click on the Tunnel Interfaces tab.

5

Choose a tunnel interface and click Properties. The Tunnel Interface (Edit) form opens.

6

Configure the Redundant Next-Hop IP Address for Static ISA Tunnels and Redundant Next-Hop IP Address for Dynamic ISA Tunnels parameters.

Note: The next hop address that you specify is used by the standby NE to shunt traffic to the master NE in the event that the standby receives traffic. The next hop address is resolved in the routing table of the corresponding service.

7

Click OK. The Tunnel Interface (Edit) form closes.

8

Close the VPRN Service (Edit) form.

9

Close the Manage Services form.

END OF STEPS

41.8 To configure MC IPsec on an IES or VPRN L3 access interface

41.8.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search, select an IES or VPRN service, and click Properties. The IES Service (Edit) or VPRN Service (Edit) form opens.

3

In the service navigation tree, expand the Sites icon, and then the icon of the Site or Routing Instance that hosts the L3 access interface.

4

Expand the L3 Access Interfaces icon and select the required L3 access interface.

5

Configure the Redundant Next-Hop IP Address for Static ISA Tunnels and Redundant Next-Hop IP Address for Dynamic ISA Tunnels parameters.

Note: The parameters are configurable only when the port on the interface is created with a tunnel SAP. For example, tunnel-x.private/public:y.
The next hop address that you specify is used by the standby NE to shunt traffic to the master NE if the standby receives traffic. The next hop address is resolved in the routing table of the corresponding service.

6

Click OK. The IES Service (Edit) or VPRN Service (Edit) form closes.

7

Close the Manage Services form.

END OF STEPS

41.9 To perform an MC IPsec switchover

41.9.1 Purpose

Perform this procedure to perform a manual switchover from an MC IPsec tunnel group and from all of the associated IPsec tunnels that are configured above the MC IPsec tunnel group.

41.9.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC IPSec Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC IPSec groups is displayed.

3

Select an MC IPsec group and click Properties. The MC IPSec Group (Edit) form opens.

4

Click on the Components tab.

5

Right-click on a site object below the MC IPsec Group object and choose Properties. The MC IPsec Tunnel Group (Edit) form opens.

6

Perform one of the following:

a. If the Last Tunnel Group Switch Status indicator in the Protection Status panel reads Nominal, click Protection Switch.

b. When the Last Tunnel Group Switch Status in the Protection Status panel for the group is Not Ready, you cannot perform the switchover.

   This is due to one of the following:

   - there is no elected standby peer
   - the synchronization of pending IPsec states must occur
   - IKE states are pending download to ISAs

7

Close the MC IPSec Tunnel Group (Edit) form.

8

Close the MC IPSec Group (Edit) form.

9

Close the Manage Node Redundancy form.

END OF STEPS

41.10 To view the unmanaged MC IPsec peer of an NE

41.10.1 Purpose

Perform this procedure to view an MC IPsec peer when only one NE in the MC IPsec peer group is managed by the NFM-P.

41.10.2 Steps

1

In the navigation tree equipment view, right-click on the NE where you want to view an MC IPsec peer and choose Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab, and then on the MC Peer sub-tab. The MC peers are listed.

3

Select an MC peer and click Properties. The MC Peer (Edit) form opens.

4

Click on the MC IPsec tab. The MC IPsec peers are listed.

5

Select an MC IPsec peer and click Properties. The MC IPsec (Edit) form opens.

6

View the information, as required.

7

Close the MC IPsec (Edit) form.

8

Close the MC Peer (Edit) form.

9

Close the Network Element (Edit) form.

END OF STEPS

41.11 To configure an MC IPsec peer

41.11.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC peer groups is displayed.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Before you can configure an MC IPsec peer, you must administratively disable the tunnel groups associated with the peer site.

As required, administratively disable the associated tunnel groups.

1. Click Properties beside the First Peer Site ID or Second Peer Site ID parameter in the MC IPsec panel. The MC IPsec (Edit) form opens.
2. Click on the Tunnel Groups tab.
3. Select a tunnel group and click Properties. The MC IPsec Tunnel Group (Edit) form opens.
4. Set the Administrative State parameter to Down.
5. Click OK. The MC IPsec Tunnel Group (Edit) form closes.

5

Click on the General tab and configure the parameters.

6

Click OK. The MC IPsec (Edit) form closes.

7

Administratively enable the disabled tunnel groups.

Perform the following:

1. Click Properties beside the First Peer Site ID or Second Peer Site ID parameter in the MC

IPsec panel. The MC IPsec (Edit) form opens.

2. Click on the Tunnel Groups tab.
3. Select a tunnel group and click Properties. The MC IPsec Tunnel Group (Edit) form opens.
4. Set the Administrative State parameter to Up.
5. Click OK. The MC IPsec Tunnel Group (Edit) form closes.

8

Close the MC Peer Group (Edit) form.

9

Close the Manage Node Redundancy form.

END OF STEPS

41.12 To configure an MC IPSec domain

41.12.1 Steps

1

In the navigation tree equipment view, right-click on the *NE* and choose Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab, and then on the MC IPSec Domain sub-tab.

3

Click Create or select an entry and click Properties. The MC IPSec Domain (Create|Edit) form opens.

4

Click on the General tab and configure the parameters.

5

Click on the Statistics tab to view and collect statistics, as required.

6

Save your changes and close the forms.

END OF STEPS

41.13 To configure an MC Peer IPSec domain

41.13.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Members tab.

5

Select a peer and click Properties. The MC Peer (Edit) form opens.

6

Click on the MC IPSec tab.

7

Select an MC IPsec peer and click Properties. The MC IPsec (Edit) form opens.

8

Click on the IPSec Domains tab.

9

Click Create or select an entry and click Properties. The MC Peer IPSec Domain (Create|Edit) form opens.

10

Configure the required parameters.

11

Associate a MC IPSec domain in the Peer panel; see 41.12 “To configure an MC IPSec domain” (p. 1352) .

12

Save your changes and close the forms.

END OF STEPS

42 MC endpoint groups

42.1 Overview

42.1.1 Purpose

This chapter describes MC endpoint groups and how to manage them.

42.1.2 Contents

MC endpoint groups overview

42.2 MC endpoint groups and MC peer groups

42.2.1 MC endpoint groups and MC endpoint peer groups

An MC endpoint group consists of two MC endpoint peers that are configured as peers in a pair of sites. Multiple VPLS endpoints can use the MC endpoint peer. The endpoints must be on a supporting NE, and must support PW redundancy. Each endpoint can be associated with different destinations, with a maximum of two spoke SDPs for each endpoint. The endpoints communicate with each other to get the associated status of the spoke SDPs, which ensures that only one spoke SDP is active at any time; the other spoke SDPs have a standby status. The grouping of multiple spoke SDPs that are associated with the two MC endpoint peers eliminates traffic loops in a VPLS or B-VPLS.

An MC endpoint group includes:

- two MC endpoint peers that can be used by a spoke SDP
- an MC protocol that is used for:
  - determination of which MC endpoint peers are active or standby
  - synchronization of the PW information between the peers
  - fault detection using centralized BFD; the MC endpoint protocol includes a keep-alive mechanism because BFD cannot detect the state of an MC endpoint peer
- T-LDP signaling, which is used to communicate whether a PW is active or standby to other gateway pairs. The other gateway pairs may not have an MC endpoint.

An MC peer group that is managed by the NFM-P contains two MC peers. Each MC peer is configured as the peer of the other MC peer. An MC peer group must be created before you create an MC endpoint group. When only one peer in an MC peer group is managed by the NFM-P, you can configure node redundancy parameters only for the managed peer and only from the NE properties form. See Chapter 40, “MC peer groups” for information about MC peer groups.

42.2.2 BFD

The MC endpoint protocol uses keep-alive mechanisms. The MC endpoint protocol also supports BFD to eliminate traffic loops. See 42.4 “To configure an MC endpoint group” (p. 1357) for information about how to configure an MC endpoint group for BFD.

MC endpoint group management procedures

42.3 Workflow to manage MC endpoint groups

42.3.1 Stages

1

Create an MC peer group; see 40.4 “To configure an MC peer group” (p. 1330) .

2

Create an MC endpoint group; see 42.4 “To configure an MC endpoint group” (p. 1357) .

3

Create an endpoint on a VPLS site; see 77.43 “To create an endpoint for redundancy (dual homing) on a VPLS site” (p. 2303) .

4

Perform one of the following; see 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

a. Create a redundant spoke SDP binding under an endpoint.

b. Create a spoke SDP binding for the site.

42.4 To configure an MC endpoint group

42.4.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu.

3

Click Search. A list of MC peer groups is displayed.

4

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

5

Click on the Associated Groups tab.

6

Right-click on MC Endpoint Group and choose Create MC Endpoint Group. The MC Endpoint Group (Create) form opens.

7

Configure the Description parameter.

8

Configure the parameters in the MC Endpoint on First Site panel.

9

Configure the parameters in the MC Endpoint on Second Site panel.

10

Click OK to save your changes and close the form.

11

Close the MC Peer Group (Edit) form.

12

Close the Manage Node Redundancy form.

END OF STEPS

42.5 To view an MC endpoint peer on one NE

42.5.1 Steps

1

In the navigation tree equipment view, right-click on the NE where you need to view an MC endpoint peer and select Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab.

3

Click on the MC Peer tab.

4

Click Search. A list of MC peers is displayed.

5

Select an MC peer in the list and click Properties. The MC Peer (Edit) form opens.

6

Click on the MC Endpoint tab.

7

Click Search. A list of MC endpoint peers is displayed.

8

Select a peer and click Properties. The MC Endpoint (Edit) form opens.

9

View the information, as required.

10

Close the MC Endpoint (Edit) form.

11

Close the MC Peer (Edit) form.

12

Close the Network Element (Edit) form.

END OF STEPS

42.6 To delete an MC endpoint group

CAUTION

Service Disruption

When you perform this procedure, the MC endpoint group is deleted from the NFM-P and member NE configurations.

In addition, deleting an MC endpoint group removes all of the MC configurations that are associated with the MC peer group.

Note: If one or more MC endpoints use the MC endpoint group, you cannot delete the MC endpoint group until you remove the MC endpoints from the group.

42.6.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Endpoint Group (Multi-Chassis) from the drop-down menu.

3

Click Search. A list of MC endpoint groups is displayed.

4

Select one or more MC endpoint groups in the list and click Delete. A dialog box appears.

5

Click View Dependencies. A Warning form opens.

6

View the dependency information.

7

Select the I understand the implications of this action check box.

8

Click Yes. The NFM-P deletes the MC endpoint group and the corresponding configuration on each member NE.

9

Close the Manage Node Redundancy form.

END OF STEPS

43 MC LAG groups

43.1 Overview

43.1.1 Purpose

This chapter describes MC LAG groups and how to manage them.

43.1.2 Contents

MC LAG groups overview

43.2 MC LAG groups overview

43.2.1 Overview

A link aggregation group, or LAG, is a group of physical ports that form one logical link between two NEs to increase bandwidth, allow load balancing, and provide seamless redundancy. LAG support over multiple devices provides NE-level redundancy in addition to link-layer redundancy using a switchover function. An MC LAG configuration provides redundant L2 access connectivity that extends beyond link-layer protection by allowing two devices to share a common LAG endpoint.

An MC LAG configuration includes one active member NE and one standby member NE. The active and standby NEs synchronize the link state information to facilitate link-layer messaging between an access node and each NE. The active and standby NE coupling provides a synchronized forwarding plane to and from the access node. LACP is used to manage the active and standby states of the available LAG links; only the links of one member NE are active at one time.

The following NEs support the creation of MC LAG groups using Ethernet ports in Access mode.

- 7210 SAS
- 7250 IXR
- 7450 ESS
- 7705 SAR
- 7750 SR
- 7950 XRS

Support for MC LAGs on the 7210 SAS and 7705 SAR varies depending on the chassis type and release; see the NE documentation for support information.

The 7705 SAR-8, 7705 SAR-18, and 7705 SAR-M support the creation of MC LAG groups to provide NE-level redundancy. The operator has the option of creating an alternative path for deploying a service on the redundant NEs. The 7705 SAR NEs support inter-chassis backup, or ICB, on MC LAG to provide redundant Epipe service paths.

Some OmniSwitch NEs also support NE redundancy using MC LAGs; see 43.2.3 “MC AOS groups” (p. 1363) in this section.

When you use the NFM-P to change the MC LAG configuration on an NE, the NFM-P automatically updates the MC LAG configuration on the other NE. When you change the MC LAG configuration on an NE using, for example, a CLI, the NFM-P detects the configuration mismatch between the NEs and raises an alarm. The alarm information includes the type of configuration mismatch.

You can create an MC LAG group only from within an existing MC peer group. When only one peer in an MC peer group is managed by the NFM-P, you can configure NE redundancy parameters only for the managed peer, and only from the NE properties form. See Chapter 40, “MC peer groups” for information about MC peer groups.

43.2.2 MC synchronization

When subscriber management is enabled on an NE in an MC LAG configuration, the NE maintains dynamic state information for each subscriber host. The active and standby NEs synchronize the information to ensure uninterrupted service delivery in the case of an MC LAG switchover. See Chapter 44, “MC synchronization groups” for information about MC synchronization.

43.2.3 MC AOS groups

CAUTION

Service Disruption

You must not run an Ethernet OAM or SAA test on an OmniSwitch that has an MC LAG configuration.

MC AOS groups enable MC LAG creation on OS 9700E and OS 9800E NEs at Release 6.4.5 or later that are equipped with XNI-U12E cards, and on Release 7.3.1 and later OS 6900 and OS 10K NEs.

MC LAG group management workflows and procedures

43.3 Workflow to manage MC LAG groups

43.3.1 Stages

1

Create one or more LAGs using Ethernet access ports, as required; see Chapter 13, “Logical group object configuration” .

2

Create an MC peer group; see Chapter 40, “MC peer groups” .

3

Create an MC LAG group; see 43.5 “To create an MC LAG group” (p. 1365) .

4

As required, configure MC LAG group members; see 43.6 “To configure an MC LAG group member” (p. 1366) .

5

As required, configure MC LAG peer on one NE; see 43.7 “To configure an MC LAG peer on one NE” (p. 1367) .

43.4 Workflow to manage MC AOS groups

43.4.1 Stages

1

Create one or more LAGs using OmniSwitch ports, as required; see Chapter 13, “Logical group object configuration” .

2

Create an MC AOS Group; see 43.8 “To create an MC AOS group” (p. 1368) .

3

Create an MC AOS VFLink group; see 43.9 “To create an MC AOS VFLink Group” (p. 1369) .

4

Create an MC AOS LAG Group; see 43.10 “To create an MC AOS LAG Group” (p. 1370) .

5

As required, configure MC AOS Group members; see 43.11 “To configure an MC AOS group member” (p. 1371) .

43.5 To create an MC LAG group

43.5.1 Purpose

Perform this procedure to create an MC LAG group. Consider the following before you create an MC LAG group:

- You can create an MC LAG member only on an Ethernet MDA.
- MC LAG member ports must be in Access mode.
- The NFM-P assigns the same LACP key, system ID, and system priority to each MC LAG member.
- MC LAGs are not supported when MAC subnetting is enabled.

43.5.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) and click Search.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on the MC Lag object in the navigation tree and choose Create MC LAG Group. The MC LAG Group (Create) form opens.

6

Configure the required parameters.

The LACP Key, System ID, and System Priority parameters must be left at the default when the Standby Signalling parameter is set to Power Off.

7

Configure the Use LACP Key and MAC LSB (hex) parameters.

Note: The parameters are configurable only on the 7750 SR, 7750 SR-c12, 7750 SR-12E, 7450 ESS, and 7950 XRS. The NE must also be configured with an IOM 3 MDA or an XMA.
The parameters are displayed and configurable only when both MC LAG members are PBB-capable.

8

Choose the LAG for the first MC LAG member in the First Site panel.

Note: If you are configuring access dual homing with local switching over PBB tunnels, the L2 access interfaces must be on LAGs that participate in the MC LAG.

9

Choose the LAG for the second MC LAG member in the Second Site panel.

10

Save your changes and close the forms.

END OF STEPS

43.6 To configure an MC LAG group member

43.6.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC LAG Group (Multi-Chassis) and click Search.

3

Select an MC LAG and click Properties. The MC LAG Group (Edit) form opens.

4

Configure the required parameters.

Note: The LACP Key, System ID, and System Priority parameters must be left at the default when the Standby Signalling parameter is set to Power Off.

5

Click on the States tab.

6

Configure the members in the First Site and Second Site panels.

Use the following steps:

1. Click Properties. The MC Peer (Edit) form opens.
2. Click on the MC LAG tab and configure the required parameters.

   Note:

   If you change the default values of the Keep-Alive Interval (deciseconds) and Lost Connection Wait Interval parameters, the MC LAG group configuration of the member sites is not removed when the MC LAG group is deleted. You can view the MC LAG member information under the LAG icon in the network view of the navigation tree.

7

Save your changes and close the forms.

END OF STEPS

43.7 To configure an MC LAG peer on one NE

43.7.1 Steps

1

In the navigation tree equipment view, right-click on an NE and select Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab, and then on the MC Peer tab.

3

Choose an MC peer and click Properties. The MC Peer (Edit) form opens.

4

Click on the MC LAG tab and configure the required parameters.

5

Click on the Members tab.

6

Choose a member and click Properties. The MC LAG (Edit) form opens.

7

View the information, as required.

8

Save your changes and close the forms.

END OF STEPS

43.8 To create an MC AOS group

43.8.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC AOS Group (Aos Redundancy) and click Create. The MC AOS Group (Create) form opens.

3

Configure the required parameters.

4

Perform one of the following.

a. Configure the parameters in the First MultiChassis (Chassis Id 1) panel.

b. Select an OmniSwitch in the First MultiChassis (Chassis Id 1) panel and click OK.

5

Perform one of the following.

a. Configure the parameters in the Second MultiChassis (Chassis Id 2) panel.

b. Select an OmniSwitch in the Second MultiChassis (Chassis Id 2) panel and click OK.

6

Perform one of the following.

a. Configure the IPC VLAN Site ID parameter.

b. Enable the Default parameter.

7

Configure the remaining parameters.

Note: The Primary Minimum Lag Id and Primary Maximum Lag Id parameters are configurable only when the Configure Primary Lag Range parameter is enabled.

The Secondary Minimum Lag Id and Secondary Maximum Lag Id parameters are configurable only when the Configure Secondary Lag Range parameter is enabled. The Minimum MCLag Id and Maximum MCLag Id parameters are configurable only when the Configure MultiChassis Lag Range parameter is enabled.

8

Right-click on MC AOS LAG Group and choose Create MC AOS LAG Group. The MC AOS LAG Group (Create) form opens.

9

Configure the Lag ID parameter.

10

Save your changes and close the forms.

END OF STEPS

43.9 To create an MC AOS VFLink Group

43.9.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC AOS Group (Aos Redundancy) and click Search.

3

Choose an MC AOS group and click Properties. The MC AOS Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on MC AOS VFLink Group and choose Create MC AOS VFLink Group. The AOS VFLink Group (Create) form opens.

6

Select a Default Vlan for VFLink and click OK.

7

Click Apply.

8

Click Properties in the First MultiChassis VFLinkEnd (Chassis Id 1) panel. The AOS VF Link End form opens.

9

Click on the VF Link End Member Port tab and click Create. The Select Port - AOS VF Link End form opens.

Note: Up to eight member ports can be added for each OmniSwitch; the ports must be access ports with a speed of 10Gb or higher.

10

Select one or more ports and click OK. The Select Port - AOS VF Link End form closes.

11

Click Properties in the Second MultiChassis VFLinkEnd (Chassis Id 2) panel. The AOS VF Link End form opens.

12

Click on the VF Link End Member Port tab and click Create. The Select Port - AOS VF Link End form opens.

Note: Up to eight member ports can be added for each OmniSwitch; the ports must be access ports with a speed of 10Gb or higher.

13

Choose one or more ports and click OK. The Select Port - AOS VF Link End form closes.

Note: This member port cannot be used as a learned port security, VLAN access interface, or LAG member. Attempting to use the member port in such a way causes a deployment failure.

14

Save your changes and close the forms.

END OF STEPS

43.10 To create an MC AOS LAG Group

43.10.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC AOS Group (Aos Redundancy) and click Search.

3

Choose an MC AOS Group and click Properties. The MC AOS Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on MC AOS LAG Group and choose Create MC AOS LAG Group. The MC AOS LAG Group (Create) form opens.

6

Configure the Lag ID parameter.

Note: Configuring the Lag ID parameter with a range of Peer or MultiChassis results in a deployment failure.

7

Save your changes and close the forms.

END OF STEPS

43.11 To configure an MC AOS group member

43.11.1 Steps

1

In the navigation tree equipment view, right-click on an OmniSwitch NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the MultiChassis tab and configure the required parameters.

3

Click on the Lag Ranges tab and configure the required parameters.

Note: The Local Minimum Lag ID and Local Maximum Lag ID parameters are configurable only when the Configure Local Lag Range parameter is selected.
The Peer Minimum Lag ID and Peer Maximum Lag ID parameters are only configurable when the Configure Peer Lag Range parameter is selected.
The Minimum MCLag ID and Maximum MCLag ID parameters are only configurable when the Configure Local Lag Range parameter is selected.

4

Click on the Loopback Detection tab.

5

Configure the Admin Status and Transmit Interval (seconds) parameters.

6

Save your changes and close the form.

END OF STEPS

44 MC synchronization groups

44.1 Overview

44.1.1 Purpose

This chapter describes MC synchronization groups and how to manage them.

44.1.2 Contents

MC synchronization groups overview

44.2 MC synchronization groups overview

44.2.1 Overview

When subscriber management is enabled on an NE, the NE maintains dynamic subscriber host state information. The state information must be synchronized between the active and standby NEs in a redundant configuration to ensure that service delivery is uninterrupted if a switchover occurs.

The 7210 SAS, 7450 ESS, 7750 SR, and 7950 XRS support the creation of MC synchronization groups using Ethernet ports in Access mode.

Support for MC synchronization groups on the 7210 SAS varies depending on the chassis type and release; see the NE documentation for support information.

Note: You can create an MC synchronization group only in an existing MC peer group. See Chapter 40, “MC peer groups” for information about MC peer groups.

MC synchronization can be used to ensure that the following dynamic state information is synchronized:

- basic and enhanced subscriber management
- IGMP snooping in VPLS
- PIM snooping in VPLS
- IGMP on IES or VPRN group interfaces
- SRRP in VPRN

You create an MC synchronization group inside an MC peer group using a unique synchronization tag to define the pair of NEs and the two ports or LAGs on which the dynamic state information is synchronized. The synchronization is applied to all SAPs on the port or LAG that have the same synchronization tag. When only one peer in an MC peer group is managed by the NFM-P, you can configure node redundancy parameters only for the managed peer, and only from the NE properties form. See Chapter 40, “MC peer groups” for information about MC peer groups.

A synchronization tag can be applied to a specified VLAN range on a port or LAG. All of the SAPs in the VLAN range are assigned the synchronization tag. The SAPs that are not in the VLAN range are not synchronized.

Note: Only ports and LAGs that use dot1q or QinQ encapsulation support MC synchronization.

MC synchronization group VLAN ranges are configurable only after MC synchronization group creation.

If the NFM-P detects a port or VLAN range configuration mismatch in an MC synchronization group during NE discovery, the NFM-P raises an alarm.

44.2.2 MC synchronization and dual-homed L2/L3 CO

MC synchronization is typically used in a dual-homed L2 or L3 CO configuration. For example, an access node that aggregates several subscriber lines can be dual-homed to a redundant pair of

NEs. Dynamic subscriber-host state information on the NE must be synchronized with the redundant peer to ensure that service delivery in unaffected if a switchover occurs. See Chapter 78, “IES management” and Chapter 79, “VPRN service management” for more information about L2 and L3 CO dual homing.

MC synchronization for dual homing requires the configuration of protocol synchronization on the MC peer group that contains the redundant NEs. See 40.4 “To configure an MC peer group” (p. 1330) for information about configuring protocol synchronization on an MC peer group.

MC synchronization groups management procedures

44.3 MC synchronization groups management workflow and procedures

44.3.1 Stages

1

Create an MC peer group; see Chapter 40, “MC peer groups” .

2

Create an MC synchronization group, including MC peer synchronization ports or VLAN ranges; see 44.4 “To create an MC synchronization group” (p. 1376) .

3

Configure protocol synchronization for the MC synchronization group; see 44.5 “To configure protocol synchronization between MC peer group members” (p. 1378) .

44.4 To create an MC synchronization group

44.4.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC peer groups is displayed.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on the MC Sync Group object in the navigation tree and choose Create MC Sync Group. The MC Sync Group (Create) form opens.

6

Configure the parameters.

7

Click Select beside the Port/LAG Name parameter in the First Element panel to specify a port, PW port, or LAG on the first NE.

Note: If you set the Sync Tag Config Level parameter to VLAN Range Level, only dot1q and QinQ-encapsulated Ethernet ports, PW ports, or LAGs in Access mode are listed.

8

Click Select beside the Port/LAG Name parameter in the Second Element panel to specify a port, PW port, or LAG on the first NE.

Note: If you set the Sync Tag Config Level parameter to VLAN Range Level, only dot1q and QinQ-encapsulated Ethernet ports, PW ports, or LAGs in Access mode are listed.

9

If you set the Sync Tag Config Level parameter to SDP Level, specify the SDP ranges that the MC synchronization group is to monitor.

Note: If the Synchronization Tag parameter is configured at the SDP level, a range cannot be configured for the SDP. The Synchronization Tag parameter should not be the same across different SDPs, whether at the SDP level or the Range level.

Use the following steps.

1. Enter the SDP ID in the SDP panel.
2. Click Apply.
3. Click on the First Site SDP Range Entries|Second Site SDP Range Entries tab.
4. Click Create. The SDP Range (Create) form opens.
5. Configure the parameters.
6. Click OK. The SDP Range (Create) form closes.
7. Repeat Step 9 4 to 6 to add another SDP range for the site, if required.
8. Repeat Step 9 for the other site, if required.

10

If you set the Sync Tag Config Level parameter to VLAN Range Level, specify the VLAN ranges that the MC synchronization group is to monitor.

Note: The VLAN ranges that you configure are applied to the configuration of each site in the MC synchronization group.

Use the following steps.

1. Click Apply.
2. Click on the First Site VLAN Entries tab.
3. Click Create. The VLAN Range (Create) form opens.

4. Configure the parameters.

   The Minimum Inner Encap Value and Maximum Inner Encap Value parameters are configurable only when the encapsulation type of the associated ports or LAGs is QinQ.

5. Click OK. The VLAN Range (Create) form closes.

11

Close the MC Sync Group (Create) form.

12

Close the MC Peer Group (Edit) form.

13

Close the Manage Node Redundancy form.

END OF STEPS

44.5 To configure protocol synchronization between MC peer group members

44.5.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC peer groups is displayed.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Synchronize Protocols tab.

5

Configure the parameters.

6

Click OK.

7

Close the Manage Node Redundancy form.

END OF STEPS

44.6 To view the unmanaged MC synchronization peer of an NE

44.6.1 Purpose

Perform this procedure to view an MC synchronization peer when only one NE in the MC synchronization group is managed by the NFM-P.

Note: This procedure assumes that you have already configured an MC peer group with an associated MC synchronization group. If only one NE is managed by the NFM-P, you must create these objects from the CLI.

44.6.2 Steps

1

On the equipment view, right-click on the NE where you want to view an MC synchronization peer and choose Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab, and then on the MC Peer tab.

3

Click Search. A list of MC peers is displayed.

4

Select an MC peer and click Properties. The MC Peer (Edit) form opens.

5

Click on the Synchronization Protocol tab.

6

Click Search. A list of MC synchronization peers is displayed.

7

Select an MC synchronization peer and click Properties. The MC Sync (Edit) form opens.

8

View the information, as required.

9

Close the MC Sync (Edit) form.

10

Close the MC Peer (Edit) form.

11

Close the Network Element (Edit) form.

END OF STEPS

44.7 To delete an MC synchronization group

44.7.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Sync Group (Multi-Chassis) from the drop-down menu.

3

Click Search. A list of MC synchronization groups is displayed.

4

CAUTION

Service Disruption

Deleting an MC synchronization group that is associated with an MC ring group disables the protocol synchronization between the member sites in the MC ring group.

Select one or more MC synchronization groups in the list and click Delete.

5

Click Yes. The NFM-P deletes the MC synchronization group and the corresponding configuration on each member site.

6

Close the Manage Node Redundancy form.

END OF STEPS

45 MC ring groups

45.1 Overview

45.1.1 Purpose

This chapter describes MC ring groups and how to configure and manage them.

45.1.2 Contents

MC ring groups overview

45.2 MC ring groups overview

45.2.1 Overview

MC ring groups are an extension of dual-homing support that provides aggregation redundancy in networks that have multiple access nodes, called ring nodes in the context of MC ring groups, that are connected in a single ring.

The 7210 SAS, 7450 ESS, 7750 SR, and 7950 XRS support the creation of MC ring groups using Ethernet ports in Access mode.

Support for MC ring groups on the 7210 SAS varies depending on the chassis type and release; see the NE documentation for support information.

The following figure shows a simple subscriber aggregation network in which a single ring of access nodes, such as DSLAMs, is connected to a BSA in a VPLS. Each BSA is connected to an IES or VPRN L3 interface on a BSR using a spoke SDP. Each BSR L3 interface aggregates the subscriber traffic in one subnet.

Figure 45-1 Simple subscriber aggregation network

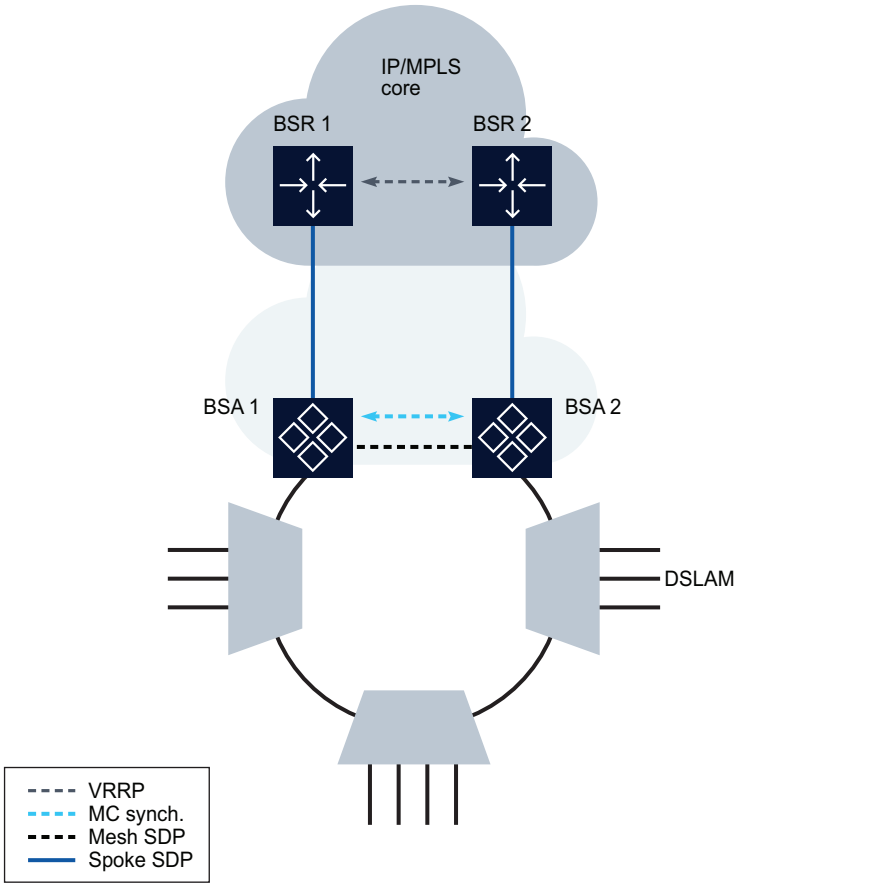

Note: BTV distribution is typically implemented in a separate VPLS that uses one SAP per access node.

An MC ring group does not support a ring that has more than one break.

45.2.2 Steady-state condition

In an MC ring group, the operation of a typical dual-homed ring is as shown in Figure 45-2, “Dual-homed ring, steady-state condition” (p. 1384) . A steady-state condition is defined by the following.

- The participating BSAs have matching MC ring configurations.
- The In-Band Ring Control Connection, or IB-RCC, is operationally up.
- The MC ring is operationally up.

Figure 45-2 Dual-homed ring, steady-state condition

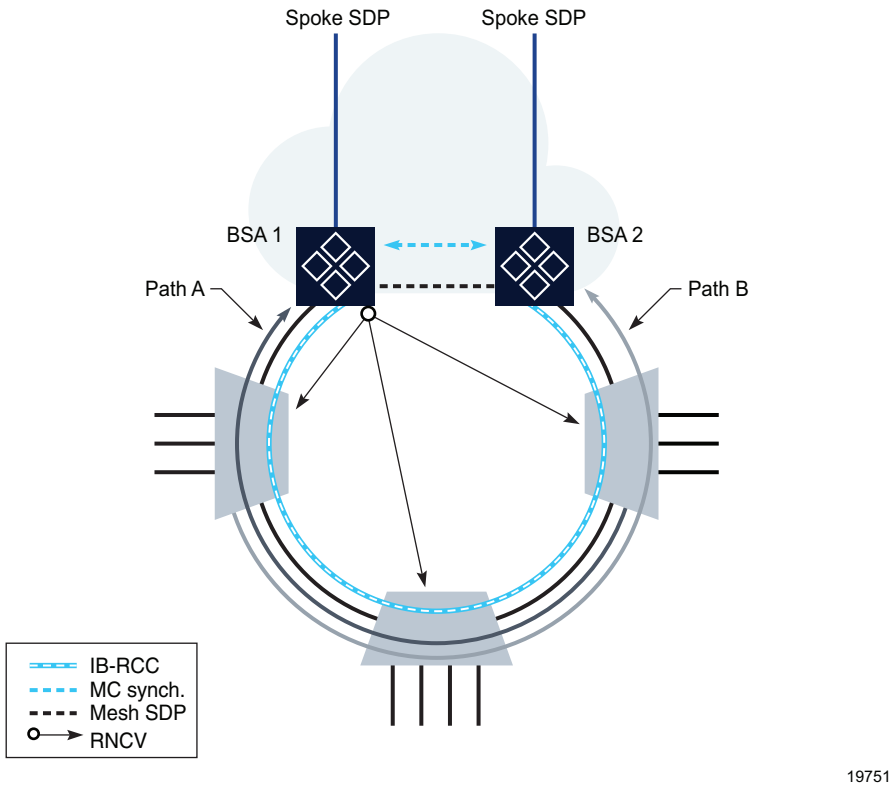

The IB-RCC is set up using a BFD session between IES or VPRN IP interfaces on BSA 1 and BSA 2. This connection requires a separate ring VLAN.

You can create an MC ring group only in an MC peer group. When only one peer in an MC peer group is managed by the NFM-P, you can configure node redundancy parameters only for the managed peer and only from the NE properties form. See Chapter 40, “MC peer groups” for information about MC peer groups.

In the steady state, the ring is fully closed and each ring node has two paths to the VPLS core; the paths are shown in as Path A and Path B in Figure 45-2, “Dual-homed ring, steady-state condition” (p. 1384) . To prevent a communication loop, a ring node can use only one path for VLAN traffic. The VLAN range can be explicitly assigned to Path B on each BSA; by default, the SAP uses Path A. The VLAN range assignment for each path must be the same on each BSA. A SAP in the conflict range is assigned to Path A, regardless of the local configuration.

Each SAP in the path that ends on the master BSA is operationally up. Each subscriber-host forwarding database entry associated with the SAP points to the SAP. Each SAP in the path that ends on the standby BSA is operationally down. Each subscriber-host forwarding database entry associated with the SAP points to the mesh SDP that connects to the master BSA.

The master BSA periodically performs a Ring Node Connectivity Verification, or RNCV, check. The loss of connectivity to a ring node does not automatically trigger a switchover to the other path; if the IB-RCC BFD session is up, the ring is considered closed, and the master and standby BSA roles do not change.

45.2.3 Broken ring condition

A broken ring condition occurs when there is a link failure or ring node failure, as shown in the following figure. In this condition, the IB-RCC is operationally down. Each ring node has only one path to the VPLS core.

Figure 45-3 Dual-homed ring, broken-ring condition

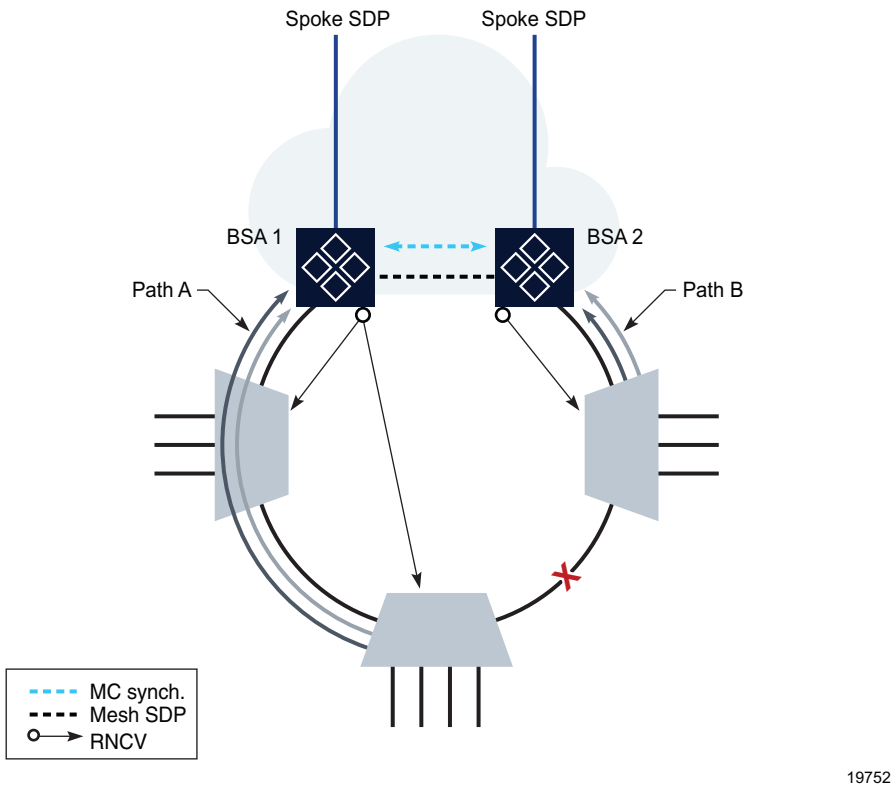

Each BSA becomes the master for the ring nodes that it can reach, and performs as described in the steady-state condition. Each L2 SAP on each BSA is operationally up, except the SAPs that are explicitly excluded from the MC ring control.

In the broken-ring condition, each BSA performs an RNCV check because each is a master. The forwarding database entry for each subscriber host behind an unreachable ring node points to the mesh SDP.

When the MC ring connectivity is restored, one BSA is the master for Path A and the other BSA is the master for Path B.

45.2.4 Object relationships

The following figure shows the MC ring and ring node object relationships with equipment and service objects.

Figure 45-4 MC ring group object relationships

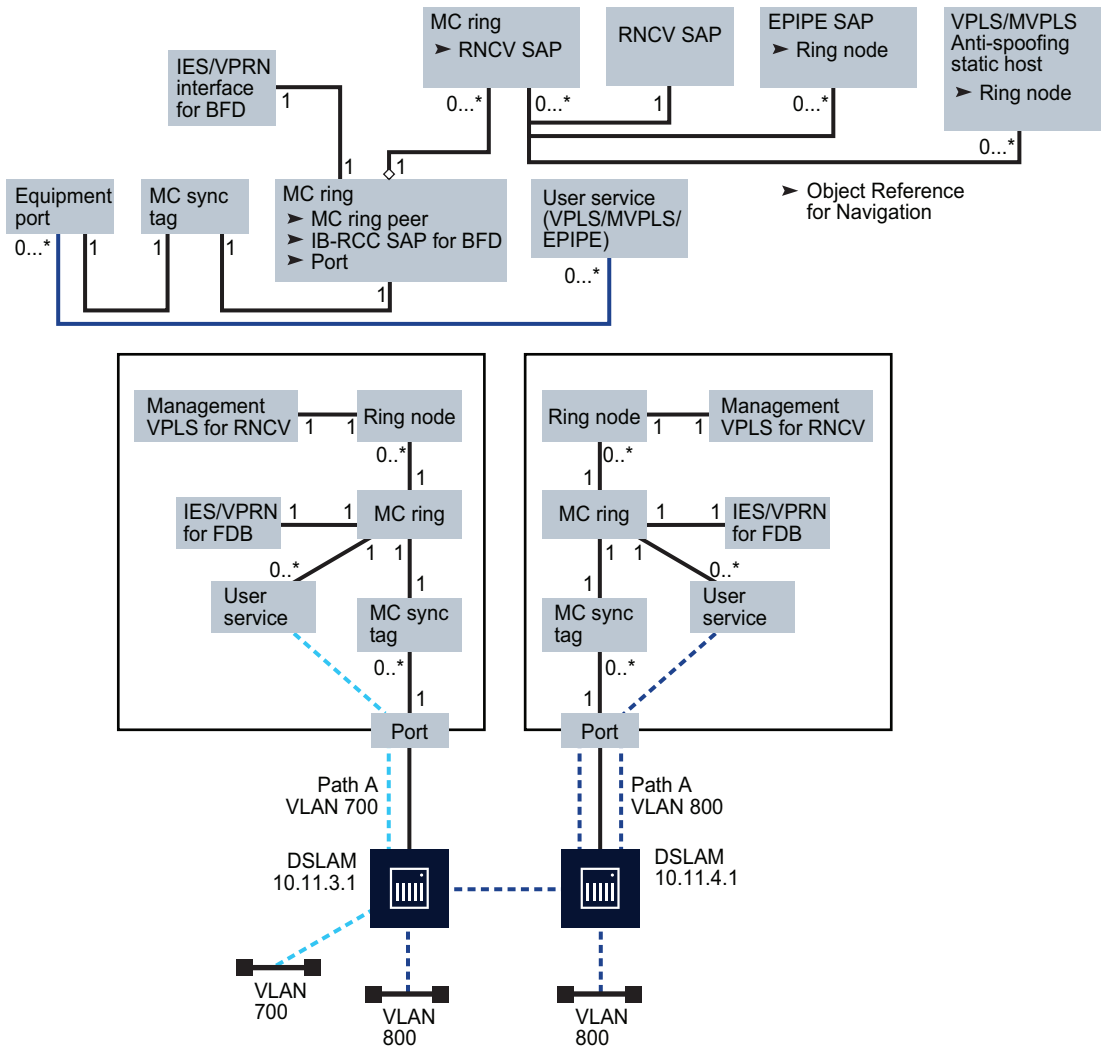

An MC ring has a one-to-one relationship with an MC synchronization tag. An MC ring uses one IES or VPRN service for IB-RCC. An MC ring can have multiple ring nodes, and each ring node can use only one SAP for RNCV.

When an MC ring group is correctly configured, each BSA L2 SAP is protected by the ring. The service that contains these protected SAPs is referred to as the user service. The service that is used for RNCV is an MVPLS that can also contain user SAPs and be a user service.

45.2.5 MC ring group and redundant VLL Epipe access

An MC ring group can connect to a VLL SAP to support redundant VLL Epipe access. The following figure shows a sample configuration using a VLL Epipe.

Figure 45-5 Dual-homed VLL Epipe configuration

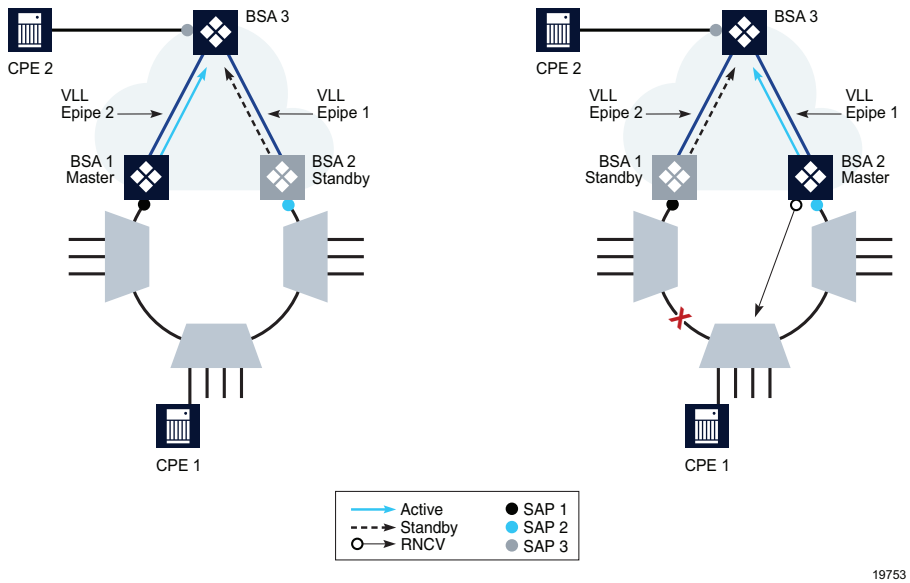

CPE 1 connects to a ring node that has access to each BSA through a VLAN that is provisioned on each ring node. The SAP on each BSA uses the same VLAN tag.

In the closed ring on the left, BSA 1 is the master and sends an active status notification to BSA 3 on a VLL Epipe. BSA 2, the standby, sends a standby status notification. Based on this information, BSA 3 chooses a path to reach CPE 2.

In the broken ring on the right, the BSA that can use RNCV to reach CPE 1 becomes the master and sends an active status notification to BSA 3 on a VLL Epipe.

In each scenario, only one SAP is operationally up at a time, and each Epipe must be operationally up. If the ring node of CPE 1 is operationally down, neither BSA can reach the ring node, so each SAP and Epipe is operationally down.

45.2.6 MC ring groups and subscriber hosts

Each subscriber host on a SAP that is protected by an MC ring group must be associated with a ring node to determine whether the subscriber host is reachable by the BSA. Each VLL Epipe used for forwarding from an MC ring group must be explicitly configured with the name of a ring node in the MC ring. See 45.6 “To configure an MC ring group for redundant VLL Epipe access” (p. 1395)

for information about configuring VLL Epipes for use with an MC ring group.

When subscriber management is enabled on a VPLS SAP, each dynamic subscriber host is automatically associated with a ring node. Static hosts on a VPLS SAP require explicit configuration. The following must be true before you can turn up an MC ring on a VPLS SAP that has one or more static hosts, or turn up a static host o. a VPLS SAP in an operational MC ring:

- Residential subscriber management is enabled.
- The following are configured on the static host:
  - the name of a ring node in the MC ring group
  - subscriber identification
  - a subscriber profile
  - an SLA profile

MC ring group management workflow and procedures

45.3 Workflow to manage MC ring groups

45.3.1 Stages

NE preconfiguration

1

Configure physical connectivity between the two BSAs and the ring nodes, as shown in Figure 45-6, “MC ring group preconfiguration” (p. 1389) . Each port must be operationally up.

Figure 45-6 MC ring group preconfiguration

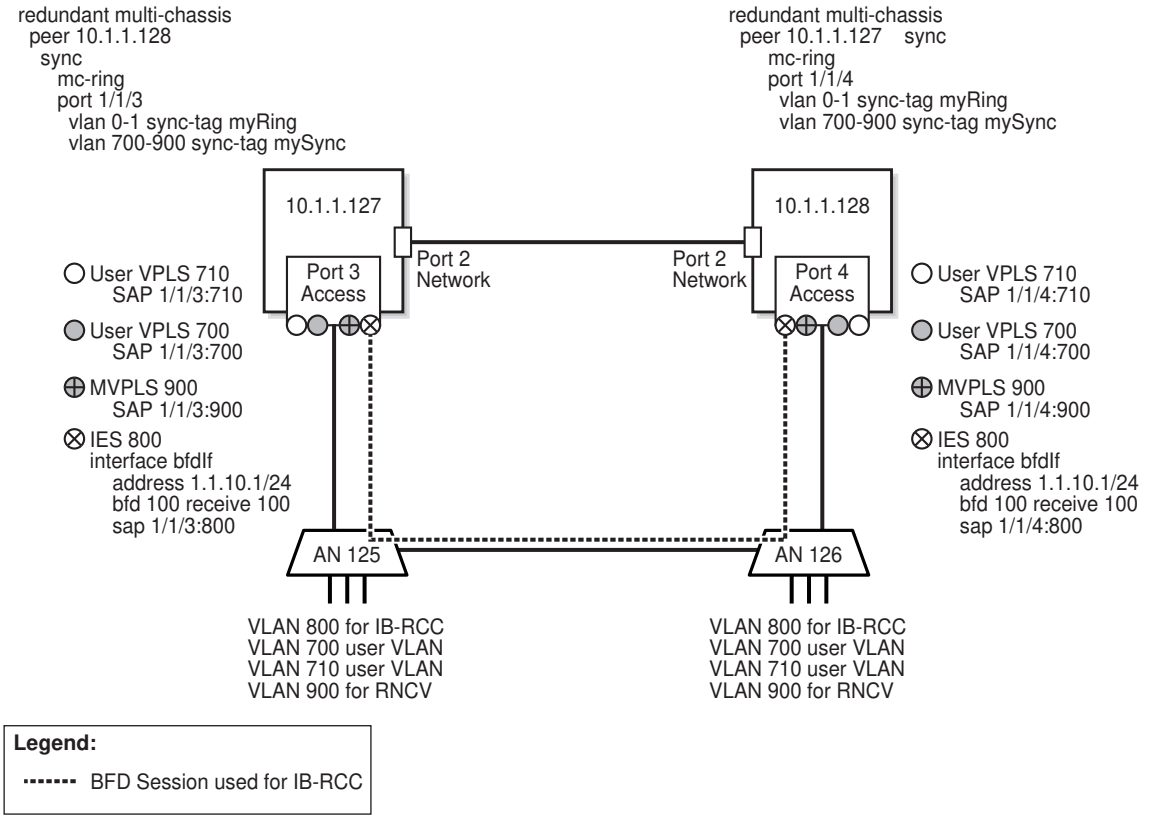

2

Configure a routing protocol on each BSA to enable IP communication; see Chapter 28, “Routing protocol configuration” for information about configuring routing protocols.

3

Configure an MVPLS for RNCV on the BSAs; see Chapter 77, “VPLS management ” for information about configuring MVPLS.

4

Create an MC peer group that has the two BSAs as members and has MC ring synchronization enabled; see Chapter 40, “MC peer groups” for information about MC peer groups.

5

In the MC peer group, create an MC synchronization group that specifies each BSA access port. The MC synchronization group can synchronize the port or a VLAN range. See Chapter 44, “MC synchronization groups” for information about MC synchronization groups.

6

Configure each ring node with the VLAN tag used for BFD to enable BFD communication through the ring and make each BSA SAP operational.

MC ring group configuration

7

Create an MC ring group that contains the BSAs, which are shown as 10.1.1.127 and 10.1.1.128 in Figure 45-7, “MC ring configuration” (p. 1391) . For each ring node within the ring group, you must specify the management IP address of a ring node as the destination IP address and assign one VLAN for connectivity checking, which is shown in Figure 45-7, “MC ring configuration” (p. 1391) as VLAN 900; see 45.4 “To create an MC ring group” (p. 1392) for more information.

Figure 45-7 MC ring configuration

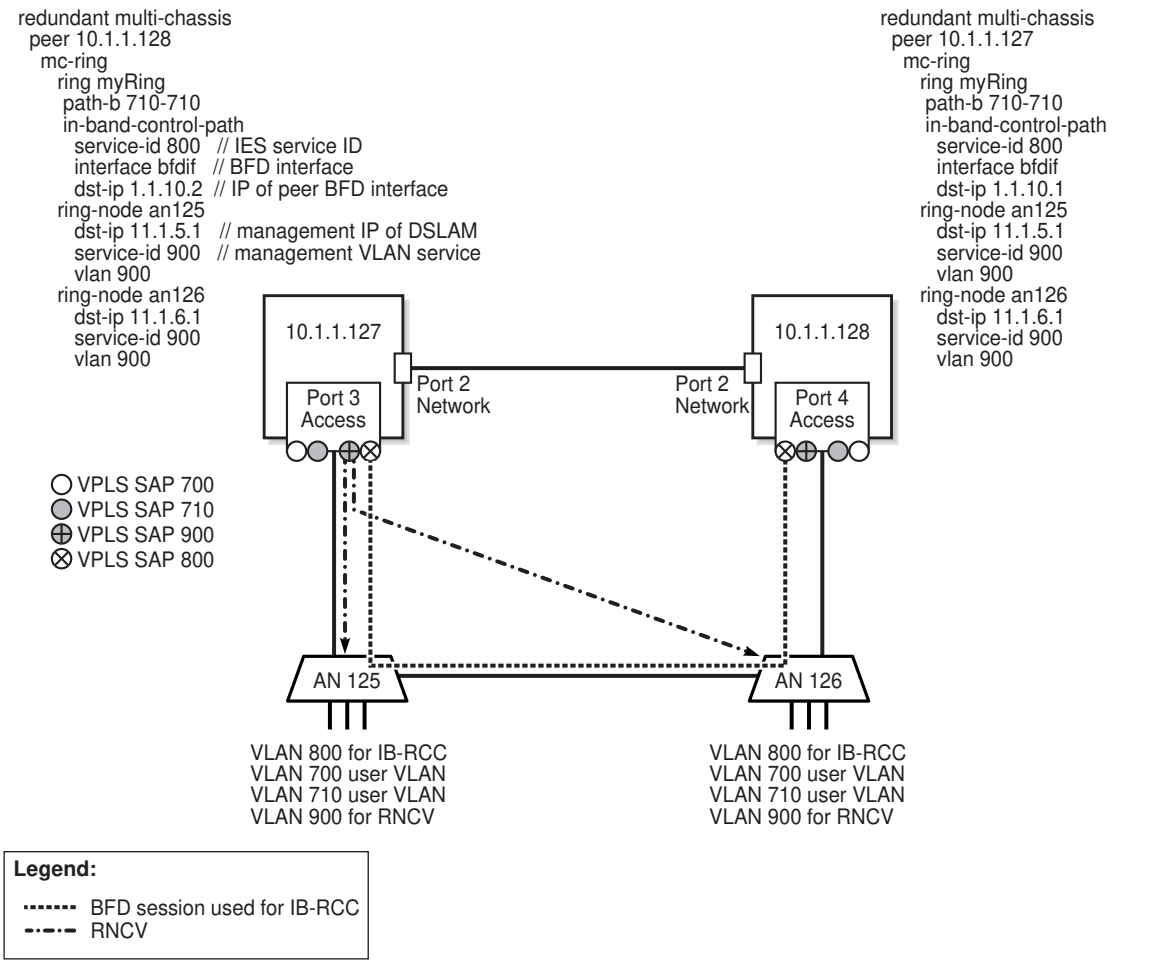

19759

8

Create the user VPLS and assign a SAP to each of the two BSA access ports. You can configure a protecting SAP on each BSA to provide SAP-level redundancy, if required; see Chapter 77, “VPLS management ” for information about protecting SAPs.

9

If the VPLS is to send traffic to an IES or a VPRN service, perform the following steps.

1. Create a spoke SDP from each VPLS site in the MC ring group to the IES or VPRN service.
2. Configure traffic forwarding to the IES or VPRN service; see 45.5 “To configure L3 forwarding from a VPLS or MVPLS to an IES or VPRN service” (p. 1394) .

10

If the VPLS is to send traffic to redundant VLL Epipes, perform the following steps.

1. Create a redundant VLL Epipe from each VPLS site in the MC ring group to the BSA that is the common endpoint of the VLL Epipes.
2. Configure traffic forwarding to each VLL Epipe service; see 45.6 “To configure an MC ring group for redundant VLL Epipe access” (p. 1395) .

11

Turn up the MC ring.

At this point, the following statements apply to the Figure 45-7, “MC ring configuration” (p. 1391) configuration.

- VPLS SAP 700 is operationally up on BSA 10.1.1.127 and operationally down on BSA 10.1.1.128.
- VPLS SAP 710 is operationally down on BSA 10.1.1.127 and operationally up on BSA 10.1.1.128.
- The operational state of the ring is Connected.

To test the configuration, you can shut down an access-node port in the ring to break the ring. As a result, the BSA SAPs should remain operationally up, but the operational state of the ring changes to Broken.

45.4 To create an MC ring group

Note: You cannot create an MC ring SAP on a VPLS site that is the endpoint of redundant spoke SDPs.

45.4.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Peer Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC peer groups is displayed.

3

Select an MC peer group and click Properties. The MC Peer Group (Edit) form opens.

4

Click on the Associated Groups tab.

5

Right-click on the MC Ring Group object and choose Create MC Ring Group. The MC Ring Group (Create) form opens.

6

Configure the parameters.

7

Click Apply.

8

Click Properties in the Multi-Chassis Ring on First Site panel. The MC Ring (Edit) form opens.

9

Configure the parameters.

10

Create one or more ring nodes, as required.

Perform the following steps:

1. Click on the Components tab.
2. Right-click on the MC Ring object and choose Create MC Ring Node. The MC Ring Node (Create) form opens.
3. Configure the parameters.
4. Click OK to save your changes and close the form.

11

Configure a VLAN range for traffic that is to use the non-default path through the MC ring, if required.

Perform the following steps:

1. Click on the Path B VLAN Range tab.
2. Click Create. The Path B VLAN Range (Create) form opens.
3. Configure the parameters.
4. Click OK to save your changes and close the form.

12

Configure a VLAN range for traffic that is not to be protected by the MC ring, if required.

Perform the following steps:

1. Click on the Exclude VLAN Range tab.
2. Click Create. The Path B VLAN Range (Create) form opens.

3. Configure the parameters.
4. Click OK to save your changes and close the form.

13

Click OK to save your changes and close the MC Ring (Edit) form.

14

Click Properties in the Multi-Chassis Ring on Second Site panel. The MC Ring (Edit) form opens.

15

Repeat Step 9 to Step 13 .

16

Click OK to save your changes and close the MC Ring Group (Create) form.

END OF STEPS

45.5 To configure L3 forwarding from a VPLS or MVPLS to an IES or VPRN service

45.5.1 Purpose

Perform this procedure to configure MC ring group traffic forwarding through a BSA spoke SDP to a BSR L3 access interface in an IES or VPRN service, as shown in Figure 45-1, “Simple subscriber aggregation network” (p. 1383) .

45.5.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search. A list of services is displayed.

3

Select a VPLS or MVPLS and click Properties. The service properties form opens.

4

Click on a site in the service navigation tree. The site properties form opens.

5

Click on the Default Gateway tab.

6

Configure the parameters, which specify the IP and MAC addresses of the default gateway to which the VPLS is to forward L3 traffic.

7

Click OK to save your changes and close the site properties form.

8

Click OK to save your changes and close the service properties form.

END OF STEPS

45.6 To configure an MC ring group for redundant VLL Epipe access

45.6.1 Purpose

Perform this procedure to configure traffic forwarding from a BSA SAP in an MC ring group through a VLL Epipe to a BSA SAP outside the ring group, as shown in Figure 45-5, “Dual-homed VLL Epipe configuration” (p. 1387) .

45.6.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search. A list of services is displayed.

3

Select the appropriate VLL Epipe service and click Properties. The Epipe Service (Edit) form opens.

4

Right-click on an L2 access interface in the service navigation tree and click Properties. The L2 Access Interface (Edit) form opens.

5

Configure the MC Ring Node parameter by clicking Select or by typing the name of a ring node that terminates on the port associated with the L2 access interface.

6

Click OK to save your changes and close the L2 Access Interface (Edit) form.

7

Click OK to save your changes and close the Epipe Service (Edit) form.

END OF STEPS

45.7 To turn up the MC rings in an MC ring group

45.7.1 Purpose

Note: You cannot administratively enable an MC ring group in one operation; you must turn up the MC rings in the MC ring group individually.
If a Problems Encountered form is displayed when you try to turn up an MC ring, you can perform one or both of the following to view information about the problem.

- Select the listed problem and click Properties.
- View the Operational State and Failure Reason indicators on the State tab of the MC ring properties form.

The NFM-P performs a series of checks to determine whether an MC ring can be turned up. The checks include verifying the following:

- The synchronization tag is configured on the BSA.
- The VLAN-level synchronization tag is valid.
- The IES or VPRN interface for the IB-RCC is configured.
- BFD is enabled on the IB-RCC interface.
- The IES or VPRN interface is operationally up.
- The IB-RCC interface is not in use by another MC ring.
- The IB-RCC interface is on the port used for MC synchronization.
- The IB-RCC destination IP address is configured.
- The IB-RCC destination IP address is not the same as the IB-RCC source IP address.
- The IB-RCC destination IP address is not in the same subnet as the IB-RCC source IP address.
- The MVPLS or VPLS site is not an endpoint for redundant spoke SDPs.

45.7.2 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Ring Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC ring groups is displayed.

3

Select an MC ring group and click Properties. The MC Ring Group (Edit) form opens.

4

Perform one of the following.

a. Turn up the MC ring using the MC ring properties form.

Note: Using this method enables you to view real-time MC ring status information while the NFM-P turns up the MC ring.

Perform the following steps:

1. Click Properties beside the Site ID parameter in the Multi-Chassis Ring on First Site panel. The MC Ring (Edit) form opens.
2. Click on the State tab.
3. Set the Administrative State parameter to Up.
4. Click Apply. The NFM-P attempts to turn up the MC ring.
5. View the dynamically updated Operational State indicator.
6. Close the MC Ring (Edit) form.
7. Click Properties beside the Site ID parameter in the Multi-Chassis Ring on Second Site panel. The MC Ring (Edit) form opens.
8. Repeat Step 4 a 2 to 6 .

b. Turn up the MC ring using the navigation tree.

Perform the following steps:

1. Click on the Components tab.
2. Right-click on the first site object and choose Turn Up. The NFM-P tries to turn up the MC ring.
3. Right-click on the second site object and click Turn Up. The NFM-P tries to turn up the MC ring.

5

Close the MC Ring Group (Edit) form.

END OF STEPS

45.8 To view the operational status of MC ring group components

45.8.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Ring Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC ring groups is displayed.

3

Select an MC ring group and click Properties. The MC Ring Group (Edit) form opens.

4

Click Properties beside a site that hosts an MC ring. The MC Ring (Edit) form opens.

5

Click on the State tab.

6

View the dynamically updated Operational State indicator.

The Operational State is one of the following:

- Unknown—cannot determine the operational state
- Shut Down—MC ring administratively down
- Config Error—MC ring misconfigured
- No Peer—peer MC ring not configured
- Connected—MC ring operational
- Broken—RNCV check fails
- Local Broken—MC ring site cannot connect to a ring node
- Conflict—MC ring configuration conflicts with another configuration
- Testing Ring—RNCV check in progress
- Peer Down—peer MC ring site inoperable
- Waiting For Peer—peer MC ring site initializing

7

Click on the Components tab.

8

Right-click on a ring node object and choose Properties. The MC Ring Node (Edit) form opens.

9

Click on the State tab.

10

View the dynamically updated information in the Local Operational State and Remote Operational State fields.

The Operational State is one of the following:

- Unknown—cannot determine operational state
- Not Provisioned—ring node configuration incomplete
- Config Error—ring node misconfigured
- Not Tested—no RNCV check result available
- Testing—RNCV check in progress
- Connected—ring node operational
- Disconnected—cannot reach ring node

11

Close the MC Ring Node (Edit) form.

12

Close the MC Ring (Edit) form.

13

Close the MC Ring Group (Edit) form.

END OF STEPS

45.9 To configure an MC ring peer on one NE

45.9.1 Purpose

Perform this procedure to configure an MC ring peer when only one NE in the MC ring is managed by the NFM-P.

Note: This procedure assumes that you have already configured an MC peer group with an associated MC ring. If only one NE is managed by the NFM-P, you must create these objects from the CLI.

45.9.2 Steps

1

On the equipment view, right-click on the NE where you want to configure an MC ring peer and choose Properties. The Network Element (Edit) form opens.

2

Click on the Redundancy tab.

3

Click on the MC Peer tab.

4

Click Search. A list of MC peers is displayed.

5

Select an MC Peer and click Properties. The MC Peer (Edit) form opens.

6

Click on the MC Rings tab.

7

Click Search. A list of MC rings is displayed.

8

Select an MC ring and click Properties. The MC Ring (Edit) form opens.

9

Configure the parameters.

10

Create one or more ring nodes, if required.

Perform the following steps:

1. Click on the Components tab.
2. Right-click on the MC Ring object and choose Create MC Ring Node. The MC Ring Node (Create) form opens.
3. Configure the parameters.
4. Click OK to save your changes and close the MC Ring Node (Create) form.

11

Configure a VLAN range for traffic that is to use the non-default path through the MC ring, if required.

Perform the following steps:

1. Click on the Path B VLAN Range tab.
2. Click Create. The Path B VLAN Range (Create) form opens.
3. Configure the parameters.
4. Click OK to save your changes and close the Path B VLAN Range (Create) form.

12

Configure a VLAN range for traffic that is not to be protected by the MC ring, if required.

Perform the following steps:

1. Click on the Exclude VLAN Range tab.
2. Click Create. The Exclude VLAN Range (Create) form opens.
3. Configure the parameters.
4. Click OK to save your changes and close Exclude VLAN Range (Create) form.

13

Close the MC Ring (Edit) form.

14

Close the MC Peer (Edit) form.

15

Close the Network Element (Edit) form.

END OF STEPS

45.10 To delete an MC ring group

CAUTION

Service Disruption

Deleting an MC ring group removes all child objects associated with the MC ring group, such as the MC rings and ring nodes.

Note: You must set the administrative state of an MC ring to Down before you can delete the MC ring.
Before you can set the administrative state of an MC ring to Down, you must administratively disable each VPLS or MVPLS SAP in the ring, regardless of the path used by each.

45.10.1 Steps

1

Choose Manage→Redundancy→Node Redundancy from the NFM-P main menu. The Manage Node Redundancy form opens.

2

Choose MC Ring Group (Multi-Chassis) from the drop-down menu and click Search. A list of MC ring groups is displayed.

3

Select an MC ring group and click Properties. The MC Ring Group (Edit) form opens.

4

Click Properties beside the Site ID parameter in the Multi-Chassis Ring on First Site panel. The MC Ring (Edit) form opens.

5

Click on the State tab.

6

Set the Administrative State parameter to Down.

7

Click Apply. The NFM-P shuts down the MC ring.

8

Close the MC Ring (Edit) form.

9

Click Properties beside the Site ID parameter in the Multi-Chassis Ring on Second Site panel. The MC Ring (Edit) form opens.

10

Repeat Step 5 to Step 8 .

11

Select the MC ring group and click Delete.

12

Click View Dependencies. A Warning form opens.

13

View the dependency information.

14

Select the I understand the implications of this action check box.

15

Click Yes. The NFM-P deletes the MC ring group and the associated configuration on each member NE.

16

Close the Manage Node Redundancy form.

END OF STEPS

46 Synchronization management

Synchronization management overview

46.1 NFM-P synchronization manager

46.1.1 Overview

The NFM-P synchronization manager provides a centralized view of timing synchronization across the network. The centralized view provides a common list of master and slave synchronization peers in a synchronization domain. The view includes a correlated list of alarms against synchronization objects in the domain, such as PTP clocks, ports, and peers.

You can use the synchronization manager to assign IP path monitors to peers between NEs configured with IEEE 1588 PTP clocks. The synchronization manager provides navigation to related NFM-P objects, such as NEs that are configured with IEEE 1588 PTP clocks, and to related CPAM objects, such as IP path monitors.

Note: The synchronization manager can be used to view and monitor IEEE 1588 PTP clocks. You can configure timing synchronization and IEEE 1588 PTP clocks through the NFM-P NE properties forms or through the device CLI.

46.1.2 Synchronization domains

Synchronization domains allow you to view and manage the PTP peers or SyncE sites in a domain. You can use synchronization domains to assign IP path monitors to PTP peers. You can also filter unmonitored peers and navigate to associated IP path monitors.

When the NFM-P discovers a PTP peer or SyncE site, it is assigned to a synchronization domain based on the domain ID. When the NFM-P discovers a PTP peer or SyncE site with a domain ID that is not shared with an existing synchronization domain, a synchronization domain with that domain ID is automatically created. The domain ID is configured on the IEEE 1588 PTP clock.

The following events can result in the creation of a synchronization domain.

- An NE is discovered with configured IEEE 1588 PTP clocks.
- An IEEE 1588 PTP clock is created on a discovered NE.
- The domain ID on an IEEE 1588 PTP clock is modified.

You cannot create synchronization domains manually, and you can delete a synchronization domain only when the domain does not contain peers or SyncE sites. Alarms generated for a PTP clock, port, or peer are propagated to the appropriate synchronization domain.

46.1.3 Synchronization groups

You can create synchronization groups to view a subset of PTP peers separately from the rest of the synchronization domain. A synchronization domain can contain up to 500 synchronization groups. When you create the group, you choose NEs configured with IEEE 1588 PTP clocks and

the NFM-P generates a list of constituent peers from the NEs. All of the path monitoring functionality that is available at the synchronization domain level is also available at the synchronization group level.

46.1.4 Peer remote site ID

When the NFM-P discovers a PTP peer, the peer IP address is used to determine the system IP address of the far-end NE. The IP address is displayed in the synchronization manager as the peer remote site ID. If the NFM-P discovers a system or network IP address that matches the peer IP address of a discovered PTP peer, the peer remote site ID is recalculated. The peer remote site ID is set to 0.0.0.0 if it cannot be mapped to a managed NE in the network or if it maps to more than one managed NE.

46.1.5 Timing synchronous links

The synchronization manager lists all timing synchronous links in the network related to a selected PTP peer or SyncE site. The list includes links that are not currently in use. You can click the Find Sync Time Sources button on the PTP peer or SyncE site properties form to view a list of timing synchronous links in the network that are currently being used by the selected site.

From the physical topology map, you can right-click an NE and select Highlight Sync Time Sources to view a graphical representation of the timing synchronous links currently being used by the selected NE.

46.1.6 IP path monitors

You can use the synchronization manager to assign unidirectional and bidirectional IP path monitors to PTP peers. The PTP Peers tabs on the synchronization domain and synchronization group forms include an Unmonitored tab that displays peers that are not assigned an IP path monitor. You can navigate to the IP path monitor object from the synchronization manager.

Alarms that are generated for IP path monitor objects, such as reachability, OAM and path change alarms, are propagated to the respective PTP peers. See the *NSP NFM-P Control Plane Assurance Manager User Guide* for more information about IP path monitors.

46.1.7 SyncE

The synchronization manager supports SyncE. The NFM-P automatically assigns NEs with Sync-E enabled to the default Sync-E synchronization domain. You can view SyncE sites and associated daughter cards, including correlated alarms. Both the far end and near end NEs must be managed by the NFM-P for SyncE to work.

The properties forms for SyncE sites includes a check box that indicates if the site is also enabled with PTP. The properties forms for PTP peers includes a check box that indicates if the peer is also enabled with SyncE. This flag is processed only if there is a valid peer remote site ID.

46.2 Synchronization topology

46.2.1 Overview

You can use the NFM-P physical topology map to view a graphical presentation of SyncE and PTP peers in your network. The synchronization topology offers snapshot view of timing sources as well

as the quality of timing in the network.

The SyncE highlights on the physical topology map are not updated dynamically.

Perform one of the following to refresh the graphical presentation:

- De-select and re-select the highlight session from the Highlight Sessions tab in the map legend
- Draw a new sync source highlight session

46.2.2 Timing sources highlight

You can use the physical topology map to highlight the timing sources for a selected NE. Right-click the NE and select Highlight Sync Time Sources. PTP timing source links are displayed as a dashed line and SyncE timing source links are displayed as a solid line. An arrow indicates the directional path of the timing source.

Note: Before using the timing source highlight feature, you must ensure that the physical topology is displaying physical links between the NEs.

Note: The topology map displays up to 20 hops for a timing source highlight. If the timing source exceeds 20 hops, you can use the Find Sync Time Sources button on the PTP peer or SyncE site properties form to generate a complete list of timing sources.

Figure 46-1 PTP timing source highlight

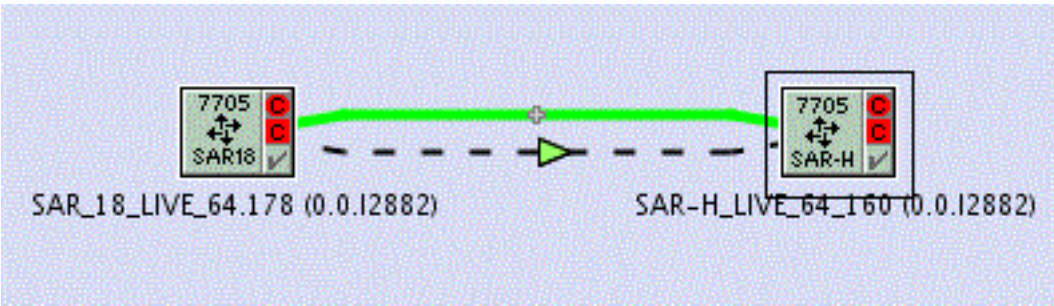

Figure 46-2 SyncE timing source highlight

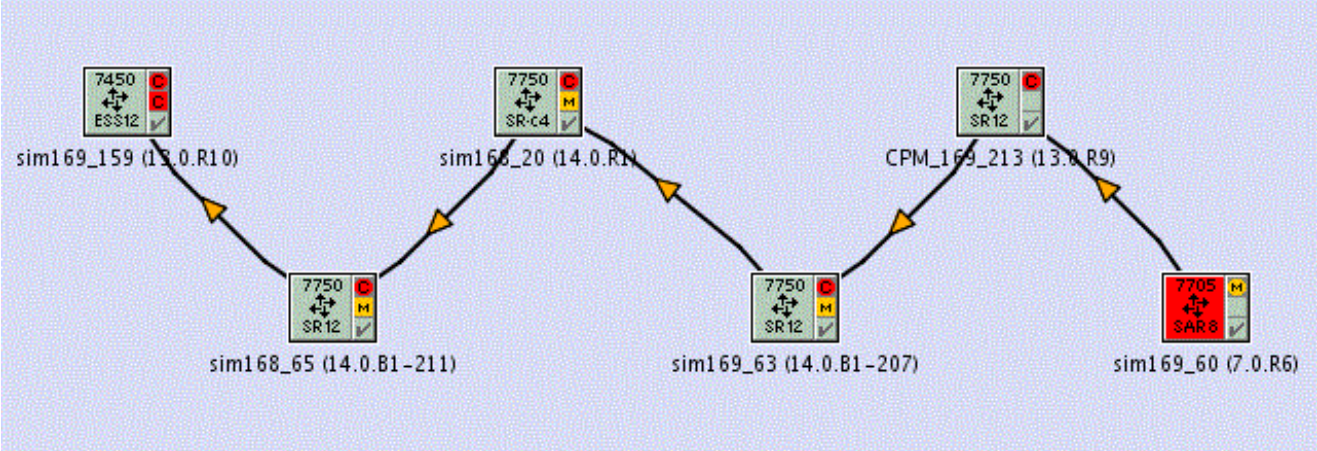

46.2.3 Timing quality

You can use the NFM-P physical topology map to view a graphical presentation of timing quality in your network. Right-click on the topology map and select Quality of Timing→Color Quality of Timing. The SyncE sites and PTP peers are colored based on the quality of the timing source.

The following colors are used to represent timing quality:

- Grey — reachable; or master-locked; or slave
- Yellow — resyncing; or master holdover; or acquiring
- Purple — suspended
- Red — unreachable; or resync failed; or master free run; or not present

In the following figure, sim169_60 is in free run, sim169_63 is in holdover, and all other NEs are master-locked.

Figure 46-3 Timing quality coloring

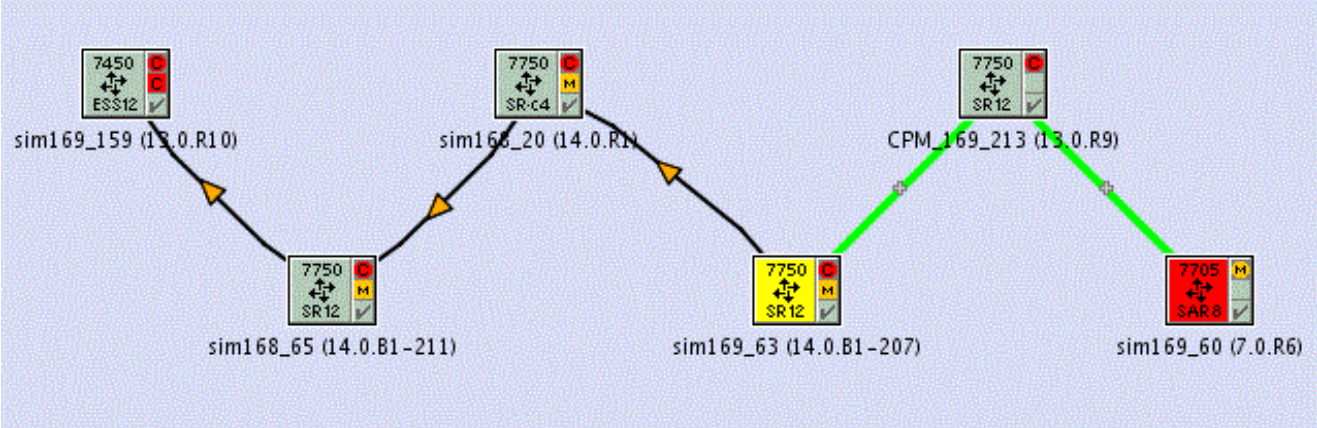

Synchronization management workflow

46.3 Workflow for synchronization management

46.3.1 General information

This workflow assumes that PTP peers are configured. See “Configuring shelf objects using the NFM-P” (p. 453) for more information about configuring timing and IEEE 1588 PTP clocks.

46.3.2 Stages

1

Configure synchronization domains and create synchronization groups. See 46.4 “To configure a synchronization domain and create synchronization groups” (p. 1411) .

2

Create unidirectional and bidirectional IP path monitors, as required. See 46.5 “To create IP path monitors for PTP peers” (p. 1412) .

3

View IP path monitors, as required. See 46.6 “To view an IP path monitor from the synchronization manager” (p. 1414) .

4

Perform path monitoring procedures, as required. See the *NSP NFM-P Control Plane Assurance Manager User Guide*.

5

Perform synchronization topology monitoring procedures, as required. See 46.8 “To use the synchronization topology map” (p. 1417).

Synchronization management procedures

46.4 To configure a synchronization domain and create synchronization groups

Note: A synchronization domain is created automatically when a new IEEE 1588 PTP clock with a unique domain ID is created or discovered. You cannot create synchronization domains manually.

46.4.1 Steps

1

Choose Tools→Synchronization Manager from the main menu. The Synchronization Manager form opens.

2

Configure the filter criteria, if required, and click on the Search button.

3

Choose a synchronization domain and click on the Properties button. The Synchronization Domain (Edit) form opens with the General tab displayed.

4

Configure the parameters:

- Displayed Name
- Description

5

Click on the Sync Groups tab button.

6

Click on the Create button. The Synchronization Group (Create) form opens with the General tab displayed.

7

Configure the parameters:

- Group ID
- Auto-Assign ID
- Displayed Name
- Description

8

Click on the Network Elements tab button.

9

To add NEs to the synchronization group:

1. Click on the Add button. The Select Network Elements form opens.
2. Configure the filter criteria, if required, and click on the Search button.

   Note:

   Only NEs with configured IEEE 1588 PTP clocks are displayed.
3. Choose one or more NEs and click on the OK button. The Select Network Elements form closes.

10

Click on the OK button. A dialog box opens.

11

Click on the OK button to confirm the changes. The Synchronization Group (Create) form closes.

12

Repeat Step 6 to Step 11 to create more synchronization groups, as required.

13

Click on the Apply button. A dialog box opens.

14

Click on the OK button. The Synchronization Domain (Edit) form closes.

15

Close the Synchronization Manager form.

END OF STEPS

46.5 To create IP path monitors for PTP peers

46.5.1 Steps

1

Choose Tools→Synchronization Manager from the main menu. The Synchronization Manager form opens.

2

Configure the filter criteria, if required, and click on the Search button.

3

Choose a synchronization domain and click on the Properties button. The Synchronization Domain (Edit) form opens with the General tab displayed.

4

Perform one of the following:

a. To create IP path monitors at the synchronization domain level, click on the PTP Peers tab button.

b. To create IP path monitors at the synchronization group level:

   1. Click on the Sync Groups tab button.
   2. Configure the filter criteria, if required, and click on the Search button.
   3. Choose a synchronization group in the list and click on the Properties button. The Synchronization Group (Edit) form opens with the General tab displayed.
   4. Click on the PTP Peers tab button.

5

Under the Unmonitored tab, perform one of the following:

a. To create unidirectional IP path monitors, click on the IP tab button.

b. To create bidirectional IP path monitors, click on the Bidirectional IP tab button.

6

Configure the filter criteria, if required, and click on the Search button.

7

Choose one or more peers.

8

Click on the Create Path Monitor button. The NFM-P creates an IP path monitor and the chosen peer is removed from the list.

Note: The IP Path Monitors are created only for PTP peers for which the peer remote site ID is resolved and mapped to a CPAM-managed router.

9

Close the Synchronization Group (Edit) form, if applicable.

10

Close the Synchronization Domain (Edit) form.

11

Close the Synchronization Manager form.

END OF STEPS

46.6 To view an IP path monitor from the synchronization manager

46.6.1 General information

Perform this procedure to navigate to an assigned IP path monitor object from the synchronization domain or synchronization group forms.

Note: You can also view the properties of an IP path monitor from the Synchronization component tree of an NE. See 46.7 “To use the Synchronization component tree” (p. 1415) for more information.

46.6.2 Steps

1

Choose Tools→Synchronization Manager from the main menu. The Synchronization Manager form opens.

2

Configure the filter criteria, if required, and click on the Search button.

3

Choose a synchronization domain and click on the Properties button. The Synchronization Domain (Edit) form opens with the General tab displayed.

4

Perform one of the following:

a. To view IP path monitors at the synchronization domain level, click on the PTP Peers tab button.

b. To view IP path monitors at the synchronization group level, perform the following:

1. Click on the Sync Groups tab button.
2. Configure the filter criteria, if required, and click on the Search button.
3. Choose a synchronization group and click on the Properties button. The Synchronization Group (Edit) form opens with the General tab displayed.
4. Click on the PTP Peers tab button.

5

Click on the All tab button.

6

Configure the filter criteria, if required, and click on the Search button.

7

Perform one of the following:

a. To view a unidirectional IP path monitor:
   1. Choose one or more monitored peers and click on the IP Path Monitors button. A dialog box opens.
   2. Click on the Yes button. The IP Path Monitor (Edit) form opens with the General tab displayed. If you specified more than one peer in 1 , several forms open.
   3. View the IP path monitor, as required. See the *NSP NFM-P Control Plane Assurance Manager User Guide* for more information about IP path monitors.
   4. Close the IP Path Monitor (Edit) form.

b. To view a bidirectional IP path monitor, perform the following:
   1. Choose one or more monitored peers and click on the Bidirectional IP Path Monitors button. A dialog box opens.
   2. Click on the Yes button. The Bidirectional IP Path Monitor (Edit) form opens with the General tab displayed. If you specified more than one peer in 1 , several forms open.
   3. View the bidirectional IP path monitor, as required. See the *NSP NFM-P Control Plane Assurance Manager User Guide* for more information about bidirectional IP path monitors.
   4. Close the Bidirectional IP Path Monitor (Edit) form.

8

Close the Synchronization Group (Edit) form, if applicable.

9

Close the Synchronization Domain (Edit) form.

10

Close the Synchronization Manager form.

END OF STEPS

46.7 To use the Synchronization component tree

46.7.1 General information

Perform this procedure to access the Synchronization component tree of an IEEE 1588 PTP synchronization-enabled NE. The component tree provides a hierarchical view of PTP Clocks, PTP Ports, PTP Peers, and Path Monitors associated with the NE.

You can also use the Synchronization component tree to create or delete IEEE 1588 PTP Clocks, as well as create IP Path Monitors.

46.7.2 Steps

1

Perform one of the following:

a. Right-click on an NE in the Equipment or Routing navigation tree or in any NE list form and choose Properties from the contextual menu. The Network Element (Edit) properties form opens.

b. Right-click on an NE in the topology map and choose Properties from the contextual menu. The Network Element (Edit) properties form opens.

c. Choose Tools→Synchronization Manager from the main menu. The Synchronization Manager form opens.

1. Configure the filter criteria, if required, and click on the Search button.
2. Choose a Synchronization Domain and click on the Properties button. The Synchronization Domain (Edit) form opens, with the General tab displayed.
3. If you want to access the Synchronization component tree from the synchronization group level, perform 4 and 5 . Otherwise proceed to 6 .
4. Click on the Sync Groups tab and click on the Search button.
5. Choose a Synchronization Group and click on the Properties button. The Synchronization Group (Edit) form opens, with the General tab displayed.
6. Click on the PTP Peers tab button. A list of IEEE 1588 PTP Peers appears in the Unmonitored sub-tab.
7. Click the All sub-tab and choose a PTP peer from the list, then click on the NE Properties button. The associated Network Element (Edit) properties form of the selected PTP Peer opens.

2

Perform one or more of the following from the Synchronization item in the navigation tree:

a. View properties of existing PTP Clocks, PTP Ports, PTP Peers, and Path Monitors associated with the NE by expanding the Synchronization item in the component tree to the required level. Click on an item in the tree to display its properties in the form. See 46.6 “To view an IP path monitor from the synchronization manager” (p. 1414) for more information on viewing IP Path Monitors.

b. To create a new PTP Clock on the NE, right-click the IEEE 1588 PTP item under Synchronization. The IEEE PTP Clock (Create) form opens. See “Configuring shelf objects using the NFM-P” (p. 453) for procedural information on creating IEEE 1588 PTP Clocks and PTP Ports.

c. To create a new PTP Peer for an IEEE PTP Clock. See “Configuring shelf objects using the NFM-P” (p. 453) for procedural information.

d. To create a new Path Monitor, right-click the required port under Synchronization and choose

either Create IP Path Monitor or Create Bidirectional IP Path Monitor from the contextual menu. See 46.5 “To create IP path monitors for PTP peers” (p. 1412) for more information.

3

Close the Synchronization Group (Edit) form, if applicable.

4

Close the Synchronization Domain (Edit) form, if applicable.

5

Close the Synchronization Manager form, if applicable.

END OF STEPS

46.8 To use the synchronization topology map

46.8.1 Steps

Open a physical topology map

1

Choose Tools→Synchronization Manager from the main menu. The Synchronization Manager form opens.

2

Click Sync Timing Topology View. The Physical Topology map opens.

View timing synchronization sources

3

Right-click an NE and select Highlight Sync Time Sources.

4

A dialog confirms that the synchronization timing source highlight displays a snapshot of network timing. Click Yes.

PTP timing source links are displayed as a dashed line and SyncE timing source links are displayed as a solid line. An arrow indicates the directional path of the timing source.

View quality of timing

5

Right-click the topology map and select Quality of Timing→Color Quality of Timing.

The SyncE sites and PTP peers are colored based on the quality of the timing source.

END OF STEPS

47 Cellular domain management

47.1 Overview

47.1.1 Cellular domains

The NFM-P supports the management of cellular domains that group together devices such as the 7705 SAR-Hm for remote management. A cellular domain groups devices that have similar cellular network characteristics. If devices with dual SIM cards are in use, the cellular domain groups devices with the same primary and backup cellular network.

All NEs in a cellular domain must be running the same NE software release.

Note: All NEs in the domain must have the same combination of carriers. If, for example, two NEs have SIM cards for the same carriers but NE1 will use Carrier A as primary and NE2 will use Carrier B as primary, the two NEs must be in separate domains.

A cellular domain provides the following:

- a management domain for a group of cellular-connected NEs with common management attributes
- automatic discovery protocol (ADP), which facilitates remote device discovery and management
- multiple head-end nodes, which are used to reach the NEs in the cellular domain
- definition of an optional in-band management VPRN service of the cellular domain used to manage the NEs
- optional NGE domain used to secure the NEs cellular interface for all NEs in the cellular domain
- wireless carrier redundancy with optional dual SIM operation
- chassis-level security using a common SIM PIN (Personal Identification Number) per carrier for all NEs in a domain

A cellular domain acts as a template for ADP device discovery within one or more predefined subnets, initial configuration, and encryption of the NEs in the domain using NGE.

After the ADP process completes, the NEs in one or more predefined subnets are included in the cellular domain and managed using the cellular domain configurations. You can add new NEs to the domain by re-enabling ADP, which enables the deployment of the required configuration, for example, BGP or NGE, to the devices.

Note: An NE can belong to only one cellular domain.

Dual SIM deployment

The NFM-P supports the use of two SIM cards in a 7705 SAR-Hm, each with a different wireless carrier, for WAN redundancy. One SIM is active at a time.

Switching from one SIM to the other can be automatic or manual. With automatic switchover, you can choose which SIM is primary and secondary, and configure SIM switchover criteria. For example, the BGP operational state associated with the cellular port can be used as a criterion for

determining when a SIM switchover should occur. If the BGP operational state is down for a specified interval, then a SIM switchover occurs.

A SIM switchover is service affecting. Overly frequent switchovers will impact continuous service operation.

Dual SIM deployment is configured at the cellular domain level. All NEs in a dual SIM domain must have two working SIM cards.

The NFM-P does not support conversion of a single SIM cellular domain to a dual SIM cellular domain.

Head-end nodes

Each cellular domain must have at least one head-end node. A single SIM domain can have from one to four head-end nodes; a dual SIM domain can have from one to eight.

Head-end nodes serve as an intermediary for communication with the NFM-P, terminate Layer 2 and Layer 3 services, forward services to other NEs, and optionally define the gateway sites used by the NGE domain associated with the cellular domain. Each head-end node in a cellular domain must also be a gateway site to the NGE domain, if NGE is in use.

A head-end node can be a VSR, 7705 SAR, or 7750 SR, and must be discovered and managed by the NFM-P to be included in a cellular domain. If NGE is required, the head-end node cannot be a 7750 SR.

47.1.2 Cellular domain operation modes

A cellular domain in single SIM layout has the following operation modes:

- static cellular system mode — in-band management via 7705 SAR-Hm PDN interface
- static cellular interface mode — in-band management via a VPRN service and a private system IP address on the 7705 SAR-Hm
- dynamic cellular interface mode — in-band management via a VPRN service and a private system IP address on the 7705 SAR-Hm

In a dual SIM layout, both SIMs automatically operate in dynamic cellular interface mode.

In any operation mode, an operator-created XML file can be used to specify the devices for discovery.

In static or dynamic interface mode, if the ADP System IP Address parameter is set to something other than Use XML, you can also specify a pool of IP addresses for assignment to the discovered devices.

Static system mode

In static system mode, ADP can be configured to do one of the following when an SNMPv2 trap is received from a 7705 SAR-Hm in an ADP subnet:

- Learn the system address and IMSI of the 7705 SAR-Hm:
  The NFM-P adds the device to the associated cellular domain and initiates the ADP.
- Verify the system address of the 7705 SAR-Hm:

If the IMSI and system address match a device specified in an operator-created XML file, the NFM-P adds the device to the cellular domain and initiates the ADP. “Specifying devices using an XML file” (p. 1422) describes how to structure and import the XML file.

Note: When ADP is used in static system mode, the system address of a 7705 SAR-Hm must match the cellular interface address.

Static interface mode

In static interface mode, ADP does one of the following:

- imports an operator-created XML file that specifies each 7705 SAR-Hm device to discover

  When an SNMPv2 trap is received from a 7705 SAR-Hm, the NFM-P initiates the ADP for the device if the IMSI value matches an entry in the file. “Specifying devices using an XML file” (p. 1422) describes how to structure and import the XML file.
- assigns an IP address from a user-specified pool of addresses

Note: The NFM-P reuses IP addresses in a pool if the IP address is not used when the ADP Domain site is deleted and assigned.

To avoid issues when reusing IP addresses, the IP address should be deleted from the 7705 SAR-Hm and the node should be unmanaged and deleted from NFM-P before deleting the ADP domain site.

When the addresses in an IP-address pool are exhausted, the NFM-P raises an alarm, and ADP discovery is halted. In such a case, you can add a new system IP pool to resume ADP discovery.

Note: If a subnet has an associated IP-address pool, and an IP address is associated with the device IMSI in an ADP XML file, the IP address in the XML file is assigned to the device.

Note: When ADP is used in static interface mode, the system address of a 7705 SAR-Hm must be unique and different from the cellular interface address.

Dynamic interface mode

In dynamic interface mode, ADP does one of the following:

- imports an operator-created XML file that specifies each 7705 SAR-Hm device to discover

  When an SNMPv2 trap is received from a 7705 SAR-Hm, the NFM-P initiates the ADP for the device if the IMSI value matches an entry in the file. “Specifying devices using an XML file” (p. 1422) describes how to structure and import the XML file.
- assigns a system IP address from a user-specified pool of addresses

Note: The NFM-P reuses IP addresses in a pool if the IP address is not used when the ADP Domain site is deleted and assigned.

To avoid issues when reusing IP addresses, the IP address should be deleted from the 7705 SAR-Hm and the node should be unmanaged and deleted from NFM-P before deleting the ADP domain site.

When the addresses in a System IP-address pool are exhausted, the NFM-P raises an alarm, and ADP discovery is halted. In such a case, you can add a new system IP pool to resume ADP discovery.

Note: If a subnet has an associated IP-address pool, and an IP address is associated with the device IMSI in an ADP XML file, the IP address in the XML file is assigned to the device.

Specifying devices using an XML file

Identifiers for 7705 SAR-Hm devices can be specified in an XML file, and imported for use in either operation mode. The XML file requires an ADP element and one node element for each device to discover. Each node element has IMSI, systemName, and systemAddress attributes.

Note: The systemAddress attribute is optional, depending on the IP allocation configuration and the operation mode of the cellular domain. The systemName attribute is also optional.

In both static and dynamic cellular interface mode, the systemAddress is the private IP address reachable via the in-band VPRN service.

The following is an example of an ADP XML file:

```
<?xml version="1.0" encoding="UTF-8"?>
<ADP>
<node IMSI="310150123456720"
systemAddress="10.10.10.20"
systemName="test"/>
</ADP>
```

47.5 “To configure a cellular domain with single SIM deployment” (p. 1440) and 47.6 “To configure a cellular domain with dual SIM deployment” (p. 1444) describe how to specify the devices for ADP discovery using an XML file.

47.1.3 Domain subnets

You can create separate subnets in a cellular domain. A domain subnet represents a group of NEs with a cellular interface IP address that exists with the specified subnet. For each subnet in a domain, a BGP dynamic neighbor is created on each head-end node in the domain. The dynamic neighbor can accept remote 7705 SAR-Hm peers as they become available. The deletion of a subnet also deletes the BGP dynamic neighbor from all head-end node BGP groups.

Note: If the head-end node is a 7705 SAR, BGP neighbors are not created automatically when subnets are added. The 7705 SAR does not support dynamic neighbors. Static neighbors must be manually created.

ADP must be enabled on a subnet that has new devices to discover. ADP is initiated based on the cellular domain operation mode; see 47.1.11 “ADP discovery process” (p. 1429).

Note: You cannot delete a subnet that has ADP enabled at the domain or subnet level.

47.1.4 7705 SAR-Hm security during ADP

For additional security, each new 7705 SAR-Hm includes an information card that names the chassis serial number and a unique administrator password. During cellular domain configuration, you must specify the serial number and password of each such device to enable ADP access to the device.

A cellular domain can also contain older devices that have a common default administrator password. To support such devices, the cellular domain configuration must include the default password for ADP communication with the older devices.

The administrative account credentials for a device in a cellular domain must meet the following requirements.

- The User Name value on the NE User properties form must match the User Name value for CLI and NETCONF access in the ADP mediation security policy.
- The Password value for console or FTP access on the NE User properties form must match the Password value for CLI and NETCONF access in the ADP mediation security policy.

  To simplify the configuration, you can use the same value for the SNMPv3 user authentication and privacy passwords.

Configuring secure 7705 SAR-Hm ADP

During cellular domain configuration, you can specify the serial numbers and passwords manually, or import the passwords from a file, as described in 47.3 “To create an ADP password mapping file” (p. 1435).

47.1.5 Cellular domain security with NGE

You can secure a cellular domain by binding the cellular domain to an NGE domain. The encryption status of PDN interfaces in the cellular domain is determined by the encryption status of the gateway interfaces of the gateway sites of the NGE domain.

Each head-end node in each cellular domain that belongs to the NGE domain must also be a gateway site in the NGE domain. The NEs in each cellular domain of an NGE domain are listed as sites in the NGE domain. Multiple cellular domains can point to the same NGE domain.

You can unbind a cellular domain from an NGE domain if the NGE domain is not encrypting. Before the cellular domain can be removed, the RI NGE encryption on each 7705 SAR-Hm NE in the cellular domain must be stopped.

47.1.6 Enhanced NE security mode

To prevent unwanted tampering with security settings on any 7705 SAR-Hm, you can enable enhanced NE security mode, which protects all 7705 SAR-Hm devices in the network using stringent security constraints that cannot be altered by an NFM-P operator. Attempts to do so are blocked, and generate NFM-P alarms.

Note: Enabling enhanced NE security mode affects each 7705 SAR-Hm in the managed network. Also, in order to disable the function, you must first unmanage each managed NE of any type in the entire network.

When enhanced NE security mode is enabled, the NFM-P enforces the following security constraints for each 7705 SAR-Hm device:

- The enabled SSH versions cannot include SSHv1.
- Exponential backoff must be enabled.
- The password policy must include specific criteria.

Note: The function does not validate any password, only the conformance of the local NE password policy.

47.4 “To enable enhanced NE security mode” (p. 1437) describes how to configure and enable the mechanism.

Note: The local and global NE password policy definitions are verified against the required password criteria before being applied to the NEs.

The SAR-Hm Enhanced Security indicator on the NFM-P System Preferences form shows whether enhanced NE security mode is enabled.

Implementation

If enhanced NE security mode is enabled, the NFM-P raises an alarm against any 7705 SAR-Hm whose configuration violates any listed security constraint. The alarm is raised regardless of whether a device is discovered before or after enhanced NE security mode is enabled. An alarm is also raised if a managed NE configuration is changed via CLI in a way that violates a constraint.

When you invoke the ADP process and enhanced NE security mode enabled, the NFM-P first verifies the password of the SNMPv3 user in the cellular domain mediation policy against the required password complexity rules. If the password violates any complexity rule, ADP does not proceed.

After ADP completes:

- Any client GUI configuration that attempts to enable SSH1 on an NE is blocked. If the configuration is performed using a CLI, the NFM-P raises an alarm.
- Any client GUI attempt to disable exponential backoff is blocked, and an attempt via CLI causes the NFM-P to raise an alarm.
- An error message is displayed if an NFM-P operator attempts to distribute an NE security policy or password policy to a 7705 SAR-Hm.

Note: The NFM-P does not initiate any configuration change to resolve an alarm raised because of a constraint violation. The alarm condition must be resolved by an NFM-P operator, or via the NE CLI, depending on the nature of the violation.

47.1.7 Management of remote 7705 SAR-Hm NEs

The 7705 SAR-Hm is a small form factor wireless router that extends IP/MPLS services over secure 3G/LTE wireless networks using cellular wireless infrastructure and WLAN technology. The 7705 SAR-Hm is available in several variants that have different cellular-interface radio capabilities. The cellular interface is the primary network port for WAN connectivity.

A 7705 SAR-Hm can be deployed in a remote location to perform wireless aggregation of traffic that is forwarded as IP packets to the cellular domain head-end node. In such a deployment, the cellular domain head-end node routes the traffic through a dedicated VPRN that you can optionally secure using NGE. See the *7705 SAR-Hm Main Configuration Guide* for additional functional, operational, and deployment information.

7705 SAR-Hm discovery, configuration, and management

You can use the NFM-P to perform the following discovery, configuration, and management functions for 7705 SAR-Hm devices.

Note: NFM-P management of remotely deployed 7705 SAR-Hm devices is limited to IPv4 only.

- Initiate ADP for the discovery of each 7705 SAR-Hm in a cellular domain. For the static cellular interface mode of operation, the NFM-P creates a management VPRN service for in-band 7705 SAR-Hm management.
- Upgrade the radio card firmware on 7705 SAR-Hm, Release 15.0 R6 and later, automatically during ADP.
- Create and manage 7705 SAR-Hm devices in cellular domains. Each 7705 SAR-Hm in a domain connects to the same head-end nodes and is part of the same NGE domain.
- Add the 7705 SAR-Hm devices that are going to be discovered to a cellular domain by importing an XML file that lists the SIM IMSI, and optionally, the system IP, of each device to discover in the domain.
- Move 7705 SAR-Hm NEs in or out of a cellular domain, or from one cellular domain to another.
- Globally apply/deploy a new security PIN to all 7705 SAR-Hm devices in a cellular domain during the ADP, or from domains that contain discovered devices, to overwrite any pre-existing/default PINs applied to the device. When a PIN is configured during the domain creation, all subsequent devices added to the domain using ADP have the same PIN applied to them.
- Create a security association between the SIM, IMEI, and the chassis identifier of each managed 7705 SAR-Hm. The NFM-P interprets a subsequent unexpected identifier change as a potential security violation, and alerts an operator.
- Configure polling in a cellular domain to monitor 7705 SAR-Hm reachability and system uptime. The polling interval is configurable in order to minimize traffic between the NFM-P and a large-scale 7705 SAR-Hm deployment.

47.1.8 In-band management using VPRN

When the cellular interface on a 7705 SAR-Hm is operating in static or dynamic cellular interface mode, the NFM-P can reach the NE system IP address through an in-band management VPRN service. For this mode of operation, the system IP address for NE management is private and differs from the cellular interface IP address. The system IP address must be advertised from the 7705 SAR-Hm to the head-end node by the in-band management VPRN service.

Routing in the private IP/MPLS network past the head-end node must allow management traffic to reach the head-end node, which then sends the management traffic over the VPRN to the 7705 SAR-Hm. Operators are responsible for configuring and ensuring connectivity to the NSP past the head-end node. This configuration is not described by this guide.

Each head-end node in a cellular domain must belong to the same VPRN service, which requires the following configuration:

- Route Distinguisher
- Route Target
- SNMP Community
- Auto-Bind Tunnel with GRE, and the resolution status set to Any Filter
- Autonomous System

Note: You can associate one VPRN service with only one cellular domain. If multiple head-end nodes are present in the domain, all head-end nodes must have the same VPRN service ID.

Figure 47-1, “Cellular domain management” (p. 1425) shows the scope of cellular domain management.

Figure 47-1 Cellular domain management

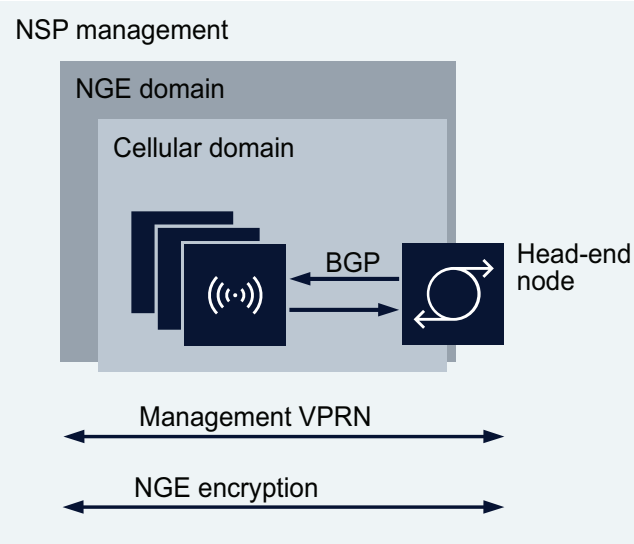

28135

47.1.9 Configuring polling for devices in cellular domains

During a system uptime poll of the NEs in a cellular domain, the NFM-P verifies the SIM information, IMEI, and chassis ID against the stored values. If a discrepancy is found, the NFM-P suspends management of the NE and disables resynchronization for the NE.

System uptime polling is performed in the following scenarios:

- during the initial ADP process
- if the dynamic BGP state of a head-end node changes

- at the operator-specified polling interval

As part of the cellular domain creation, a BGP group is configured on each head-end node. To monitor NE reachability in a cellular domain, the NFM-P polls the status of each BGP session between the head-end nodes and the managed NEs in the domain. Such a reachability check limits the traffic between the NFM-P and the managed NEs.

47.1.10 Device discovery and deployment using ADP

The NFM-P uses ADP, which is called ADP-Hm in the device documentation, to discover the remote devices in a cellular domain subnet. ADP provides all initialization and commissioning functions automatically for a newly installed device. After one or more SIMs is installed on a device and the device is turned up, ADP configures the cellular interface, establishes connectivity to the NFM-P, and waits for the NFM-P to complete the discovery and configuration of the device.

ADP automatically creates an NFM-P discovery rule to track the managed state of each NE in a cellular domain, and to initiate ADP when new devices in a cellular domain subnet are available for discovery. The NFM-P scans the network periodically for new devices, as specified by the discovery rule scan interval, which is the time between scans. To reduce the amount of network-management traffic in a cellular domain, you can configure the scan interval in the discovery rule for the subnet to be greater than the global scan interval defined in the NFM-P mediation configuration. The scan interval in a discovery rule overrides the scan interval in the NFM-P mediation configuration.

7705 SAR-Hm discovery prerequisites

The prerequisites for NFM-P discovery of a 7705 SAR-Hm using ADP are the following.

- The NFM-P manages one or more VSR, 7750 SR, or 7705 SAR head-end nodes that are to be included in the cellular domain.
- One or more valid SIM cards are inserted into the 7705 SAR-Hm.
- If dual SIM layout will be used:
  - the SIMs must be from two different carriers, with different HPLMNs
  - a PDN must be configured for each SIM

    The system assigns the PDN profile of the carrier to the port if the carrier IMSI prefix matches the port IMSI.

    For example, if carrier 2 has IMSI prefix 2121 configured and Port 1/1/1 IMSI is 2121xxxxxxxx the PDN profile configured in carrier 2 will be assigned to Port 1/1/1.
  - firmware for both carriers must be available in the NFM-P
- You have determined which ADP discovery method to use; the following are available:
  - one-step — all device configuration is performed automatically at the remote site after the device is turned up
  - two-step — NSP security configuration is performed at the customer staging facility; network configuration is performed automatically at the remote site after the device is turned up
- A route exists from the VPRN service to the cellular domain head-end node that is reachable by the NFM-P.
- A route exists from the cellular domain head-end node to each device that is to be discovered. For initial installation of a 7705 SAR-Hm cellular domain, IP addresses are typically allocated from a /24 or /18 address range.

- A default APN or Virtual Private Network (VPN) service has been procured from the service provider by the operator for the SIMs that are installed in the 7705 SAR-Hm. For a private cellular network, the operator can choose to use an APN if needed.

  If a static IP address is required for the IMSI associated with a SIM, the address can be allocated using one of the following methods:
  - by direct Home Subscriber Server (HSS) allocation, such as when a mobile carrier assigns an IP address
  - by deferred IP allocation; when the 7705 SAR-Hm first connects and authenticates with an HSS, the default APN associated with the service indicates that the IP allocation is deferred to an enterprise RADIUS AAA DHCP server. After the PGW learns the static IP address from the server, the PGW sends the address to the 7705 SAR-Hm in the PDP address IE after the default bearer is established.
- The PGW to which the 7705 SAR-Hm connects is configured with additional Protocol Configuration Options (PCOs) for the APN.

  The PCO must include the following two values:
  - dns-server-ipv4 primary – for example, config/mobile/pdn/apn/pco/dnsserveripv4/primary
  - dns-server-ipv4 secondary – for example, config/mobile/pdn/apn/pco/dnsserveripv4/backup
- Primary and secondary DNS servers are configured to resolve the NFM-P primary and standby main server IP addresses.
- An NFM-P cellular domain is configured with the required ADP operational settings and subnets for 7705 SAR-Hm discovery; see 47.1.1 “Cellular domains” (p. 1419) for information.

Offline NE handling during the ADP process

The ADP process cannot be completed when the 7705 SAR-Hm or the domain head-end node has an SNMP timeout or is not reachable.

When ADP starts, the online status of all the configured head-end nodes is checked. The ADP cannot start if none of the head-end nodes is online. ADP remembers the list of online head-end nodes and this list is used for the entire ADP process, rather than the list of configured head-end nodes. For ADP to succeed, at least one head-end node must be up during the entire ADP process. If all head-end nodes on the online head-end nodes list go down during the ADP process, ADP will fail.

When ADP with NGE configured is enabled, inbound ACL entries are created on the head-end nodes. The list of online head-end nodes is only updated when ADP is started and when outbound entries are added or removed. If any head-end node goes offline when ADP is in an initiating state, for example, adding inbound ACL entries to head-end nodes, the list of online head-end nodes may not be updated. This may cause ADP to fail or not start.

An OfflineDuringAdp alarm will be raised if a head-end node becomes offline during ADP. When a head-end node is marked as offline, its status will not be changed during the ADP process. If any offline head-end node becomes online during ADP, this head-end node will have no affect on the ongoing ADP process.

When the head-end node comes back online, you must manually fix the mismatched configuration and clear the alarm.

If any 7705 SAR-Hm becomes offline during ADP, ADP will fail for that 7705 SAR-Hm only. ADP will continue for all other online 7705 SAR-Hm NEs.

47.1.11 ADP discovery process

The following are the ADP operational phases:

- Network Discovery
- NFM-P Discovery
- 7705 SAR-Hm Discovery
- NFM-P Configuration

Phase 1 — Network Discovery

When a 7705 SAR-Hm initially boots, it runs the application load, executes the configuration file, which is empty, then checks the BOF to determine if ADP is enabled and needs to run.

If ADP is enabled on the NE, the NE performs the following:

- initializes the cellular interface using SIM1 for connectivity
- configures a PDN context with the default PDN profile
- after the cellular interface connects to the network, configures a PDN router interface to operate in dynamic cellular interface mode
- creates a loopback interface with a default name for the PDN interface, for example, “pdn1-loopback”; no IP address is assigned to the interface because it operates in dynamic cellular interface mode
- uses the loopback interface as the unnumbered interface for the PDN router interface

If the LTE network authenticates and accepts the new NE, a default bearer is established and the following are provided to the NE for the default APN to which the NE connects:

- cellular interface IP address
- DNS server IP addresses

Phase 2 — NFM-P Discovery

During the NFM-P discovery phase, the 7705 SAR-Hm sends DNS query messages to the DNS server addresses discovered in the previous phase.

The following NFM-P URLs are set in the BOF by default for the auto-discover function:

- auto-discover private.nokia.nsp.primary.nms
- auto-discover private.nokia.nsp.secondary.nms

Note: The names can also be set to the following:

- a different URL, if required
- an IP address, which eliminates the requirement for a DNS server

The 7705 SAR-Hm regularly sends a DNS query message until a DNS query response message that contains an NFM-P main server IP address is received. If no DNS query response message is received, ADP times out and reboots the device, after which ADP restarts the network discovery process.

Phase 3 — 7705 SAR-Hm Discovery

After the 7705 SAR-Hm receives one or more NFM-P server IP address, the 7705 SAR-Hm configures SNMPv2 trap destinations to the NFM-P server addresses using log ID 1.

ADP enables NetConf over SSHv2 and searches the user database for a user with access to NetConf. If none is found, NetConf access is given to the admin user.

The 7705 SAR-Hm initiates an SNMP trap poll that sends a notification to the NFM-P every 15 seconds for 30 minutes. If the ADP is not completed within the 30 minute interval, ADP will time out and begin again.

The 7705 SAR-Hm then sends an SNMPv2 Hello request, after which the NFM-P completes the device configuration, as described in the next phase.

Phase 4— NFM-P Configuration

In this phase, the NFM-P secures the 7705 SAR-Hm and completes the device configuration. During the configuration process, the 7705 SAR-Hm regularly sends an SNMPv3 trap to the NFM-P. When the configuration is complete, the NFM-P disables ADP on the NE.

47.1.12 ADP discovery methods

To meet differing security requirements, the following ADP discovery methods are available:

- one-step — all device configuration is performed automatically at the remote site after the device is turned up; the ADP process is active at the remote site from start to finish
- two-step:
  - NFM-P configures critical network and security parameters on device at the customer staging facility
  - remaining NE configuration is performed automatically after the device is installed and powered on at the remote site

During the ADP process with Dynamic Cellular IP mode, the NE can become unreachable at any time because the IP address that was used during the auto-discover process may change. The NE is at risk until the default in-band managed service is enabled and its configuration saved on the NE. Until then, the NFM-P relies on the IMSI value as the identifier for a particular NE. If the NE reboots during ADP and comes back, the SNMP trap hello message will indicate the IMSI and the cellular interface IP the NFM-P should be using to reach the NE and complete ADP.

Note: The NFM-P performs device configuration saves frequently during the process, regardless of the method used.

When the actions associated with either method are complete, the NFM-P does the following:

- stops the ADP process on the NE by executing an “ADP complete” command on the NE
- disables ADP in the NE BOF and clears the DNS entries on the NE so that a new discovery process cannot occur

One-step ADP discovery

In one-step ADP discovery, the 7705 SAR-Hm is turned up and ADP on the device completes the entire discovery and configuration process.

After the NFM-P receives an SNMPv2 trap and verifies the IMSI, and optionally, the system IP, the NFM-P uses NetConf over SSHv2 to configure the SNMPv3 user and parameters, including the required encryption and authentication keys. The configuration is based on the NFM-P mediation security policy associated with the cellular domain.

The NFM-P then completes the remainder of the device configuration:

- creates a strict security association between the 7705 SAR-Hm chassis information, IMEI, and SIM; the SIM cannot be inserted into another NE and managed by the NFM-P without operator intervention
- configures user names, passwords, scopes of control, and associated profiles
- downloads the required 7705 SAR-Hm software load, and resets the NE to apply the new load
- If NGE is required, configures NGE on the cellular interface by binding it into the NGE domain. If NGE is not required, this step is skipped. If the NGE domain is not encrypting, then the cellular interface is not enabled for encryption.
- if the cellular domain operation mode is Static Cellular Interface Mode or Dynamic Cellular Interface Mode, performs the following on each head-end node in the cellular domain to establish an in-band management service:
  - configures a BGP session to each head-end node in the cellular domain
  - configures an in-band management VPRN service used by the NFM-P to manage the 7705 SAR-Hm in-band over the GRE tunnels in the cellular network; the VPRN service can optionally be NGE-encrypted for secure NE management
  - if the 7705 SAR-Hm is running a release prior to 19.5, the NFM-P configures the GRT with static routes towards the VPRN over the PXC port to reach the primary and secondary NSP servers. The 7705 SAR- Hm system IP needs to be imported into BGP by configuring a static route in the VPRN for the system IP. This will allow the system IP to be reachable via the VPRN service from the head-end nodes. The VPRN static route for the system IP will point towards the PXC that will reach the GRT. At this point all NFM-P traffic can be routed over the VPRN service once the NFM-P starts using the new system IP address to manage the NE instead of the cellular interface IP address.
  - if the 7705 SAR-Hm is running release 19.5 or later, the NFM-P configures a GRT leak in the VPRN to leak the 7705 SAR-Hm system IP to the base routing of the 7705 SAR-Hm. This will allow the system IP to be reachable via the VPRN service from the head-end nodes. At this point all NFM-P traffic can be routed over the VPRN service once the NFM-P starts using the new system IP address to manage the NE instead of the cellular interface IP address.

After the NFM-P completes the ADP process, the 7705 SAR-Hm Status and Alarm LEDs indicate that the ADP process is complete. The NE is securely managed by the NFM-P and ready for service.

Two-step ADP discovery

In two-step ADP discovery, the 7705 SAR-Hm is powered on first in a staging area for the initial NFM-P security configuration, then a second time at the remote site to complete the remaining configuration tasks, as described in the following sequence:

1. The 7705 SAR-Hm is powered on for the first time and the NFM-P does the following:

   Note: ADP for the subnet must be enabled on the NFM-P during this step.

- creates a strict security association between the 7705 SAR-Hm chassis information, IMEI, and SIM; the SIM cannot be inserted into another NE and managed by the NFM-P without operator intervention
- configures user names, passwords, scopes of command, and associated profiles
- downloads the required 7705 SAR-Hm software load, and resets the NE to apply the new load
- stops the ADP process on the NE by executing an “ADP complete” command on the NE

The 7705 SAR-Hm Status LED turns solid green and the Alarm LED continues to blink. The 7705 SAR-Hm has completed step one and can be powered off and shipped to the remote site for installation.

2. After the 7705 SAR-Hm is installed and powered on at the remote site, the following occur:

   Note: ADP for the subnet must be enabled on the NFM-P during this step.

- The 7705 SAR-Hm regularly sends SNMPv3 traps to the NFM-P to indicate that the ADP process can resume.
- The NFM-P downloads the NGE key group of the NGE domain associated with the cellular domain, if NGE is to be used, and configures the key group on the 7705 SAR-Hm cellular interface.
- If the cellular domain mode is Static Cellular Interface Mode or Dynamic Cellular Interface Mode, the NFM-P performs the following on each head-end node in the cellular domain to establish an in-band management service:
  - configures a BGP session to each head-end node in the cellular domain
  - configures an in-band management VPRN service used by the NFM-P to manage the 7705 SAR-Hm in-band over the GRE-IMPLS tunnels in the cellular network; the VPRN service can optionally be NGE-encrypted for secure NE management

After the NFM-P completes the ADP process, the 7705 SAR-Hm Status and Alarm LEDs indicate that the ADP process is complete. The NE is securely managed by the NFM-P and ready for service.

47.2 Workflow to manage cellular domain devices

47.2.1 Stages

1

Review and complete the required prerequisites to allow the NFM-P to discover the required devices using ADP; see “7705 SAR-Hm discovery prerequisites” (p. 1427).

2

For the NEs in a cellular domain that are assigned a unique administrator password, create an XML file that maps each chassis serial number to each unique password, as described in 47.3 “To create an ADP password mapping file” (p. 1435).

3

If you want to automatically apply a specific configuration to discovered NEs in a cellular domain, create a post-discovery action. See 9.5 “Post-discovery actions on discovered NEs” (p. 282) in Chapter 9, “Device discovery” for information.

4

Create an SNMPv3 NE user account; see the procedure to configure a user account on a manage device in the *NSP System Administrator Guide*.

Note: The user requires console, FTP, SNMP, and NetConf access.

Note: If enhanced NE security mode is enabled, or you intend to enable the function, the user password must adhere to the password criteria specified in 47.4 “To enable enhanced NE security mode” (p. 1437).

Note: You must not specify AES-256 as the Privacy Protocol in the user SNMPv3 configuration.

5

Create a software upgrade policy in which the Software Download and Activate Image parameters are enabled, and the Software Upgrade parameter is disabled; see 26.5 “To configure a software upgrade policy” (p. 776).

6

Create a mediation security policy that specifies the following; see 9.17 “To configure device mediation” (p. 301) for information:

- SNMPv3 user: the new user created in Stage 4.
- CLI user: must be the same as the SNMPv3 user
- SSH for CLI access
- Secure FTP

Note: You must enable NetConf in the mediation policy, and specify the SNMPv3 user in the NetConf section of the mediation policy. The NetConf credentials are available from technical support.

7

If required, enable enhanced NE security mode; see

8

Create a carrier VPRN with static routes to and from the 7705 SAR-Hm; see 79.5 “To create a VPRN service” (p. 2534) and 79.57 “To configure static routes on a VPRN site” (p. 2621); ensure that the service is configured as described in 47.1.8 “In-band management using VPRN” (p. 1425).

9

Configure a PDN profile policy that specifies each cellular interface used to connect 7705 SAR-Hm mobile data users to the network; see 69.2 “To configure a PDN Profile policy for a 7705 SAR-Hm” (p. 1901).

10

If NGE security is required, create an NGE domain that includes the cellular domain head-end nodes as the NGE domain gateway sites; see 82.10 “Workflow for NGE management using NFM-P” (p. 2727).

11

If required, upgrade 7705 SAR-Hm radio card firmware; see 26.12 “To upgrade 7705 SAR-Hm radio card firmware” (p. 790).

12

Create a cellular domain that has the cellular domain head-end node as the head end, and includes each 7705 SAR-Hm device; see 47.5 “To configure a cellular domain with single SIM deployment” (p. 1440) or 47.6 “To configure a cellular domain with dual SIM deployment” (p. 1444).

13

If you intend to discover the devices listed in an XML file, import the file to the cellular domain that is to include the devices, as described in 47.5 “To configure a cellular domain with single SIM deployment” (p. 1440) or 47.6 “To configure a cellular domain with dual SIM deployment” (p. 1444); see “Specifying devices using an XML file” (p. 1422) for information about the XML file format.

14

If the head-end node is a 7705 SAR, create BGP peers in the domain BGP group; see 28.33 “To configure peer-level BGP” (p. 927).

15

Enable ADP on one or more selected cellular domain subnets, and monitor the discovery process; see 47.7 “To enable and monitor the ADP discovery process” (p. 1448).

16

If you are using the two-step ADP method, perform the following steps after the first step ADP is completed:

1. Check the LED status of each device to be discovered, as described in the device documentation, to ensure that the first ADP step is complete.
2. Disable ADP on the cellular domain subnet that includes the devices; see 47.11 “To disable ADP” (p. 1454).
3. Deliver the devices to the remote installation site.
4. Enable ADP on the cellular domain subnet that contains the devices.

5. Install and turn up each 7705 SAR-Hm that is to complete the ADP process at the remote site.
6. Monitor the ADP process using the NFM-P, as described in 47.7 “To enable and monitor the ADP discovery process” (p. 1448), and by checking the device LEDs, as described in the device documentation.

17

Disable ADP on the cellular domain subnets that no longer require ADP for device discovery; see 47.11 “To disable ADP” (p. 1454).

18

Apply a new security PIN to all NEs in a cellular domain, as required; see 47.10 “To configure a new PIN value for a cellular carrier” (p. 1452).

19

If dual SIM deployment is in use, configure the Cellular tab of the MDA on each 7705 SAR-Hm; see 15.78 “To configure an MDA” (p. 536).

20

Configure polling for the NEs in the cellular domain; see 8.12 “To configure polling for a 7705 SAR-Hm” (p. 264).

47.3 To create an ADP password mapping file

47.3.1 Purpose

Perform this procedure to create a file that facilitates the entry of 7705 SAR-Hm serial number and password information for the devices in a cellular domain.

The XML file that you create is imported during cellular domain configuration. The file maps the chassis serial number of each device in the cellular domain to the unique administrator password of the device.

47.3.2 Steps

1

Use a plain-text editor to create a file that contains the following XML header:

```
<?xml version="1.0" encoding="UTF-8"?>
<ADPSerialNumbers>
```

2

For each device that has a unique password associated with the chassis serial number, add a line below the header in the following format:

```
<node serialNumber="serial_number" password="password"/>
```

where

serial_number is the chassis serial number

password is the unique administrator password

3

When you are finished adding entries, add the following XML footer to the end of the file:

```
</ADPSerialNumbers>
```

The file reads as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<ADPSerialNumbers>
<node serialNumber="SN_1" password="password_1"/>
<node serialNumber="SN_2" password="password_2"/>
.
.
.
<node serialNumber="SN_n" password="password_n"/>
</ADPSerialNumbers>
```

where

SN_1 to *SN_n* are the chassis serial numbers

password_1 to *password_n* are the administrator passwords

4

Save and close the file using a .xml filename extension.

5

Transfer the XML file to the client GUI station that is to be used for the cellular domain configuration.

END OF STEPS

47.4 To enable enhanced NE security mode

47.4.1 Purpose

CAUTION

Irreversible Configuration Change

Enabling enhanced NE security mode affects each 7705 SAR-Hm in the managed network; disabling the function requires that you unmanage each managed NE of any type in the entire network.

If you are not absolutely certain that you want to enable enhanced NE security mode, do not perform the procedure.

CAUTION

Service Disruption

Modifying the server configuration can have serious consequences including service disruption.

Contact technical support before you attempt to modify the server configuration.

Perform this procedure to enable the stringent security measures of NFM-P enhanced NE security mode.

47.4.2 Steps

Configure SNMPv3 user password

1

For enhanced NE security mode, the password of the SNMPv3 NE user must conform to specific criteria.

If the user password does not currently conform to the following criteria, change the password to meet the criteria:

- 10 or more characters in length
- does not include the username
- includes 3 or fewer consecutive instances of the same character
- includes 1 or more lower-case characters
- includes 1 or more upper-case characters
- includes 1 or more numeric characters

1. Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens.
2. Select the ADP SNMPv3 user and click Properties. The NE User (Edit) form opens.
3. Configure the Password and Confirm Password parameters using a value that meets the required criteria.

4. Click OK to save your changes and close the form.

2

The password for SSH and FTP access in the ADP mediation security policy must match the password of the SNMPv3 NE user.

If the passwords do not match, update the password in the mediation security policy.

1. Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.
2. Select the ADP mediation policy and click Properties. The Mediation Policy (Edit) form opens.
3. Configure the User Password and Confirm Password parameters in the CLI panel using the SNMPv3 NE user password.
4. Configure the User Password and Confirm Password parameters in the FTP panel using the SNMPv3 NE user password.
5. Click OK to save your changes and close the form.

Configure NE password policy

3

Choose Administration→Security→NE Password Policy from the NFM-P main menu. The NE Password Policy form opens.

4

Configure the following parameters using the values shown:

- Properties panel:
  - Lockout Time (minutes)—10
  - Minimum Length—10
  - Maximum Attempts—10 or fewer
  - Maximum Attempts Time (minutes)—10 or less
- Complexity Rules panel:
  - User Name Allowed in Password—False
  - Maximum Number Of Times Character Can Used Consecutively—3 or fewer
  - Minimum Number Of LowerCase Characters—1 or more
  - Minimum Number Of UpperCase Characters—1 or more
  - Minimum Number Of Numeric—1 or more

5

Click OK to save your changes and close the form.

Configure NE security policy

6

Choose Administration→Security→NE System Security from the NFM-P main menu. The NE System Security form opens.

7

Select one or more NEs and click Properties. A properties form opens.

8

Select the Exponential BackOff parameter.

9

Click on the Servers Configuration tab.

10

In the SSH Configuration panel, configure the SSH Version parameter by selecting Version 2 and deselecting the following:

- Version 1
- Version 1-2

11

Click OK to save your changes and close the form.

Configure NFM-P main servers

12

Perform the following steps on each main server station.

1. Log in as the nsp user on the main server station.
2. Open a console window.
3. Navigate to the /opt/nsp/nfmp/server/nms/config directory.
4. Create a backup copy of the nms-server.xml file.
5. Open the nms-server.xml file using a plain-text editor such as vi.
6. Locate the section that begins with following XML tag:

   ```
   <policyConfig
   ```

7. Insert the following line before the last line of the section, which ends with the /> tag:

   ```
   enhancedNESecurityMode="true"
   ```

8. Save and close the file.

13

On the standalone main server, or the primary main server in a redundant system, enter the following:

```
bash$ /opt/nsp/nfmp/server/nms/bin/nmsserver.bash read_config ↵
```

The NFM-P puts the configuration change into effect.

14

Close the open console windows.

END OF STEPS

47.5 To configure a cellular domain with single SIM deployment

47.5.1 Steps

1

Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Click Create. The Cellular Domain [Create | Edit] form opens.

3

Configure the Name parameter.

4

Configure the Sim Layout parameter to Single SIM.

For dual SIM, see 47.6 “To configure a cellular domain with dual SIM deployment” (p. 1444).

5

Configure the Operation Mode parameter. See 47.1.2 “Cellular domain operation modes” (p. 1420) for information about operation modes.

6

Configure the ADP Process parameter to specify one-step or two-step discovery.

7

Configure the ADP System IP Address parameter.

8

If the cellular domain includes any older NEs that are factory-shipped with a default administrator password, configure the Default ADP-Hm admin Password parameter using the administrator password that is common to the NEs.

9

Specify one or more head-end nodes for the cellular domain, as required.

1. Click Create in the Head-End Nodes panel. The Head-End Node [Create] form opens.
2. Click Select and choose a managed NE.
3. Configure the parameters.
4. Click OK. The Head-End Node [Create] form closes, and the NE is listed in the Head-End Nodes List.

10

Click Select in the Mediation Security panel and choose the required SNMPv3 mediation security policy.

11

Click Select in the PDN Profile panel and choose the required PDN Profile policy.

12

Click Select in the Software panel and choose the required software and firmware upgrade policy.

During software upgrade, if the 7705 SAR-Hm has a later image than that mentioned in the domain, the node software is not changed with the version mentioned in the domain since it is already running a latest image.

Unlike software upgrade, there is no version check of the current 7705 SAR-Hm firmware. The firmware image mentioned in the domain is pushed down to the node even if it is an older image.

Leave the firmware image empty if you want to skip the firmware upgrade.

13

If required, configure the Firmware Version parameter.

For the 7705 SAR-Hm, use a format that is aligned with the Sierra Wireless web portal, [VENDOR NAME]<space>[FIRMWARE_SUFFIX]<space>[PRI]. For example: GENERIC 02.24.03.00 002.026_000

For the 7705 SAR-Hmc, use a format such as this example: CAT4_GS_BYPASS_0.3.3.9_V1.9

14

Configure the Software Version parameter using the string that identifies the software release to which the devices are to be upgraded; for example, TiMOS-B-*RR*.*r* R*n*.

15

If NGE is to be used, click Select in the NGE Domain panel and choose the required NGE domain.

16

Configure the parameters in the PIN Security panel.

17

Configure the parameters in the BGP Configuration panel.

18

If required, use the Select button in the Post Discovery Action panel to choose a post-discovery action for NEs in the cellular domain.

Note: The post-discovery action is performed on the NEs only when the ADP process for each NE in the cellular domain is complete.

19

If the Operation Mode parameter is set to Static Cellular Interface or Dynamic Cellular Interface, perform the following steps.

1. Click on the VPRN tab.
2. Use the Select button to specify the VPRN service to associate with the cellular domain.

20

Click Apply.

21

Click OK. The Cellular Domain [Edit] form refreshes and displays additional tabs.

22

If the Operation Mode is set to Static Cellular System, add new devices to the cellular domain, as required.

1. Click on the Sites tab.
2. Click Add→Add New Sites. The Domain Site [Create] form opens.
3. Configure the parameters.

   You can configure the System Name parameter. When the 7705 SAR-Hm is discovered, the system name appears as the shelf name in the NFM-P GUI. However, all main references such as alarms are based on the site name.
4. Click OK. The Domain Site [Create] form closes.

23

Add subnets to the domain, as required.

1. Click on the Subnets tab.

2. Configure the parameters.
3. Click OK.

24

If any NEs in the cellular domain have unique administrator passwords, click on the Serial Number Data tab. Otherwise, go to Step 26.

25

Perform one of the following.

a. Import the passwords from a file created as described in 47.3 “To create an ADP password mapping file” (p. 1435).
   1. Click Import Serial Numbers. A file browser form opens.
   2. Use the form to choose the password file and click Open. The passwords are imported, and the NE serial numbers are listed on the form.

b. Manually add one or more passwords; perform the following steps for each NE.

1. Click Create. The Serial Number Data (Create) form opens.
2. Configure the Serial Number and Password parameters using the NE values.
3. Save the changes and close the form.

26

If any head-end node in the domain is a 7705 SAR, configure BGP peers on the head-end node BGP group; see 28.33 “To configure peer-level BGP” (p. 927).

27

If the Operation Mode is set to Static Cellular Interface or Dynamic Cellular Interface, or the Operation Mode is set to Static Cellular System and the ADP System IP Address parameter is set to Use XML, import a list of devices.

1. Create an XML file that lists each 7705 SAR-Hm to be discovered. See “Specifying devices using an XML file” (p. 1422) for information about the file format.
2. Click on the Sites tab.
3. Click on the Import Sites button. The Open File browser form opens.
4. Navigate to the XML file that specifies the devices to discover.
5. Select the file and click Open. The devices specified in the file are added to the Domain Site list.

28

If the Operation Mode is set to Static Cellular Interface or Dynamic Cellular Interface, and the ADP System IP Address parameter is set to Verify and Allocate from Pool, define one or more System IP address pools to provide a sufficient number of IP addresses for all devices that are to be discovered in the cellular domain subnets.

Note: The NFM-P does not reuse any IP address in a pool. Ensure that the System Prefix value that you specify allocates a sufficient range of addresses for the subnet.

1. Click on the System IP Pool tab.
2. Click Create. The System IP Pool [Create] form opens.
3. Configure the System Address and System Prefix parameters.

29

Click OK to save your changes and close the open forms, as required.

END OF STEPS

47.6 To configure a cellular domain with dual SIM deployment

47.6.1 Steps

1

Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Click Create. The Cellular Domain [Create | Edit] form opens.

3

Configure the Name parameter.

4

Configure the Sim Layout parameter to Dual SIM.

For single SIM, see 47.5 “To configure a cellular domain with single SIM deployment” (p. 1440).

5

Configure the ADP Process parameter to specify one-step or two-step discovery.

6

Configure the ADP System IP Address parameter.

7

Configure the carrier details. Perform the following for each carrier:

1. Click Create in the Carriers panel. The Cellular Carrier (Create) form opens.
2. Configure the Name parameter.
3. Click Select in the PDN Profile panel and choose the required PDN Profile policy.
4. Configure the Firmware Version parameter.

For the 7705 SAR-Hm, use a format that is aligned with the Sierra Wireless web portal, [VENDOR NAME]<space>[FIRMWARE_SUFFIX]<space>[PRI]. For example: GENERIC 02.24.03.00 002.026_000

For the 7705 SAR-Hmc, use a format such as this example: CAT4_GS_BYPASS_0.3.3.9_V1.9

5. Configure the parameters in the PIN Security panel.
6. Configure other parameters in the General tab as required.
7. In the IMSI Prefix tab, enter the IMSI prefixes that need to be matched for the carrier.
8. Click OK to close the form.
9. Repeat this step to configure the second carrier.

8

Specify the head-end nodes. An NE can serve as the head-end node for one carrier or both.

Perform the following for each head-end node:

1. Click Create. The Head-End Node [Create] form opens.
2. Click Select in the Head-End Node panel and choose a managed NE.
3. Configure the parameters.
4. Click OK. The Head-End Node [Create] form closes, and the NE is listed in the Head-End Nodes List.

9

Click Select in the Mediation Security panel and choose the required SNMPv3 mediation security policy.

10

Configure the Software Version parameter using the string that identifies the software release to which the devices are to be upgraded; for example, TiMOS-B-*RR*.*r* R*n*.

11

Click Select in the Software and Firmware Upgrade Policy panel and choose the required software and firmware upgrade policy.

Associate a non default SR Software Upgrade Policy. Ensure that Software Download and Activate Image is Enabled under SR Based Setting and Firmware Download and Firmware Activate is Enabled under Firmware Settings.

During software upgrade, if the 7705 SAR-Hm has a later image than that mentioned in the domain, the node software is not changed with the version mentioned in the domain since it is already running a latest image.

Unlike software upgrade, there is no version check of the current 7705 SAR-Hm firmware. The firmware image mentioned in the domain is pushed down to the node even if it is an older image.

Leave the firmware image empty if you want to skip the firmware upgrade.

12

Verify that the correct software and firmware images are imported.

13

If NGE is to be used, click Select in the NGE Domain panel and choose the required NGE domain.

14

Configure the parameters in the BGP Configuration panel.

15

If required, use the Select button in the Post Discovery Action panel to choose a post-discovery action for NEs in the cellular domain.

Note: The post-discovery action is performed on the NEs only when the ADP process for each NE in the cellular domain is complete.

16

In the VPRN tab, use the Select button to specify the VPRN service to associate with the cellular domain.

17

Click Apply.

18

Click OK. The Cellular Domain [Edit] form refreshes and displays additional tabs.

19

If the ADP System IP Address parameter is set to something other than Use XML, define one or more IP address pools to provide a sufficient number of IP addresses for all devices that are to be discovered in the cellular domain subnets.

Note: The NFM-P does not reuse any IP address in a pool. Ensure that the System Prefix value that you specify allocates a sufficient range of addresses for the subnet.

1. Click on the System IP Pool tab.
2. Click Create. The System IP Pool [Create] form opens.
3. Configure the System Address and System Prefix parameters.

20

Add subnets to the domain as required. Perform the following for each carrier:

1. Click on the Subnets tab.
2. Click Create and configure the parameters.
3. Click OK. The subnet is created with ADP off.

21

Add new devices to the cellular domain, as required.

1. Select the carrier to which the subnet is applicable.
2. Click on the Sites tab.
3. Click Add→Add New Sites. The Domain Site [Create] form opens.
4. Configure the parameters.

   You can configure the System Name parameter. When the 7705 SAR-Hm is discovered, the system name appears as the shelf name in the NFM-P GUI. However, all main references such as alarms are based on the site name.
5. Click OK. The Domain Site [Create] form closes.

22

If any NEs in the cellular domain have unique administrator passwords, click on the Serial Number Data tab. Otherwise, go to Step 24.

23

Perform one of the following.

a. Import the passwords from a file created as described in 47.3 “To create an ADP password mapping file” (p. 1435).
   1. Click Import Serial Numbers. A file browser form opens.
   2. Use the form to choose the password file and click Open. The passwords are imported, and the NE serial numbers are listed on the form.

b. Manually add one or more passwords; perform the following steps for each NE.

1. Click Create. The Serial Number Data (Create) form opens.
2. Configure the Serial Number and Password parameters using the NE values.
3. Save the changes and close the form.

24

If any head-end node in the domain is a 7705 SAR, configure BGP peers on the head-end node BGP group; see 28.33 “To configure peer-level BGP” (p. 927).

25

If the ADP System IP Address parameter is set to Use XML, import a list of devices.

1. Create an XML file that lists each 7705 SAR-Hm to be discovered. See “Specifying devices using an XML file” (p. 1422) for information about the file format.
2. Click on the Sites tab.
3. Click on the Import Sites button. The Open File browser form opens.
4. Navigate to the XML file that specifies the devices to discover.
5. Select the file and click Open. The devices specified in the file are added to the Domain Site list.

26

Click OK to save your changes and close the open forms, as required.

END OF STEPS

47.7 To enable and monitor the ADP discovery process

47.7.1 Before you begin

Ensure the NEs support a valid combination of authentication and privacy protocols. See the NE documentation or 9.11 “To enable SNMPv3 management of a device” (p. 291) to review the disallowed combinations.

47.7.2 Steps

1

Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Click Search and choose a cellular domain. The Cellular Domain [Edit] form opens.

3

Select the Enable ADP parameter.

4

Click Apply. The ADP process begins.

5

Monitor the ADP process for a site, if required.

1. Click on the Sites tab. Each 7705 SAR-Hm in the cellular domain is listed.
2. Select a site and click Properties. The Domain Site [Edit] form opens.
3. View the Operation and Status indicators in the Status panel. When the ADP process for the site is complete, the Operation indicator displays ADP Complete, and the Status indicator reads Success or Failure, depending on the outcome.

   If the process fails, see 47.8 “Workflow to recover from ADP discovery failure” (p. 1449).

   If the process is successful, the State indicator in the ADP panel reads Managed.
4. Close the Domain Site [Edit] form.

6

Close the open forms, as required.

END OF STEPS

47.8 Workflow to recover from ADP discovery failure

47.8.1 Process

If the NSP host server restarts or switches over to a redundant server during ADP, the ADP process will fail. In this scenario, NFM-P will show ADP Off. To recover from an ADP failure due to a restart or switchover, perform Stage 1 through Stage 3. For any other ADP failure, perform Stage 4 through Stage 6.

Scenario 1: ADP failure caused by an NSP host server restart or switchover

1

Perform the following in the NFM-P:

a. If an NGE domain is created in the cellular domain and the NGE domain is encrypted, create inbound and outbound ACL IP exception filters on the head-end nodes; see 51.7 “To configure an ACL IP exception filter policy” (p. 1683). Apply the new ACL IP exception filters to the gateway interfaces in the NGE domain.

- For an inbound IP exception filter, the Source IP parameter must be set to the subnet IP address of the cellular domain. The Destination IP must be set to the IP address of the active NSP host server.
- For an outbound IP exception filter, the Source IP parameter is set to the IP address of the active NSP host server. The Destination IP is set to the IP address of the 7705 SAR-Hm to be discovered.

b.

Delete the failed sites from the cellular domain:

1. Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.
2. Click Search and choose the cellular domain that initiated the failed ADP operation. The Cellular Domain [Edit] form opens.
3. Click on the Sites tab.
4. Choose a site from the list and click Delete.
5. Repeat for all sites in the domain for which ADP failed.

c. If the NE is already in a managed state, delete it from the NFM-P discovery manager; see 9.37 “To delete a device from the managed network” (p. 325).

d. Delete the rule element corresponding to the NE from discovery rules created by ADP for the domain; see 9.23 “To configure a discovery rule” (p. 310).

e. If the ADP System IP Address parameter is set to Verify, reload the XML file to the cellular domain.

2

Perform the following on the NE. See the NE documentation for more information.

a. Establish an SSHv2 session to the NE from either a terminal window or the NFM-P.

b. Restore the BOF file to factory default.

c. Delete the existing 7705 SAR-Hm config file.

d. Reboot the NE.

3

Perform the following in the NFM-P:

a.

If ACL IP exception filters were created in Stage 1 a

1. Remove the ACL IP exception filter from the gateway interfaces in the NGE domain.
2. Delete both inbound and outbound ACL IP exception filters; see 51.7 “To configure an ACL IP exception filter policy” (p. 1683).

b. Enable ADP; see 47.7 “To enable and monitor the ADP discovery process” (p. 1448).

Scenario 2: Other ADP failure

4

Perform the following in the NFM-P:

a. Disable ADP; see 47.11 “To disable ADP” (p. 1454).

b. If an NGE domain is created in the cellular domain and the NGE domain is encrypted, create inbound and outbound ACL IP exception filters on head-end nodes; see 51.7 “To configure an ACL IP exception filter policy” (p. 1683). Apply the new ACL IP exception filters to the gateway interfaces in the NGE domain.

- For an inbound IP exception filter, the Source IP parameter must be set to the subnet IP address of the cellular domain. The Destination IP must be set to the IP address of the active NSP host server.
- For an outbound IP exception filter, the Source IP parameter is set to the IP address of the active NSP host server. The Destination IP is set to the IP address of the 7705 SAR-Hm to be discovered.

c.

Delete the failed sites from the cellular domain:

1. Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.
2. Click Search and choose the cellular domain that initiated the failed ADP operation. The Cellular Domain [Edit] form opens.
3. Click on the Sites tab.
4. Choose a site from the list and click Delete.
5. Repeat for all sites in the domain for which ADP failed.

d. If the NE is already in a managed state, delete it from the NFM-P discovery manager; see 9.37 “To delete a device from the managed network” (p. 325).

e. Delete the rule element corresponding to the NE from discovery rules created by ADP for the domain; see 9.23 “To configure a discovery rule” (p. 310).

f. If the ADP System IP Address parameter is set to Verify, reload the XML file to the cellular domain.

5

Perform the following on the NE. See the NE documentation for more information.

a. Establish an SSHv2 session to the NE from either a terminal window or the NFM-P.

b. Restore the BOF file to factory default.

c. Delete the existing 7705 SAR-Hm config file.

d. Reboot the NE.

6

Perform the following in the NFM-P:

a.

If ACL IP exception filters were created in Stage 4b:

1. Remove the ACL IP exception filter from the gateway interfaces in the NGE domain.
2. Delete both inbound and outbound ACL IP exception filters; see 51.7 “To configure an ACL IP exception filter policy” (p. 1683).

b. Enable ADP; see 47.7 “To enable and monitor the ADP discovery process” (p. 1448).

47.9 To move a 7705 SAR-Hm NE in or out of a cellular domain

47.9.1 Purpose

Use this procedure to add or remove a 7705 SAR-Hm NE from a cellular domain. Adding an NE may be required when a managed NE is to be included in the domain, but has been deployed using a method other than ADP, or when an NE is to be moved from one cellular domain to another.

47.9.2 Steps

1

Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Select a cellular domain and click Properties. The Cellular Domain [Edit] form opens.

3

Click on the Sites tab.

4

To add an NE to the domain, click Add→Add Managed Sites. The Add Managed Sites-Select Sites form opens.

Click Search to list all available NEs, select an NE, and click OK.

Note: The following conditions are required to successfully add the 7705 SAR-Hm NE to the cellular domain:

- The NE must be in one of the domain subnets and not already in a cellular domain.
- If the domain is in Cellular Interface mode, the NE must be part of the VPRN service for in-band management.
- The NE must use the same mediation security policy, PDN profile policy, and software version that is used in cellular domain configuration.
- The NE must have the domain BGP Group configured on its routing instance, with the same AS number as the domain.
- The NE must have a BGP peer to the domain head-end node.
- If NGE is in use, the NE and its PDN interface must be in the NGE domain. The NE must appear in the Sites tab, and the PDN interface must appear in the L3 interfaces tab.
- If dual SIM is in use, the NE must have the same combination of cellular carriers as the other NEs in the domain.

5

To remove an NE from the domain, click Search to list all available NEs, select an NE, and click Delete.

6

Close the open forms, as required.

END OF STEPS

47.10 To configure a new PIN value for a cellular carrier

47.10.1 When to use

Each new 7750 SAR-Hm node that is discovered by the NFM-P may have a default PIN applied to each SIM for security and access rights.

The PIN is configured on a per carrier basis. If a single SIM layout is in use, a single PIN is configured for the domain and applied to all sites. In a dual SIM layout, a PIN is configured for each carrier and applied to all sites for the SIM that applies to that carrier.

You can globally apply and deploy a new security PIN during the ADP or from domains that contain discovered nodes. When a PIN is configured during the domain creation, all subsequent nodes added to the domain through the ADP will have the new PIN applied for the carrier.

Note: Before starting this procedure, you must first create the cellular domain and at least one 7705-SAR Hm node must be present in the domain; see 47.5 “To configure a cellular domain with single SIM deployment” (p. 1440) or 47.6 “To configure a cellular domain with dual SIM deployment” (p. 1444)..

47.10.2 Steps

1

Choose Manage→Cellular Domain from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Click Search and choose the cellular domain that you want to configure a new PIN value for. The Cellular Domain [Edit] form opens.

3

On the General tab, disable the Enable ADP parameter.

4

If Dual SIM is in use, choose a carrier and click Properties. The Cellular Carrier (Edit) form opens.

5

Enter a new value for the New PIN parameter.

6

Click Apply to save the form and deploy the new PIN entry to all 7705 SAR-Hm nodes in the domain.

7

On the General tab, click the Enable ADP parameter if you want all subsequent nodes added to the domain to have the same PIN applied to them.

8

Close the form.

END OF STEPS

47.11 To disable ADP

47.11.1 Purpose

Perform this procedure after the ADP process has completed, or as part of recovery from ADP failure.

Caution: Do not disable ADP while ADP is in progress for any NEs. This could cause ADP to enter an indefinite state, requiring the failure recovery process to be restarted.

47.11.2 Steps

1

Choose Manage→Cellular Domains from the NFM-P main menu. The Cellular Domain Manager form opens.

2

Click Search and choose a cellular domain. The Cellular Domain [Edit] form opens.

3

Deselect the Enable ADP parameter.

4

Click Apply. The ADP process is disabled for the cellular domain.

5

Close the open forms, as required.

END OF STEPS

48 MACsec

48.1 MACsec

48.1.1 Overview

Media Access Control Security (MACsec) provides point-to-point, point-to-multipoint, and in-band security on Ethernet links between directly connected NEs or NEs connected over a Layer 2 cloud. MACsec uses MACsec Key Agreement (MKA) to signal data path encryption keys known as Security Association Keys (SAKs).

To use MACsec, you need to configure a global and a local connectivity association. A connectivity association creates secure channels for inbound and outbound traffic.

MACsec in NFM-P supports the static Connectivity Association Key (CAK) feature, which uses MKA and a Pre-Shared Key (PSK) to discover and authenticate peers. Static CAK uses two security keys to secure the Ethernet link: a CAK that secures control plane traffic, and an SAK that secures data plane traffic. Both keys are regularly exchanged between devices on each end of the Ethernet link to ensure link security. PSK is also used for encrypting SAKs between the key server and other peers in the MACsec connectivity association.

48.1.2 MACsec configuration components

The following need to be configured to use MACsec:

- Connectivity association: a global connectivity association configuration. The global configuration contains the static CAK.
- Connectivity association site: a local connectivity association configuration. The site contains parameters that must be synchronized to the global configuration, and parameters that can be changed on a local basis.
- Static CAK: the CAK contains the PSK and related information. A maximum of two PSKs can be configured at one time, and only one can be active at a time.
- PSKs: the PSKs contain the connectivity association name and CAK.
  - The global PSK is created under the global static CAK object.
  - The local PSK is created as part of the global PSK creation if the connectivity association site exists under the global connectivity association, or as part of distributing PSKs to sites.

  The NFM-P manages PSK generation and rekeying operations. PSK distribution uses SSH2.

48.1.3 MACsec interfaces

A MACsec port is an Ethernet port on a card that supports MACsec. A MACsec interface is a virtual sub-port, that is, a VLAN, on a MACsec port. MACsec port properties apply to all MACsec interfaces on the port. MACsec interface properties apply only to the interface.

MACsec is activated when MACsec interfaces are added to a connectivity association, and the interfaces, the ports, and the connectivity association are all in an administrative up state. If the

port, interface, or connectivity association is down, the MKA session providing encryption will end. A new session will be established when the port, interface, and connectivity association are all back up.

48.1.4 Key updates with offline nodes

A node is in offline mode if the NFM-P has lost connectivity to the node at the time of the action (global PSK creation, manual PSK distribution, scheduled rekey). A node is not in offline mode if it responds to SSH and SNMP requests.

By default, the key update process skips offline nodes. When the node is back online, you can manually distribute the new keys to the node or create a global key to be automatically distributed to all sites in the CA.

48.2 MACsec statistics

48.2.1 MACsec statistics overview

You can collect MKA session statistics for a port or subport on demand, and schedule statistics collection using a MIB entry policy. The statistics are displayed on the Statistics tab of the MACsec interface. See Chapter 9, “Device discovery” for information about configuring MIB entry policies.

MACsec statistics can be viewed and collected from the Statistics tab of the following object properties forms:

- MACsec interface
- MKA session
- MKA peer

See the Statistics search tool for detailed information about the statistics.

48.3 Workflow to configure MACsec

48.3.1 Stages

1

Verify that the discovery rule for each participating NE includes an SSH2 security access mediation policy; see 9.23 “To configure a discovery rule” (p. 310).

2

Create a global MACsec connectivity association; see 48.4 “To configure a global MACsec connectivity association” (p. 1457).

3

Create a local MACsec connectivity association; see 48.5 “To configure a local connectivity association” (p. 1458).

4

Create a global PSK; see 48.6 “To create a global PSK” (p. 1459).

5

Create a rekeying schedule; see 48.7 “To configure a rekeying schedule” (p. 1460).

6

Add interfaces to the MACsec connectivity association; see 48.8 “To add an interface to a MACsec connectivity association” (p. 1461).

After MACsec interfaces are added, the NE establishes an MKA session, discovers MACsec peers, and begins encrypting traffic between peers.

48.4 To configure a global MACsec connectivity association

48.4.1 Purpose

Perform this procedure to create a connectivity association for HSM key management.

Note: Before you can use an HSM for key management, you must add the HSM to the NFM-P configuration; see the procedure to add an HSM to the NFM-P in the *NSP System Administrator Guide*.

Note: The NFM-P does not store CAKs generated by an HSM.

48.4.2 Steps

1

Choose Manage → MACsec from the NFM-P main menu. The Manage MACsec form opens.

2

Click Create → Connectivity Association, or choose a connectivity association and click Properties. The Connectivity Association (Create|Edit) form opens.

3

Configure the required General parameters.

4

In the Keying Parameters panel, configure the Key Source parameter.

If you choose HSM, click the Select button and select an HSM.

5

Configure the required parameters for Static CAK.

6

Close the forms.

End of steps

48.5 To configure a local connectivity association

48.5.1 Purpose

Perform this procedure to add sites to a connectivity association and to configure local parameters. Parameter values are inherited from the global connectivity association by default. However, parameters with the Inherit Value check box can be configured with a value that is only applicable to the site.

48.5.2 Steps

1

Choose Manage → MACsec from the NFM-P main menu. The Manage MACsec form opens.

2

Choose a connectivity association and click Properties. The Connectivity Association (Edit) form opens.

3

Click on the Sites tab.

4

Click Add Site. The Select Site form opens. Click Search to populate the list.

5

Choose a site and click OK. The site is added to the connectivity association and appears in the list of sites.

6

To configure local parameters, choose the site from the Sites tab and click Properties. The Connectivity Association Site (Edit) form opens.

7

Configure local parameters as required.

Select the ANYSEC Reserved check box to enable ANYsec encryption. See 12.30 “To configure ANYsec encryption on an NE” (p. 366)

8

Click OK to close the Connectivity Association Site (Edit) form.

9

From the Sites tab, select the site and click Distribute PSKs to Sites.

Existing global PSKs are distributed to the site.

10

Close the forms.

END OF STEPS

48.6 To create a global PSK

48.6.1 Before you begin

Note: Distribution of a global PSK requires the `allow-immediate` parameter in the NE configuration to be set. If the parameter is set to `no-allow-immediate`, the PSK distribution will fail. See the NE documentation for more information about NE configuration.

48.6.2 Steps

1

Choose Manage → MACsec from the NFM-P main menu. The Manage MACsec form opens.

2

Choose a connectivity association and click Properties. The Connectivity Association (Edit) form opens.

3

Click on the Static CAK tab.

4

Click Create and choose an encryption type. The PSK is created.

The PSK will be automatically distributed to all the associated sites in the Connectivity Association.

5

Close the forms.

END OF STEPS

48.7 To configure a rekeying schedule

48.7.1 Purpose

Perform this procedure to configure PSK hitless rekeying.

A rekeying operation has the following stages:

- delete the old, inactive key, if it exists
- create and deploy a new key
- set the new key as active

When a rekeying operation is completed, the new key is displayed under the PSK list in the connectivity association. The status of the operation is displayed on the Rekey Schedule form Results tab. The status of the last operation is displayed on the Connectivity Association form General tab.

48.7.2 Steps

1

Choose Manage → MACsec from the NFM-P main menu. The Manage MACsec form opens.

2

Choose a connectivity association and click Properties. The Connectivity Association (Edit) form opens.

3

Click on the Rekey Schedule tab.

4

Click Create. The Rekey Schedule (Create) form opens.

5

Configure the required parameters.

6

In the Schedule panel, click Select and use the form that opens to choose or create a schedule.

7

Click OK. A rekeying scheduled task is created.

8

Close the forms.

END OF STEPS

48.8 To add an interface to a MACsec connectivity association

48.8.1 Steps

1

Choose Manage → MACsec from the NFM-P main menu. The Manage MACsec form opens.

2

Choose a connectivity association and click Properties. The Connectivity Association (Edit) form opens.

3

Click on the MACsec Interface tab.

4

Click Add. The Select MACsec Interface form opens.

5

Click Search to populate the list.

6

Perform one of the following:

a. Select an interface and click OK.

b. Create an interface.

1. Click Create. The MACsec Interface (Create) form opens.
2. Click Select to choose the site ID.
3. Click Select to choose the MACsec port.
4. Configure additional parameters as needed.
5. Click OK.

The new or selected MACsec interface is added to the connectivity association.

7

Close the forms.

END OF STEPS

Part V: Policy management

Overview

Purpose

This part provides information about policy management using the NFM-P.

Contents

49 Policies overview

49.1 Policies

49.1.1 Overview

The NFM-P supports the template-based creation of rules that govern how network traffic is handled and prioritized. These rules are referred to as policies. Policies provide a consistent, centralized configuration of common characteristics across multiple objects in the network. An NFM-P operator can distribute or delete policies on sites and NEs that are within their span of control.

You can assign policies to resources during service creation or modification. You can also assign policies to resources before or after you configure the service by choosing and modifying the resource from the Manage Equipment or Manage Services form.

Service and routing management policies are globally and seamlessly distributed to devices when they are used by resources on the device. They can also be manually distributed to devices. Subsequent changes to policies are distributed and affect all participating resources. Policy configurations can also be changed locally when you configure a network resource, for example, during service configuration or modification. These changes do not affect the global policy.

49.1.2 Global versus local policies

Policies are Global or Local in scope. Global policies are created using the NFM-P and are available for use throughout the network. Local policies are instances of global policies that are assigned to individual NEs and may contain properties applicable only to that NE.

When you modify and distribute a global policy using the NFM-P, the local policy definition is also updated, provided that its distribution mode is set to Sync With Global. This ensures that the policy instances are synchronized. If a local policy differs from the corresponding global policy because of changes to the global policy, a warning alarm is raised against the local policy. After a global policy is updated and distributed to the participating NEs, the NFM-P clears the mismatch alarms associated with the local policy.

When there is no global policy associated with a local policy, the NFM-P automatically creates a global policy that is identical to the first discovered local policy. If the local policy is incomplete, the NFM-P creates an incomplete global policy.

49.1.3 Global policies states

When you initially select a policy type under the main Policies menu, a selection window is displayed that lists all occurrences. You can choose to display Global or Local policies here. When you select Global, the menu bar over the list displays filterable and read-only parameters pertinent to the chosen policy type. The following read-only parameters describe the initialization phases of a Global policy:

- Discovery State

  The Discovery State parameter describes the state of the global policy during node discovery, or if a local policy on an NE is created or modified using CLI after the NE is discovered. The following may be displayed:
  - In Progress—the global policy is created as part of node discovery, and the policy definition is not yet complete.
  - Completed—the global policy is created as part of node discovery, and the global policy definition is successful and complete.
  - Failed—the global policy is created as part of node discovery, and the policy definition is not successful. This may occur because of an NFM-P server failure or activity switch. A failed global policy is updated during the next successful full NE resynchronization, or the policy can be manually synchronized with a specific local policy.
  - Initialized—the global policy is created by SNMP trap notification from a local policy that was created or modified on the NE by CLI after NE discovery. User may need to synchronize the global policy manually with the specific local policy. After such a policy synchronization on an “Initialized” global policy (using the NFM-P), the Discovery State for the global policy changes to “N/A”.
  - N/A (default)—the global policy is created or modified by an NFM-P user, and not as a result of NE discovery or CLI change.
- Last Sync Time

  The Last Sync Time field provides date and time information concerning the current state of an existing global policy that you are viewing.

  Note:

  This field displays “N/A” when you are first creating the global policy in the NFM-P.

  The Last Sync Time information field may denote any of the following timestamps:
  - the date and time when a policy that was created on an NE (using the CLI) was discovered by the NFM-P
  - the most recent date and time that the global policy was synchronized by the NFM-P with a corresponding local definition on an NE
  - N/A, indicating that the NFM-P has undergone an upgrade since the creation time of the global policy
- Last Sync From

  The Last Sync From field provides information concerning the origin of an existing global policy that you are viewing.

  Note:

This field displays “N/A” when you are first creating the global policy in the NFM-P, and then “admin” once the policy is initially applied.

The Last Sync From information field may denote any of the following origins:

- NFM-P user name (for instance, admin), indicating that the policy was created from the NFM-P
- System ID, indicating the NE from which the policy was learned by the NFM-P
- System ID, indicating the NE with which the policy was most recently synchronized by the NFM-P
- N/A, indicating that the NFM-P has undergone an upgrade since the creation of the global policy

Note: If you attempt to modify a global policy with an “In Progress” Discovery State, the following occurs:

- When discovery of the NE is ongoing, the modification fails.
- When discovery of the NE is completed, modifications using OSS proceed, although a warning appears and allows you to cancel the modification. If the modification is not canceled, the Discovery State is changed to N/A, and the global policy is not updated to match the local policy.
- When a global policy is modified, the Discovery State is reset to N/A.

At any time during discovery or a full resynchronization, you can use the SyncTo method from OSS or synchronize the global policy from any local policy.

When a discovery of an NE fails and there is no NFM-P server failure or activity switch, all global policies that are waiting for completion are changed to Failed for the Discovery State. If you start another full resynchronization of the failed NE, the NFM-P tries again to complete the global policies from the NE that is part of the resynchronization and which has the Failed or In Progress Discovery State. The NFM-P also attempts completion of the synchronization of global policies that have a Failed or an In Progress State until a server failure or an activity switch occurs.

The following table describes the recommended operator actions to recover global policies when the Discovery State remains In Progress or Failed.

Table 49-1 Actions to recover global policies

Discovery State	NFM-P server failure	NE state	Recovery action
In Progress/Failed	No	Resync Failed	Retry a full resynchronization or synchronization from another local policy
In Progress/Failed	No/Yes	Removed/Unmanaged	Synchronization from another local policy
In Progress/Failed	Yes	Resync Completed	Synchronization from another local policy
In Progress/Failed	Yes	Resync Failed	Retry a full resynchronization and synchronization from another local policy

49.1.4 Default global policy creation

When an NE is discovered, if a given global policy on the NE is not present in the NFM-P, then a non-empty global policy is created in the NFM-P. It has the same content as the originating NE's

policy. The local definition specific to the NE is set to “Sync With Global”, unless the originating policy is a routing policy, in which case the local definition is set to “Local Edit Only”.

When a global policy (other than a routing policy) is created using CLI on an NE already known to the NFM-P, the behavior is different. In this case, the NFM-P creates a new resynched global policy with only default mandatory values (for example, default queues). The NFM-P global policy does not contain objects specific to the NE-originated policy, such as filter entries. However, such objects are contained in the local definition specific to the NE. This local definition is set to “Local Edit Only”. In the case of a 7210 and 1830 Network policy, the local definition is set to “Sync With Global”, and the global policy created will be a copy of the local policy. For a CLI-created routing policy on a discovered NE, a non-empty global policy is created in the NFM-P. The global policy will be in sync with the local routing policy.

49.1.5 Policy audits

Nokia recommends that you change a policy on an NE using only the NFM-P or an OSS client. This ensures that there is no mismatch in policy IDs or policy configurations between the managed NEs and the NFM-P.

Creating or modifying a policy using a CLI may lead to inconsistencies in the policy configuration throughout the network.

Note: For policies that require you to enter a name to identify the policy, you cannot use “N/A” for the policy name; the term N/A is restricted.

Local changes to a policy may occur where users have permissions to access local policies using CLI. To identify local and global policy mismatches, you can perform a policy audit that compares the local and global instances of a policy, and then review the differences. Note however that the audit functionality is not applicable to non-modifiable default NFM-P policies. An audit can be performed on local policies of NEs that are in the user span of control.

You can audit all policies in the network or on selected NEs in the following ways:

- by policy group
- by policy type
- by global policy

If an entry exists in both the global and a local version of a policy but a field within that entry is different between the versions, then the audit indicates exactly which field is different. This includes missing entries as well. See 49.21 “To perform a policy audit for global policy” (p. 1498)
and 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503) for information about identifying policy discrepancies using a policy audit.

If there is a mismatch between a local policy instance and the global instance, an alarm is raised against the local policy. The NFM-P clears the alarms from previous audits if the mismatch condition no longer exists.

Depending on the results of an audit, the mode of the local policy may change. If the audit discovers differences between the local and global policies, you can change the mode to allow only local configuration. When the audit discovers that the local policy and global policy match, the local mode can be changed to synchronize with global policies. For example, when you need the NFM-P to discover new NEs, set the mode parameter to local configuration only to ensure that unique

policies are not modified when a global policy is updated. The audit results can be reviewed to determine whether a policy should remain local.

49.1.6 Policy synchronization

Changes to a local policy made using CLI will not result in a change to the related global policy. If you perform a policy audit of managed NEs and identify that the policy configuration is modified and the policies are out of sync, you can initiate a policy synchronization of the global policy with a selected local policy.

You can also do the reverse action, namely, update the global policy with a local policy definition and then release and re-distribute the policy to all deployed NEs. To prevent a catastrophic deployment of an invalid policy, the NFM-P changes the global policy to draft mode, and it is not available for use until the policy is released. See 49.13 “To synchronize a policy” (p. 1487) for information about how to synchronize a policy.

49.1.7 Policy overrides

With some policy types, you can use the NFM-P to configure a policy override that allows you to change or override the default settings associated with the policy. Policy overrides are typically configured on the object to which the policy is assigned.

For policy types that allow multiple overrides to be configured, the Override Policy Items tab on the policy creation form displays all of the subordinate override policies that are configured on the policy. Each override policy is displayed on their own sub-tabs under the Override Policy Items tab. The override policies that appear on the GUI are dynamic, based on the policy type to be configured.

49.2 Policy distribution

49.2.1 General Information

The NFM-P creates a global policy in Draft mode. This is the mode used to configure a policy before distributing it to the NEs. Once a policy is configured and ready for distribution, you change its configuration mode to Released. Whether you’re dealing with a new or existing policy, changing its mode to Released opens the Distribute window. If the policy was previously distributed to a group of NEs, then only those NEs are shown in the Distribute window. By default, all the NEs are selected in the window. You can deselect the NEs you want, or use the current selection to distribute the policy. If the policy was not previously distributed, or its local versions are non-existent (deleted either by a user or an NE), then the Distribute window displays a list of compatible NEs. You must then select which NEs to distribute the policy to. Alternatively, you can use the drop-down menu in the Distribute window to select a Policy Distribute Group for the distribution.

The NFM-P also supports the partial distribution of global policies. Before a global policy is distributed to the NEs, the NFM-P determines whether the global policy, policy properties, and the policy entries are applicable for each NE to which the policy is to be applied. If an NE does not support the policy, property, or entry defined in a global policy, the policy is partially distributed to the NE. The inapplicable policy, property, or entry is not distributed to the NE. An alarm is not raised if a policy is partially distributed. When you use the NFM-P GUI or OSS, failure of distribution to an NE does not affect distribution to other NEs.

You can monitor policy distribution either in the Distribute window or by using the Task Manager. The NFM-P saves the latest released version of the global policy. If you select a Policy Distribute Group, the NEs contained within the group will be shown in the Selected Objects panel once the distribution begins. This allows you to monitor the distribution progress at the individual NE level.

You can also interrupt or completely stop an ongoing policy distribution. See 49.12 “To stop a policy distribution currently in progress” (p. 1486) for more information.

49.2.2 Distribution considerations

Consider the following before distributing policies using the NFM-P:

- A policy must first be released before it can be distributed to an NE.
- The default behavior when you release a policy for distribution is that the Distribute window opens and allows you to select the NEs required for distribution. Alternatively, you can enable a parameter setting in the System Preferences that will upon release of the global policy, automatically distribute it to NEs that already have local versions of the policy. See the *NSP System Administrator Guide* for more information on the use of the Auto Distribute Global Policy when Released parameter.
- Local NE versions of policies that use the Sync With Global distribution mode will allow the NE to receive the distribution of a global policy.
- Local NE versions of policies that use the Local Edit Only distribution mode will not allow the NE to receive the distribution of a global policy. You must ensure that the policy distribution mode for the local policy is set to Sync With Global if you want the NE to receive the distribution of a global policy.
- Local SR-family NE versions of policies that are changed using CLI can have their distribution modes automatically set to either Sync with Global or Local Edit Only. This is governed by the setting of the Switch Distribution Mode to Local Edit Only on CLI Change parameter in System Preferences. See the *NSP System Administrator Guide* for more information about the Local Edit Only function.
- When you distribute a policy to a 7705 SAR, all values within that policy must be supported by that 7705 SAR; otherwise, the distribution of the policy to that 7705 SAR is blocked.

49.2.3 Scaling policy deployments

The NFM-P allows policies to be distributed to numerous NE sites in a single operation. To accomplish this effectively, the NFM-P can use multiple deployers to distribute the policies. A deployer in this context is a thread within the NFM-P that executes a task. Releasing a large global policy to multiple NEs using only one deployer may degrade system performance. Scaling policy deployments using multiple deployers helps to maintain system performance.

To avoid or minimize system degradation, an operator can configure the maximum number of managed objects that a single deployer is allowed to send during policy distribution. The default value is 10,000 managed objects per deployer. A managed object in this context is basically a configuration entry within a policy. For example, in an ACL IP filter policy, each IP filter entry is considered to be one managed object. Therefore a single global policy can often contain numerous managed objects.

When an operator sets the “policyDistributionMaxObjectsPerDeployer” parameter to a desired value, the NFM-P uses this value, along with the actual number of managed objects in the policy to

be distributed, to calculate the maximum number of NE sites that can receive the policy per deployer. If the number of sites selected for the policy's distribution exceeds this number, additional deployer requirements are automatically calculated and used by the NFM-P. This applies to the distribution of single as well as multiple global policies. For a single global policy, the values are calculated as follows:

Maximum number of NE sites allowed per deployer = Maximum number of objects per deployer / Number of objects in the global policy to be distributed

Number of deployers used to distribute the global policy = Total number of local NE sites to receive the global policy / Maximum number of sites per deployer

For the distribution of multiple policies, the number of deployers required is determined on a per policy basis.

Table 49-2, "Example of policy deployment" (p. 1470) provides an example of deploying a single global policy that contains 1,000 managed objects. Note that if the maximum number of objects per deployer setting is less than the actual number of managed objects in the policy, the NFM-P will only use one deployer per site. However, if the setting for the maximum number of objects per deployer is very large, the NFM-P might use a single deployer for a very large number of sites, which could degrade system performance. See 49.10 "To configure the maximum number of policy objects per deployer" (p. 1483) for more information on configuring the maximum number of managed objects allowed per deployer.

Table 49-2 Example of policy deployment

Maximum number of managed objects allowed per deployer	Number of managed objects in a single global policy	Maximum number of NEs (sites) per deployer
0	—	All sites
1	—	1 site
500	1,000	1 site
1,000	1,000	1 site
5,000	1,000	5 sites
10,000	1,000	10 sites
100,000	1,000	100 sites

When configuring the "policyDistributionMaxObjectsPerDeployer" parameter, the user also has the option to set the value to "0" or "1". A value of "0" means that only one deployer will be used for all NE sites that are to receive the policy. A value of "1" means that one deployer per site will be used. When distributing global policies with a large number of managed objects to a large number of NEs, Nokia recommends setting a value between 2,000 and 100,000 for this parameter.

49.3 Policy types

49.3.1 General information

The following categories define the high-level policy management functionality on the NFM-P:

- **Service management policies**

  Service management policies specify how service traffic is handled by network resources such as interfaces, ports, daughter cards, and circuits. The policies can be used by multiple resources on multiple services. Examples of service management policies include QoS policies, access ingress/egress policies, MSAP policies, and network policies.

- **Routing management policies**

  Routing management policies specify routing configuration according to specifically defined parameters that override the default routing protocol decisions normally handled by the router. Routing management policies control the size and content of routing table, advertised routes, and best routes. Examples of routing management policies include routing policies (statement, prefix list, community, and damping), MPLS administrative group, and VRRP priority control-policies.

- **Network management policies**

  Network management policies specify how the NFM-P communicates with network resources, handles alarms, manages statistics used for billing, and stores information. Examples of network management policies include alarm, mediation, and accounting policies.

- **Security policies**

  Security policies specify how the NFM-P handles the user security measures required to protect all NFM-P data, software, and hardware and monitor the system/network for any security threats. This includes creating policies required to set up NFM-P user accounts and user groups with the required scope of command roles and span of control permissions and the ongoing monitoring and management of those accounts. Examples of security policies include NE DoS/DDoS protection policies, NE password policies, RADIUS, TACAC+, and AOS/PKI authentication policies, and format and range policies.

49.3.2 Policy information map

The following table provides a high-level navigation aid to help you locate policy information contained in this guide and other NFM-P customer documentation.

Table 49-3 Location of NFM-P policy information

Applicable device	Policy type	Location
All NEs supported by the NFM-P, except for the devices listed below	Device backup, deployment, and software upgrade Service management Routing management, and residential subscriber QoS, Filter, and Multicast Event	*NSP NFM-P Classic Management User Guide*
	NFM-P user security Device security Database File Size constraint, and format and range Ageout constraint IPDR File transfer Alarm settings	*NSP System Administrator Guide*
	Statistics Accounting and file	*NSP NFM-P Statistics Management Guide*
CPAM and 7701 CPAA	Backup MIB Non-routed edge discovery Size constraint, and retention OAM diagnostics, and RCA audits	*NSP NFM-P Control Plane Assurance Manager User Guide*
Wavence SM and Wavence SA	Wavence QoS Multicast CAC interface Radio interface queue map	*NSP Wavence Device Support Guide*

49.4 Workflow to configure, distribute, and manage policies

49.4.1 Core tasks

1

Before creating policies, review the default setting of NFM-P policy-related system preferences and customize them as required. For example, you can display or hide the policy names on policy configuration forms for certain types of local policies or enforce a restriction to the distribution mode for certain types of local policies that will permit local editing only. See the *NSP System Administrator Guide* for information about configuring policy-related system preferences.

2

Create the required policy; see Table 49-3, “Location of NFM-P policy information” (p. 1473) to help locate the appropriate policy information.

Note: For policies that require you to enter a name to identify the policy, you cannot use “N/A”, “n/a”, or “@” for the policy name; the term “N/A” is restricted. Creating a policy with “@” in the

name results in Template Inconsistent/Configuration Mismatch alarms, which do not clear after a synchronization with the global policy. Also, updating the name using node CLI as “N/A” is not valid.

3

Release and distribute the policies:

a. To release and distribute select policies; see 49.6 “To release and distribute a policy” (p. 1476).

b. As required, create a Policy Distribute Group to specify a group of NEs to distribute a policy to; see 49.7 “To create a policy distribution group” (p. 1478).

c. To distribute multiple policies using a Policy Distribute Group in one action; see 49.8 “To distribute multiple policies” (p. 1479).

4

As required, change the deployment settings for policy distribution:

a. Change the distribution mode of local policies to synchronize with a global policy or to change a global policy back to a local edit distribution mode; see 49.9 “To change the distribution mode of a policy” (p. 1482).

b. Change the maximum number of managed policy objects that a distribution deployer is allowed to send; see 49.10 “To configure the maximum number of policy objects per deployer” (p. 1483).

c. Change the default maximum number of tasks involved when a group of policies is distributed; see 49.11 “To configure the maximum number of tasks for distribution” (p. 1484).

5

As required, stop the distribution of the policy while it is in progress; see 49.12 “To stop a policy distribution currently in progress” (p. 1486).

6

As required, synchronize policies across all deployed NEs; see 49.13 “To synchronize a policy” (p. 1487).

7

Assign policies to resources during service configuration or modification. See Chapter 70, “Service management and QoS”for more information.

8

As required, configure a policy override that allows you to change default settings associated for the following policies:

a. AAA subscriber policy overrides; see 13.8 “To configure an AA subscriber policy override on an ISA-AA group or partition” (p. 421).

b. Port scheduler policy overrides; see 16.24 “To configure Ethernet ports” (p. 599).

c. HSMDA policy overrides; see 16.40 “To configure an HSMDA override” (p. 632).

d.

QoS policy overrides on:

- L2 or L3 access interfaces; see 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654)
- access ingress meters for 7210 SAS; see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657).
- access ingress queues for 7210 SAS-X; see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659).

e. Epipe or Ipipe L2 access interface policy overrides; see 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174).

Incidental tasks

9

As required, manage NFM-P policies:

a. Enable or disable an exclusive policy editing restriction on one or more NEs or remove any exclusive edit locks that were previously set; see 12.18 “To configure an exclusive policy editing restriction on an NE” (p. 354).

b. Copy a policy; see 49.14 “To copy a policy” (p. 1488).

c. Modify a policy; see 49.15 “To modify a policy” (p. 1489) and 49.16 “To view local policy contents” (p. 1491).

d. Export existing global policies to a file in order to make them available for import elsewhere; see 49.17 “To export a policy” (p. 1492).

e. Import global policies from an exported policy file; see 49.18 “To import a policy” (p. 1493).

f. Delete a policy; see 49.19 “To delete a policy” (p. 1494).

49.5 Workflow to perform a policy audit

49.5.1 Stages

1

Complete one of the following to perform a policy audit:

a. To compare all local policies with the associated global policies for policy groups and types; see 49.20 “To perform a policy audit for policy groups and types” (p. 1495).

b. To compare all local policies with the associated global policy; see 49.21 “To perform a policy audit for global policy” (p. 1498).

c. To compare multiple global policies with same type; see 49.22 “To perform a policy audit for multiple global policies with same type” (p. 1499).

2

To schedule a policy audit to perform regular audits on selected policies and NEs; see 49.23 “To schedule a policy audit” (p. 1501).

3

As required, identify the differences between a global policy and a local policy or two local policies; see 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503).

49.6 To release and distribute a policy

49.6.1 Purpose

Perform this procedure to release a policy for distribution when a policy is in the Draft configuration mode and to distribute a policy to individual NEs or to multiple NEs in a Policy Distribution Group.

Review the 49.2.2 “Distribution considerations” (p. 1470) information in this chapter before starting this procedure. This procedure assumes you have created the required policy before you start this procedure. Perform 49.7 “To create a policy distribution group” (p. 1478) first if you intend to distribute the policy using a Policy Distribution Group.

Note: You do not need to explicitly distribute a policy; a policy is distributed to a device when it is assigned to a resource on the device.

You can configure the NFM-P to automatically distribute policies to applicable NEs upon release. See the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide* for more information.

49.6.2 Steps

Select a policy for release and distribution

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to release and distribute. The appropriate policy manager form opens.

2

Click Search and select a policy.

3

Double-click on the policy to view or edit the current policy parameters, if required.

Note: If you change any parameters on a policy that is in a Released Configuration mode, the configuration mode changes back to Draft. You need to re-release the policy in order to distribute the policy changes.

4

Perform one of the following:

a. If the policy is in a Draft Configuration Mode, go to Step 5 .

b. If the policy is in a Released Configuration Mode, click Distribute. The Distribute Instance form opens. Go to Step 6 .

Note: The policy configuration mode is displayed in the Configuration Mode column on the policy manager form.

Release the policy

5

Perform one of the following to release the policy for distribution:

a. Click Properties. The policy configuration form opens.

1. Click Switch Mode beside the Configuration Mode parameter. A dialog box appears.
2. Click Yes. The policy configuration mode changes to Released and the Release Policy Instance form opens.

b. Click Release. The policy configuration mode changes to Released and the Release Policy Instance form opens.

Each NE that has a local definition of the policy with its Distribution Mode set to Sync with Global is listed in the Selected Objects panel. These are ready for distribution.

Additional NEs that are eligible for policy distribution are listed in the Available Objects panel.

Select the NEs for distribution

6

Select the NEs for distribution in the Available Object panel and click on the right-arrow button. The chosen object(s) moves to the Selected Objects panel and is ready for distribution.

Note: The Available Objects and Selected Objects panels contain a column titled Local Policy Exists. This indicates if local definitions of the policy already exist on the eligible NEs.

Distribute the policy

7

Perform one of the following to distribute the policy:

a. To distribute the policy to the individual NEs, go to Step 8.

b. To distribute the policy to multiple NEs in a Policy Distribution Group:

Note: If you use the Policy Distribution Group method to distribute the policy, no additional NEs can be added to the policy distribution unless you modify the Policy Distribution Group boundaries.

See 49.7 “To create a policy distribution group” (p. 1477) for information about configuring Policy Distribution Groups.

1. Select the Policy Distribute Group option from the drop-down menu and click Search. The available Policy Distribute Groups are displayed in the list.
2. Select the appropriate Policy Distribution Group; go to Step 8.

8

Close the form. The policy manager form reappears.

Verify the policy distribution

9

To view the NEs to which the policy has been distributed:

1. Choose a policy.
2. Click Properties. The policy configuration form opens.
3. Click on the Local Definitions tab to view the list of NEs to which the policy has been distributed.

10

Close the policy manager form.

END OF STEPS

49.7 To create a policy distribution group

49.7.1 Purpose

Perform this procedure to create a Policy Distribution Group to specify a group of NEs to which you want to distribute a policy to. Using a Policy Distribution Group simplifies the distribution process for a set group of NEs, and allows you to easily add or remove members.

49.7.2 Steps

1

Choose Policies→Policy Distribute Group from the NFM-P main menu.The Policy Distribute Group form opens.

2

Click Create. The Policy Distribute Group (Create) form opens.

3

Configure the required parameters.

4

Click on the Network Elements tab.

5

Click Add and choose the NEs that you want to include in the distribution group and click OK.

6

Save your changes and close the form.

END OF STEPS

49.8 To distribute multiple policies

49.8.1 Purpose

Use this procedure to distribute a group of policies in one action. When you distribute a group of global policies:

- The policies must all be of the same type.
- The same conditions and restrictions that apply to the distribution of a single policy exist. See 49.6 “To release and distribute a policy” (p. 1476) for more information.
- You can set a limit to the maximum number of tasks involved in the bulk policy distribution. Perform 49.11 “To configure the maximum number of tasks for distribution” (p. 1484) to set this limit.
- The default behavior when you release a policy for distribution is the Distribute window opens, and you select the NEs required for distribution. Alternatively, you can enable the Auto Distribute Global Policy when Released parameter in the System Preferences to automatically distribute the Global policy to NEs that already have a local version of the policy. See the *NSP System Administrator Guide* for more information.
- Local NE versions of policies that use the Sync With Global distribution mode will allow the NE to receive the distribution of a global policy. However, local NE versions of policies that use the Local Edit Only distribution mode will not allow the NE to receive the global policy. You must ensure that the policy distribution mode for the local policy is set to Sync With Global if you want the NE to receive the global policy distribution.
- Multiple distribution is currently not supported for routing or firewall policies.

49.8.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to distribute. The appropriate policies manager form opens.

2

Choose Global from the Policy scope drop-down menu.

3

Click Search and select the policies that you want to release or distribute.

Note: If you select a combination of policies that are in Draft and Release mode, the Release and Distribute button will be disabled. The policies you select must be in one mode or the other.

4

Perform one of the following:

a. If the selected policies are in the Draft Configuration Mode, go to Step 5.

b. If the selected policies are in the Released Configuration Mode, go to Step 6.

5

When the selected policies are in the Draft Configuration Mode, the policies cannot be distributed to the NEs. You must first release the policies for distribution. Click Release. The Release form opens. Go to Step 9.

6

When the selected policies have previously been released (they are in the Released Configuration Mode) and you want to redistribute them, click Distribute. The Distribute form opens. The Network Elements (Network) option from the drop-down menu is displayed by default.

7

Perform one of the following:

a. If you want to select individual NEs for the distribution, go to Step 9.

b. If you want to use a policy distribution group for the distribution, go to Step 10 .

8

Click Distribute. The policy is distributed to the NEs.

A progress bar at the lower right of the GUI displays the overall distribution progress. Distribution states for individual NEs are shown in the Status column of the Selected Objects panel.

Note: Depending on the size of the policy, the number of NEs the policy is being distributed to, and the number of deployers used to distribute the policy, the progress bar status may not be noticeable and the policy may take considerable time to distribute. The progress bar reflects more than just the policy deployment; it also tracks the progression of the following events:

- polls the initial state of the NE to determine a baseline which is used to track the overall progression of the policy deployment on the NE
- the deployment of the policy to the NEs
- the resynchronization of the MIB tables that were changed on the NE

Note: If you are using a Policy Distribution Group for the distribution after you click the Distribute button, the NEs in the group are listed in the Available Objects panel. This allows you to monitor the distribution progress for specific NEs.

Note: Distribution failures are reported in a pop-up message. Distribution failures are also reported using JMS.

9

Select the NEs to release/distribute the policies to and click on the right-arrow. This moves the required row entries between the Available Objects and Selected Objects panels.

Observe the following behavior:

- The Available Objects panel is populated with a list of compatible NEs to which all the selected policies can be released/distributed.
- When local definitions exist, the list of NEs to which the policies have already been released and that have their Distribution Mode set to Sync with Global will appear in the Selected Objects panel. These are ready for distribution. Go to Step 11.

10

Select the Policy Distribute Group option from the drop-down menu and click Search. The available Policy Distribute Groups are displayed in the list.

Note: You cannot use a Policy Distribute Group and then add further individual NEs. See 49.7 “To create a policy distribution group” (p. 1478) for more information on Policy Distribute Groups.

11

Click Distribute to distribute the global policies to the selected devices.

If you are using a Policy Distribute Group for the distribution, once you click Distribute, the NEs contained in the group are listed in the Available Objects panel. This allows you to monitor the distribution progress for any specific NE.

A progress bar displays the overall percentage of distribution progress completion to the selected sites. Distribution states (In Progress, Succeeded, or Failed) for individual sites are shown in the Status column for the Selected Objects.

Note: Distribution failures are reported in a pop-up message. You can view the error message to determine the cause of the failure and which NE or NEs were affected. Distribution failures are also reported using JMS.

12

To stop the distribution of the policies, perform 49.12 “To stop a policy distribution currently in progress” (p. 1486).

13

Close the policy manager form.

END OF STEPS

49.9 To change the distribution mode of a policy

49.9.1 Purpose

Perform this procedure to change the distribution mode of local policies to synchronize them with a global policy or to change a global policy back to a local edit distribution mode.

49.9.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu where *Policy Type* is the type policy that you want to change the distribution mode of. The appropriate policy manager form opens.

2

Click Search and select a policy.

Note: The current policy distribution mode is displayed in the Distribution Mode column on the policy manager form.

3

Click Switch Distribution Mode. The Switch Distribution Mode - *Policy* form opens.

4

Choose one or more local definitions from the Available Local Policies list and click on the right-arrow. The definitions move to the Selected Local Policies panel.

5

Click Sync With Global or Local Edit Only to change the distribution mode of the policy. The Distribution Mode of the local definitions in the Selected Local Policies panel changes accordingly.

6

Save you changes and close the form. As required, distribute the policy to all participating NEs, as described in 49.6 “To release and distribute a policy” (p. 1476).

7

Close the policy manager form.

END OF STEPS

49.10 To configure the maximum number of policy objects per deployer

49.10.1 Purpose

Perform this procedure to set the maximum number of managed objects that can be distributed by one deployer. This in turn will determine how many deployers are used to distribute a global policy to a given set of NEs.

CAUTION

Service Disruption

Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file.

Modifying the nms-server.xml file can have serious consequences that can include service disruption.

49.10.2 Steps

1

Log in to the NFM-P main server station as an admin user.

2

Navigate to the /opt/nsp/nfmp/server/nms/config directory.

3

Create a backup copy of the nms-server.xml file.

4

Open the nms-server.xml file using a plain-text editor.

5

Search for the following XML tag:

```
<policyConfig policyDistributionMaxObjectsPerDeployer="10000"/>
```

6

Change the number as required. An integer value of “0” or less specifies that only one deployer will be used to distribute a policy to all the NE sites. A value of “1” specifies that one deployer will be used per site. The default value is 10,000 managed objects per deployer. Nokia recommends that you configure the value between 2,000 and 100,000.

Refer to the section 49.2.3 “Scaling policy deployments” (p. 1470) for more information on setting the maximum number of managed objects per deployer.

7

Save and close the nms-server.xml file.

8

Open a console window.

9

Navigate to the /opt/nsp/nfmp/server/nms/bin directory.

10

If the main server is a standalone server, or the primary server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash read_config ↲
```

The main server reads the nms-server.xml file and puts the configuration change into effect.

11

If the main server is the standby server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash force_restart ↲
```

The main server restarts and puts the configuration change into effect.

12

Close the console window.

END OF STEPS

49.11 To configure the maximum number of tasks for distribution

49.11.1 Purpose

The NFM-P allows you to select multiple policies of the same type for bulk distribution. Perform this procedure to change the default maximum number of tasks involved when a group of policies is distributed.

CAUTION

Service Disruption

Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file.

Modifying the nms-server.xml file can have serious consequences that can include service disruption.

49.11.2 Steps

1

Log in to the NFM-P main server station as the nsp user.

2

Navigate to the /opt/nsp/nfmp/server/nms/config directory.

3

Create a backup copy of the nms-server.xml file.

4

Open the nms-server.xml file using a plain-text editor.

5

Search for the following XML tag:

```
<policyConfig maxTasksForBulkDistribute="200000"/>
```

6

Change the number, as required. The default is 200,000. Nokia recommends that you configure the value between 1 and 200,000.

7

Save and close the nms-server.xml file.

8

Open a console window.

9

Navigate to the /opt/nsp/nfmp/server/nms/bin directory.

10

If the main server is a standalone server, or the primary server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash read_config ↵
```

The main server reads the nms-server.xml file and puts the configuration change into effect.

11

If the main server is the standby server in a redundant deployment, enter the following:

```
bash$ ./nmsserver.bash force_restart ↵
```

The main server restarts and puts the configuration change into effect.

12

Close the console window.

END OF STEPS

49.12 To stop a policy distribution currently in progress

49.12.1 Purpose

This procedure presumes you have initiated the distribution of a policy as described in 49.6 “To release and distribute a policy” (p. 1476) . To stop an ongoing policy distribution process, you must be in the Release or Distribute window, depending on whether you are currently releasing a draft policy or distributing a released policy. Stopping the release and distribution of multiple policies is also supported using this procedure.

49.12.2 Steps

1

Click Stop Distribute in the Release or Distribute window to halt the currently ongoing process. A confirmation message appears.

2

Click Yes to stop the distribution.

When the distribution is stopped, the Status column in the Selected Objects pane is updated with the current task status.

The status will indicate Succeeded for the policies that were successfully deployed to the targeted NEs.

The status will indicate Aborted for those policies that were not distributed.

3

If you want to continue the distribution, click Re-enabled Distribute. The NFM-P takes the list of aborted tasks and invokes the distribute function on them again.

When the distribution action is complete, the Status column in the Selected Objects pane is updated.

4

Close the policy manager form.

END OF STEPS

49.13 To synchronize a policy

49.13.1 Purpose

You can use the Synchronize function to specify a local policy as global and re-distribute it to all deployed NEs. For example, if a policy is distributed to a wide range of devices and on one of these devices the policy was changed, you can use the Synchronize command to synchronize the policy on the device and on the NFM-P. You must then re-distribute the policy to all participating devices, as described in 49.6 “To release and distribute a policy” (p. 1476).

Note: When you synchronize a policy that is in Draft mode, the policy is not distributed to existing local definitions. When you synchronize a policy that is in Released mode, the NFM-P sets the policy to Draft mode. You must set the Configuration Mode parameter to Released to distribute the policy to existing local definitions. Only local policies that have the Distribution Mode parameter set to Sync With Global are affected.

49.13.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to synchronize.

The Manage *Policy Type* form opens.

2

Choose Local or Global from the Policy scope drop-down menu. When you set it to Local, you can specify an NE using the Local Node IP Address parameter.

3

Click Search and choose the policy or policies that you want to synchronize.

4

Click Synchronize. The Synchronized - *Policy Type* form opens.

Note: The Synchronize function is available under the More Actions button. Do not confuse this with the Resync button. The Resync button on any property form is used to resynchronize the data in the NFM-P network model with the current state of the corresponding object.

5

Choose the NE to which the policy is to be synchronized from the Available Local Policies list and click on the right-arrow. The chosen NE moves to the Selected Source Local Policy panel of the form.

6

Configure the Resync Local Policy before synchronize? parameter as required.

7

Click Synchronize. The status of the synchronization process is displayed on the form.

8

Close the Synchronize - *Policy Type* form.

END OF STEPS

49.14 To copy a policy

49.14.1 Purpose

You can copy an existing policy and assign the copied version a unique policy ID. However, you cannot use this procedure to overwrite an existing policy.

49.14.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy you want to copy. The *Policy Type* manager form opens.

2

Choose Global or Local from the Policy scope drop-down menu. If you set it to Local, you must specify a Local Node IP Address by clicking Select and then choosing the required NE. Click OK.

3

Click Search and choose the policy you want to copy.

4

Click Properties. The *Policy Type* (Edit) form opens.

5

Choose Copy and configure the warning message as required. The copied *Policy Type* (Create) form opens. The form is set as a global policy in draft mode, regardless of whether you initially chose a global or local policy to copy.

6

Configure the parameters as required.

Note: If you need to configure the policy ID manually, you must assign a unique value. The NFM-P does not allow you to use the same ID as the policy you originally copied.

7

Save your changes and close the forms.

END OF STEPS

49.15 To modify a policy

49.15.1 Purpose

Use this procedure to change existing policies and entries within policies using the NFM-P client GUI.

49.15.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to edit. The *Policy Type* manager form opens.

2

Choose Global from the Policy scope drop-down menu.

3

Click Search and choose the policy that you want to modify.

4

Click Properties and perform one of the following:

a. If the policy configuration mode is set to Released, go to Step 5.

b. If the policy configuration mode is set to Draft, configure the parameters, as required, and go to Step 10.

5

Configure the parameters, if required. When you modify the configuration, the NFM-P changes the policy configuration mode to Draft.

6

When you modify filter entries for policies with filters, click on the appropriate tab and click Refresh to list the filter entries.

7

Save your changes.

8

Compare the global policy with the existing local definitions to verify whether you need to reset the policy configuration to the previously released global policy configuration. Perform 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503) if required.

9

Click on the Action button and choose Reset to Released to reset the global policy configuration to the last released configuration of the policy and to cancel any modifications you entered in Step 5.

10

Click Switch Mode beside the Configuration Mode parameter. A dialog box appears.

11

Click Yes. The policy configuration mode changes to Released and the Release form opens.

The NEs that have a local definition of the policy set to Sync with Global are shown in the Selected Objects panel and are ready for distribution. You can select an NE and move it to the Available Objects panel by using the left-arrow. This will exclude it from distribution.

12

Click Distribute to distribute the policy locally to the devices listed in the Selected Objects panel.

A progress bar at the lower right of the GUI displays the overall percentage of distribution progress completion to the selected sites. Distribution states (In Progress, Succeeded, or Failed) for individual sites are shown in the Status column for the Selected Objects.

Note: Distribution failures are reported in a pop-up message. You can view the error message to determine the cause of the failure and which NE or NEs were affected. Distribution failures are also reported using JMS.

13

Click on the Local Definitions tab to view the local instances of the global policy, if required.

14

Save your changes and close the form.

END OF STEPS

49.16 To view local policy contents

49.16.1 Purpose

Use this procedure to view various contents of an existing local policy definition and any modifications you make to it before deploying the changes.

Use this procedure to view various contents of an existing local policy definition and any modifications you make to it before deploying the changes.

- QoS SAP ingress and QoS SAP egress polices
- ACL IP/IPv6 filters and ACL MAC filter policies
- Port list and Prefix lists
- All routing policies

  Note:

  Current and modified contents to VPLS (see 77.10 “To view VPLS contents” (p. 2260)), VLL (see 76.9 “To view VLL service contents” (p. 2126)), and VPRN services (see 79.8 “To view VPRN service contents” (p. 2537)) may also be viewed in a similar manner as described in this procedure for policies.

Two information views are available. The Committed info view displays the current contents of a policy, and the Committed menu item is always enabled. The Modified info view allows you to review any changes you make to a policy before committing them. Modified, created, and deleted attributes and objects are displayed.

49.16.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to edit. The *Policy Type* manager form opens.

2

Choose Local from the Policy scope drop-down menu.

3

Click Search and select the policy that you want to view or modify.

4

Click Properties and perform one of the following:

a. To view the currently committed policy contents, go to Step 5 .

b. To view modified policy contents before committing any changes, go to Step 6 .

5

Click on the Action button and choose Show Info→Committed. A Committed Values form is displayed and shows various policy attributes that have been previously configured and saved

in the local policy. The information displayed in the form is similar to the information retrieved by running the “configure>*Policy Type #*”, where *Policy Type #* is the type of policy (followed by its Service ID) that you want to query on an NE.

6

Modify any policy parameters or objects, as required. Otherwise go to Step 9 .

7

Click on the Action button and choose Show Info→Modified. A Modified Values form is displayed. The table lists modified, created, and deleted actions, as well as specific attributes and objects, along with their old value, new value, and tab location. The Attribute Title corresponds to the attribute or object name acted upon by your current modifications. For created objects, the values of mandatory attributes are shown in comma-separated format.

8

Select an item in the Modified Values form and then click Show on Form. The policy form tab containing the changed item is displayed and the modified attribute is highlighted in blue.

9

Save your changes if required, and close the form.

END OF STEPS

49.17 To export a policy

49.17.1 Purpose

You can export existing global policies to a file in order to make them available for import elsewhere, for example, on another NFM-P server. The following types are available for export:

- SAP access ingress and access egress policies
- ACL filter policies (both IPv4 and IPv6)
- Accounting policies
- Scheduler policies
- All SR-based QoS policies
- All routing policies

See 49.18 “To import a policy” (p. 1493) for information about importing policies.

49.17.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy you want to export. The *Policy Type* manager form opens.

2

Choose Global from the Policy scope drop-down menu.

3

Click Search and choose one or more policies you want to export.

4

Click Export Policy. The Export Global Policy Details form opens.

5

Choose the location to save the exported policy:

1. Choose Server or Local from the File Location menu.
   Policies saved to the server will be saved to /opt/nsp/nfmp/server/nms/bin/policyExport.
2. To save the policy locally, click Select and choose the file path.

6

Assign an Export File Name to the exported policies. You can also enter an optional description.

Note: If you selected more than one policy in Step 3, the exported xml file will contain all the selected policies.

7

Save your changes and close the form.

END OF STEPS

49.18 To import a policy

49.18.1 Purpose

You can import global policies from an exported policy file. The following types are available for import:

- SAP access ingress and access egress policies
- ACL filter policies (both IPv4 and IPv6)
- Accounting policies
- Scheduler policies
- All SR-based QoS policies
- All routing policies

Note: You can import global policies that have been exported on the same or lower release of the NFM-P.

See 49.17 “To export a policy” (p. 1492) for information about exporting policies.

49.18.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy you want to import. The *Policy Type* manager form opens.

2

Choose Global from the Policy scope drop-down menu.

3

Click Import Policy and select an Import File Name.

4

Choose the location of the policy to import:

1. Choose Server or Local from the File Location menu.
2. If the policy is saved locally, click Select and choose the file path.

5

Click Preview List to see the global policies contained in the selected file.

6

Click the Override existing policies check box if you want the imported policies to overwrite existing global policies.

7

Save your changes and close the forms.

END OF STEPS

49.19 To delete a policy

49.19.1 Purpose

Each service SAP and network interface is associated, by default, with the appropriate ingress, egress, or network policy (ID 1). You can replace the default policy with a customer-configured policy, but you cannot entirely remove a QoS policy. When you remove a QoS policy from a SAP or IP interface, the policy association reverts to the default policy (ingress or egress policy ID 1). The default policy cannot be deleted.

A QoS, ACL, or an AA transit IP policy cannot be deleted until it is removed from the all SAPs or network ports where it is applied. When a policy is deleted, it is removed from the NFM-P, including the database and all devices.

Be aware when deleting a routing policy from an NE that is associated with one or more services. Multiple services on the NE may use the same policy.

49.19.2 Steps

1

Choose Policies→*Policy Type* from the NFM-P main menu, where *Policy Type* is the type of policy that you want to delete. The Manage *Policy Type* form opens.

2

Choose Local or Global from the Policy scope drop-down menu.

3

Click Search and choose the policy that you want to delete.

4

Click Delete. A confirmation message appears.

If you are trying to delete a global routing policy, the message informs you that the policy may currently be used by services or routing protocols. Some routing policies, such AS Paths and routing statements, contain a Policy Usage tab in their properties form. That tab allows you to view the services and routing protocols that currently use the policy you are attempting to delete (see 54.21 “To view routing policy usage” (p. 1767) for more information). If you proceed to delete the global policy, then all such local definitions of the policy will also be deleted.

The confirmation message also contains a check box that allows you to exclude a global policy that is currently used by a service or routing protocol. Configure this as required.

5

Click Yes or No to continue. The policy manager form reappears. If you chose to delete the policy, it is removed from the policy list.

END OF STEPS

49.20 To perform a policy audit for policy groups and types

49.20.1 Purpose

An audit compares all local policies with the associated global policies. A policy audit can be performed on all NEs in the network, or limited to a specific NE or group of NEs. All NEs are included in the audit unless they are explicitly chosen. A policy audit can be performed on policy groups and policy types. A policy audit cannot be started when another policy audit is in progress.

CAUTION

Service Disruption

If a policy type of global policy is not specified, all policies in the user span of control are included in the policy audit.

If NEs are not specified, all NEs that are in the users span of control are included in the policy audit.

Performing a policy audit may take one or more hours to complete.

Note: Validation is performed to prevent policies from being audited if the use span of control prohibits the operation.
If a user performs an audit on objects that are not within their span of control, the audit is aborted and a error message is displayed. The error message indicates which objects are prohibited.

49.20.2 Steps

1

Choose Tools→Policies Audit from the NFM-P main menu. The Policy Audit - audit status (Edit) form opens.

2

Configure the required parameters.

3

Perform one of the following steps:

a. To choose policy groups and policy types to be included in the audit, go to Step 4.

b. If all policy groups and policy types are included in the audit, go to Step 10 to configure which NEs to include in the audit.

c. If all policy groups and policy types and the entire network is included in the audit, go to Step 11.

4

Click on the Policy Group Selection tab.

5

Right-click on the Policy Audit icon and choose Select Policy Groups from the contextual menu. The Select Policy Groups form opens.

6

Select policy groups as required and click OK.The chosen policy group keys appear below the Policy Audit icon.

7

Perform one of the following steps:

a. To choose which policy types within the selected groups are included in the audit, go to Step 8.

b. If all policy types within the selected groups are included in the audit, go to Step 10.

8

Right-click on a policy group icon and choose Select Policy types. The Select Policy types - Policy Audit form opens.

9

Select the policy types as required and click OK. The chosen policy type icons appear below the Policy Audit icon.

10

Select NEs to include in the audit, if required:

1. Click on the Network Element Selection tab to select an NE, or group of NEs to be audit.
2. Click Add. The Select Site form opens.
3. Click Search and choose the NEs to be included in the audit.
4. Save your changes and close the form.

11

Click Start Audit to proceed with the policy audit. Click Stop Audit to interrupt the audit while it is in progress. A discrepancy between a global policy and a local instance of the policy generates an alarm. Monitor the dynamic alarm window to view alarms that are generated as a result of the policy audit.

12

Double-click on the alarm in the alarm window. The Alarm Info form opens.

13

Click View Alarmed Object. The Global Policy (Edit) form opens.

14

Perform Step 6 to Step 8 of 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503) to locate the discrepancies between the local and global policies.

15

Close all open forms.

END OF STEPS

49.21 To perform a policy audit for global policy

49.21.1 Purpose

An audit compares all local policies with the associated global policy. A policy audit can be performed on all NEs in the network, or limited to a specific NE or group of NEs. All NEs are included in the audit unless they are explicitly chosen. A policy audit cannot be started when another policy audit is in progress.

CAUTION

Service Disruption

Performing this procedure is potentially service-affecting.

If NEs are not specified, all NEs that are in the users span of control are included in the policy audit.

Note: Validation is performed to prevent policies from being audited if the user span of control prohibits the operation.

If a user performs an audit on objects that are not within their span of control, the audit is aborted and a error message is displayed. The error message indicates which objects are prohibited.

49.21.2 Steps

1

Choose Policies→ *Policy_Type* from the NFM-P main menu, where *Policy_Type* is the type of policy that you want to audit. The appropriate policies manager form opens.

2

Click Search and choose a policy in the list.

3

Click Properties. The *Policy_Name* Global Policy (Edit) form opens.

4

Click on the Local Definitions tab. A list of local policies is displayed.

5

Perform one of the following:

a. To audit a specific local policy or a group of local policies, go to Step 6.

b. If all local policies in the list are included in the audit, go to Step 7.

6

Select one or multiple local policies that you want to include in the audit.

7

Click Policy Audit. A Policy Audit (Edit) form opens.

8

Click on the General tab.

9

Configure the required parameters.

10

Click Start Audit to proceed with the policy audit. Click Stop Audit to interrupt the audit while it is in progress. A discrepancy between a global policy and a local instance of the policy generates an alarm. Monitor the dynamic alarm window to view alarms that are generated as a result of the policy audit.

11

Double-click on the alarm in the alarm window. The Alarm Info form opens.

12

Click View Alarmed Object. The Global Policy (Edit) form opens.

13

Perform Step 6 to Step 8 of 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503) to locate the discrepancies between the local and global policies.

14

Close all open forms.

END OF STEPS

49.22 To perform a policy audit for multiple global policies with same type

49.22.1 Purpose

An audit compares all local policies with the associated global policies. A policy audit can be performed on all NEs in the network, or limited to a specific NE or group of NEs. All NEs are included in the audit unless they are explicitly chosen. A policy audit can be performed on one or multiple global policies with the same type. A policy audit cannot be started when another policy audit is in progress.

CAUTION

Service Disruption

Performing this procedure is potentially service-affecting.

If NEs are not specified, all NEs that are in the users span of control are included in the policy audit.

Note: Validation is performed to prevent policies from being audited if the user span of control prohibits the operation.

If a user performs an audit on objects that are not within their span of control, the audit is aborted and a error message is displayed. The error message indicates which objects are prohibited.

49.22.2 Steps

1

Choose Policies→ *Policy_Type* from the NFM-P main menu, where *Policy_Type* is the type of policy that you want to audit. The appropriate policies manager form opens.

2

Click Search and choose one or multiple policies in the list.

3

Click Policy Audit. A Policy Audit (Edit) form opens.

4

Click Add and select a policy or multiple policies in the list.

5

Click OK. The Policy Audit (Edit) form refreshes with the selected policies.

6

Click on the General tab.

7

Configure the required parameters.

8

Perform one of the following:

a. To audit selected global policies on selected NEs, go to Step 9.

b. To audit selected global policies on the entire network, go to Step 10.

9

Select NEs to include in the audit, if required.

1. Click on the Network Element Selection tab to select an NE, or group of NEs to be audit.
2. Click Add.The Select Site form opens.
3. Click Search and choose the NEs to be included in the audit.
4. Save your changes and close the form.

10

Click Start Audit to proceed with the policy audit. Click Stop Audit to interrupt the audit while it is in progress. A discrepancy between a global policy and a local instance of the policy generates an alarm. Monitor the dynamic alarm window to view alarms that are generated as a result of the policy audit.

11

Double-click on the alarm in the alarm window. The Alarm Info form opens.

12

Click View Alarmed Object. The Global Policy (Edit) form opens.

13

Perform Step 6 to Step 8 of 49.24 “To identify differences between a global and local policy or two local policies” (p. 1503) to locate the discrepancies between the local and global policies.

14

Close all open forms.

END OF STEPS

49.23 To schedule a policy audit

49.23.1 Purpose

Perform this procedure to schedule a policy audit to perform regular audits on selected policies and NEs. Prior to scheduling a policy audit, you must first create the NFM-P-based schedule; see 5.6 “To configure an NFM-P-based schedule” (p. 193) .

An audit compares all local policies with the associated global policy. A policy audit can be performed on all NEs in the network, or limited to a specific NE or group of NEs. All NEs are included in the audit unless they are explicitly chosen. A policy audit cannot be started when another policy audit is in progress.

CAUTION

Service Disruption

Performing this procedure is potentially service-affecting.

If NEs are not specified, all NEs that are in the users span of control are included in the policy audit.

Note: Validation is performed to prevent policies from being audited if the user span of control prohibits the operation.

If a user performs an audit on objects that are not within their span of control, the audit is aborted and a error message is displayed. The error message indicates which objects are prohibited.

Note: You can configure the maximum number of scheduled audits that are stored for a local policy by setting the Max Number Of Audit Result stored per Local Policy parameter in the System Preferences. See the *NSP System Administrator Guide* for more information.

49.23.2 Steps

1

Choose Tools→Policies Audit from the NFM-P main menu. The Policy Audit (Edit) form opens.

2

Configure the required parameters.

3

Click Schedule Audit and choose Create. The Policy Audit Scheduled Tasks (Create) form opens.

4

Configure the required parameters.

5

Click Select in the Schedule panel and choose a schedule from the Select Schedule - Policy Audit Scheduled tasks form.

6

Click OK and close the form. The Policy Audit Scheduled Tasks (Create) form reappears.

7

Click on the Audit tab.

8

Select one or more policy types that you want to include in the audit from the Selected Policies panel.

9

Select one or more NEs that you want to include in the audit from the Selected NEs panel.

10

Click Apply. The Policy Audit Scheduled Tasks (Create) form is updated with the changes.

11

Click Start Audit or Stop Audit operations, as required. The User Activity displays the creation, modification, start and stop operations of the scheduled audit.

12

Once a scheduled policy audit is complete, click Schedule Audit and choose View from the menu to access the audit results. The Policy Audit Schedule Tasks list form opens.

13

Select a scheduled task from the list and click Properties.The Policy Audit Scheduled Tasks (Edit) form opens.

14

Click on the Audit Scheduled Task Results tab to view a list of results.

15

Select an entry and click Properties to review specific results. The Viewing Old Audit Result indicator determines if you are viewing the latest audit result or not.

16

Close all open forms.

END OF STEPS

49.24 To identify differences between a global and local policy or two local policies

49.24.1 Purpose

Global policies can be distributed to many varying NE types. Local policy audits are performed based on absolute database records. In some cases, this means that when you audit a local policy, a change may be indicated on the associated global policy that does not apply specifically to this local policy. The changed global property may simply not be relevant to the NE associated with this local policy.

For example, a Subscriber Profile global policy contains an HSMDA QoS tab. If a change has been made to a property on this tab in the global policy, but the local policy you are auditing applies to an NE that does not support HSMDA, then the following will occur. When you perform the local audit against the global policy, the global policy form will indicate that there is a difference between the two by displaying an arrow icon on its HSMDA QoS tab. However, since the local policy's NE does not support HSMDA, the local policy form will not display an HSMDA tab for you to compare with. Any such indicated global changes are irrelevant to the local policy and can be ignored.

49.24.2 Steps

1

Choose Policies→*Policy_type* from the NFM-P main menu, where *Policy_type* is the type of policy that you want to search for global and local differences. The Manage *Policy_type* form opens.

2

Choose Local from the Policy scope drop-down menu.

3

Click Select and choose a NE in the list and click OK.

4

Click Search and choose the local policy that you want to compare with another policy.

5

Click Properties. The *Policy_type* (Edit) form opens.

6

Click Local Audit On.

Note: You can cancel the local audit at any time by clicking on the Action button and choosing Local Audit Off on the *Policy_type* (Edit) form.

7

From the Policy scope drop-down menu:

a. Choose Global and click OK. The global opens for comparison.

b. Choose Local and perform the following:

1. Click Select to choose an NE. The Select a Network Element form opens.
2. Select the required NE from the list and click OK. The local policy opens for comparison.

8

Perform one of the following steps, as applicable:

a. View the differences between the policies by clicking on the tabs that are highlighted with an

arrow icon, which indicates that differences exist between the forms. A purple arrow indicates that an attribute is set differently between the two audited forms. A pink arrow indicates that an attribute is missing. A purple arrow in a pink background rectangle indicates that both conditions are occurring. In lists, new entries are highlighted in pink and modified entries are highlighted in purple. Messages are also provided above a list to help isolate the differences on the affected form. The example shown in Figure 49-1, “Policy audit mismatches” (p. 1504) demonstrates some of these various conditions and indications.

Figure 49-1 Policy audit mismatches

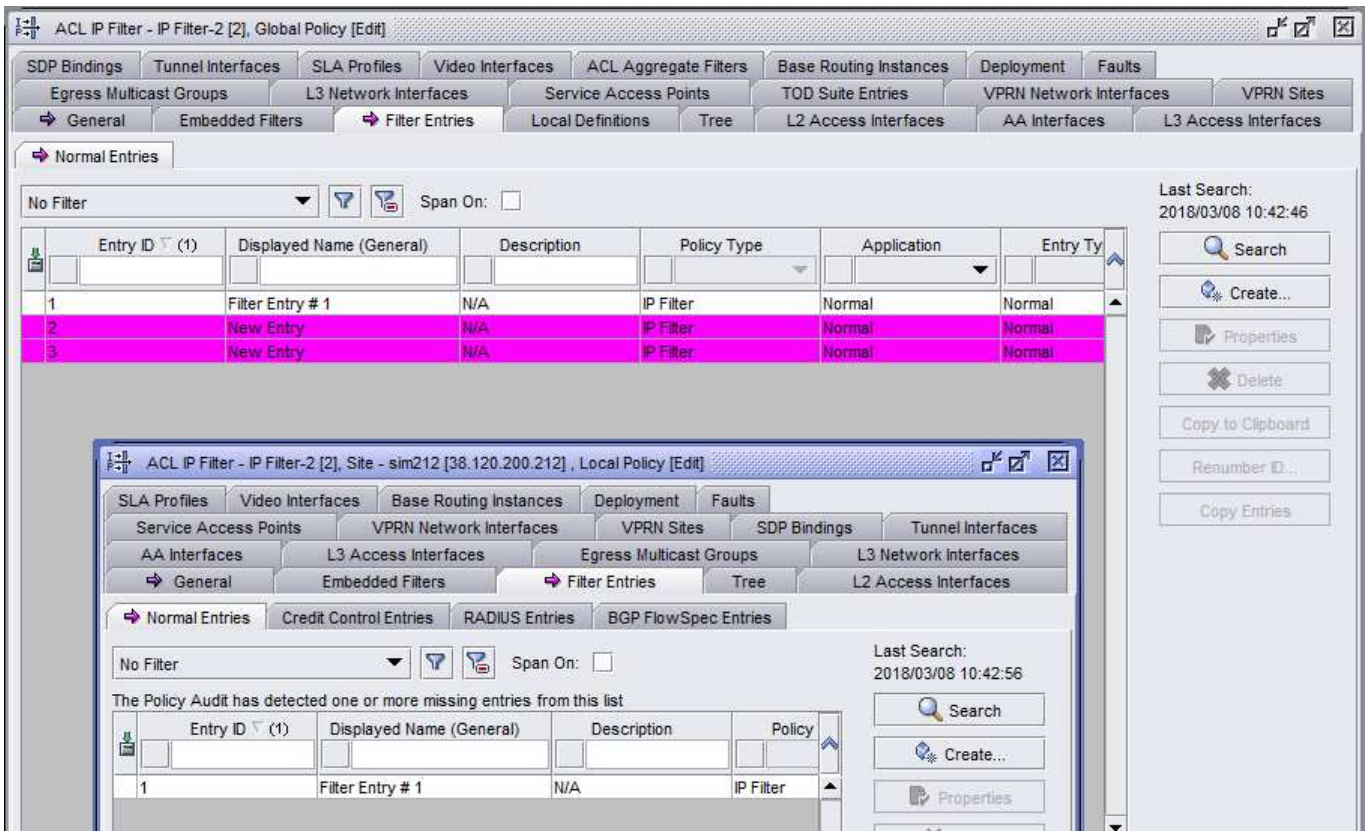

b. For certain policy types, you can click on a Tree tab on the local or global policy form. New items are highlighted in pink text. Modified items in the tree are highlighted in purple text.

9

Close the local and global *Policy_type* (Edit) forms.

END OF STEPS

50 QoS policies

QoS policy types

50.1 Overview

50.1.1 General Information

QoS policies define how network traffic is shaped and queued on one or more NEs. You can use the NFM-P to create QoS policies that regulate data throughput on the following:

- equipment; for example, ports and MDAs
- routing and forwarding points; for example, access and network interfaces

The NFM-P provides QoS policies that are specific to 7210 SAS and 7250 IXR NEs.
See 50.23 “7210 SAS QoS policies” (p. 1528) and 50.24 “7250 IXR QoS policies” (p. 1534) in this chapter.

50.2 SAP access ingress policies

50.2.1 General Information

SAP access ingress policies are applied to access interfaces and specify QoS on ingress.

SAP access ingress policies define ingress service forwarding class queues and map flows to those queues. When an access ingress policy is created, it always has two queues defined that cannot be deleted: one for the default unicast traffic and one for the default multipoint traffic. These queues exist within the definition of the policy. The queues only get instantiated in hardware when the policy is applied to an access interface. In the case where the service does not have multipoint traffic, the multipoint queue is not instantiated.

In the simplest access ingress policy, all traffic is treated as a single flow and mapped to a single queue, and all flooded traffic is treated with a single multipoint queue.

The required SAP access ingress policy elements include:

- a unique access ingress policy ID
- an exclusive scope for one-time use, or a template scope for use with multiple SAPs and interfaces
- at least one default unicast forwarding class queue
- at least one multipoint forwarding class queue

The optional SAP access ingress policy elements include:

- additional unicast queues up to a total of 8 for each of the 8 forwarding classes
- 8 HSMDA queues that are automatically created when the policy is created; HSMDA queues are only used by HSMDA SAPs

- additional multipoint queues up to 3 per forwarding class for each type of multipoint traffic (broadcast, multicast and destination unknown unicast)
- QoS policy match criteria to map packets to a forwarding class
- A policer to control traffic flow rate
- A dynamic policer to specify traffic flow rate for dynamically-spawned policers

Each queue can have unique queue parameters to allow individual policing and rate shaping of the flow mapped to the forwarding class. Mapping flows to forwarding classes is controlled by comparing each packet to the match criteria in the policy.

There is one default SAP access ingress policy. The default policy gives all traffic equal priority with the same chance of being sent or dropped during periods of congestion.

Note: The NFM-P supports the configuration of HQoS scheduling mechanisms. HQoS provides the ability to rate limit across multiple queues from single or multiple access interfaces for a specific customer. The building blocks for HQoS include access ingress, access egress, and scheduler policies.

See 70.12 “Sample network QoS configuration” (p. 1942) for a sample service configuration using HQoS.

50.2.2 Policy override

You can override some or all settings associated with an access ingress policy on an L2 or L3 access interface, SLA profile, or subscriber profile. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654).

Note: You can override SAP access ingress policies that have the Scope parameter set to template; see 50.28 “To configure a SAP access ingress policy” (p. 1538).

Tagged IP/IPv6 match entry overrides

In a SAP access ingress policy, you can assign a numbered tag identifier to IP and IPv6 match criteria entries. When the SAP access ingress policy is assigned to an L2 or L3 SAP, you can configure a QoS policy override to use the tag identifiers to select tagged entries with the same identifier. Any untagged entries in the policy are still valid on the SAP. The IP or IPv6 Criteria Type in the SAP access ingress policy must be set to Tagged Entries.

The following task flow lists the basic steps required to enable tagged IP/IPv6 match entry overrides.

1. Plan the match criteria and tag mapping as required.
2. Create a SAP access ingress policy and set the IP or IPv6 Criteria Type parameters in the policy to Tagged-Entries; see 50.28 “To configure a SAP access ingress policy” (p. 1538).
3. In the same policy, create IP or IPv6 match criteria entries and assign tag IDs to the entries.
4. Distribute the policy to NEs that support tagged match entries.
5. For the required L2/L3 SAPs, configure a QoS policy override to define the tag IDs to use on the SAP; see50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654).

50.2.3 Forwarding classes

The NFM-P supports the configuration of eight forwarding classes and class-based queuing or policing on the managed devices. Each forwarding class is only important in relation to other forwarding classes. A forwarding class provides NEs with a method to determine the relative importance of one packet over another packet in a different forwarding class.

Queues are created for a specific forwarding class to determine how the queue output is scheduled into the switch fabric and the type of parameters that the queue accepts. The forwarding class of the packet, and the in-profile and out-of-profile states, determine how the packet is queued and handled at each hop along its path to a destination egress point. Forwarding classes may also be associated with policers instead of queues. Eight forwarding classes are supported. The table below lists the default definitions for the supported forwarding classes.

Although all forwarding classes support profile marking, it is a good network engineering practise to ensure that all high priority forwarding classes are in-profile (CIR=PIR) and all low priority forwarding classes are out-of-profile (PIR > CIR=0). This way, distinguishing packets as in-profile or out-of-profile only occurs for assured class types.

Table 50-1 Forwarding classes

Forwarding class ID	Forwarding class name	Forwarding class designation	DiffServ name	Class type	Intended
7	Network control	nc	nc2	High priority	For network control traffic
6	High-1	h1	nc1		For a second network control class or delay/jitter sensitive traffic
5	Expedited	ef	ef		For delay/jitter sensitive traffic
4	High-2	h2	h2		For delay/jitter sensitive traffic
3	Low-1	l1	af2	Assured	For assured traffic; default priority for network management traffic
2	Assured	af	af1		For assured traffic
1	Low-2	l2	cs1	Best effort	For best effort traffic
0	be	be			

50.2.4 Forwarding subclasses

You can use forwarding subclasses for additional access ingress packet classification. One or more subclasses can be associated with each forwarding class. The designations for forwarding subclasses are the same as those for the forwarding classes listed in Table 50-1, “Forwarding classes” (p. 1509) . Each subclass assumes the behavior of its parent forwarding class. and in combination with the forwarding class provides a greater range of access ingress QoS classification possibilities.

50.2.5 Enabling ETH-LMM Y.173 statistics based on QoS forwarding classes

The ETH-LMM Y.1731 approach to Ethernet frame loss measurements allows for the collection of frame counters in order to determine the unidirectional frame loss between point-to-point ETH-CFM MEP peers. See the supported NE OAM and Diagnostics guide for more information about the ETH-LMM Y.1731 protocol.

You can enable the collection of ETH-LMM Y.1731 statistics on the NFM-P based on eight forwarding classes (either profile aware, or profile unaware) for the following Service types: EPIPE, VPLS, VPRN, and IES. Use the Collect ETH-LMM FC Stats parameter or Collect ETH-LMM FC Stats in Profile parameter on the supported Services navigation tree toolbar to enable ETH-LMM Y.1731 statistics collection as follows:

- Service Type→Sites→L2 or L3 Access Interfaces→SAP. On the Port OAM tab, ETH-CFM sub-tab, select the LMM Session Stats Collection panel and configure the parameters.
- Service Type→Sites→Spoke SDP Bindings. On the Service OAM tab, ETH-CFM sub-tab, select the LMM Session Stats Collection panel and configure the parameters.

See the XML API Reference for information about configurable parameters, their applicability, and the NFM-P GUI forms from which the parameters can be accessed.

50.2.6 Policers

You can add a policer to an access ingress policy to provide traffic flow limiting. Policers are associated with the forwarding classes defined in the access ingress policy. Policers can also be linked to a policer control policy, which maintains a hierarchy of multiple policer objects in the NFM-P system. Policers are linked to policer control policies by means of an arbiter. For more information, see 50.16 “Policer control policies” (p. 1524).

When policers are configured in an access ingress policy, they can be used to generate statistics. The type of statistics generated depends on the Stats Mode. The following table lists the options for counter allocation when generating statistics:

Table 50-2 Stats Mode options

Option	Description
No Override	The system does not override the value of Stats Mode set in the policer control policy.
Minimal	The system creates one offered-output counter in the network processor and one discard counter in the QChip. Packet priority, initial profile, and CIR profile output are ignored, and are not individually visible in the policer statistics. Use the Minimal option when only basic policer accounting is required.
No Stats	The system does not allocate any counters to the policer. The absence of counters does not affect the operation of the policer. Use the No Stats option when policer accounting is not required. You cannot enable this option if the policer has a parent arbiter assigned to it.

Table 50-2 Stats Mode options (continued)

Option	Description
Offered Limited Capped CIR	The system creates two offered-output counters in the network processor and two discard counters in the QChip. Use the Offered Limited Capped CIR option when ingress packets are being classified as either in-profile, or ignored when the traffic is in excess of CIR. Packet priority is ignored, and is not separated in the accumulated statistics.
Offered Limited Profile CIR	The system creates three offered-output counters in the network processor and three discard counters in the QChip. Use the Offered Limited Profile CIR option when ingress color-aware profiling is in use but packets are not being classified as in-profile. In most color-aware instances, only undefined and explicit out-of-profile packets are used in service offerings. If an in-profile classified packet is offered to the policer, it is included in the offered-undefined statistic. Packet priority is ignored, and is not separated in the accumulated statistics.
Offered Priority CIR	The system creates four offered-output counters in the network processor and four discard counters in the QChip. Use the Offered Priority CIR option when the ingress policer is used on a non-color-aware mode (all ingress packets have an undefined initial profile), but packet priority input and CIR state output visibility is required.
Offered Priority No CIR	The system creates two offered-output counters in the network processor and two discard counters in the QChip. Use the Offered Priority No CIR option when ingress packet priority (high profile and low profile classification) accounting is the primary requirement. The initial and output profile of the packets offered to the policer is ignored in the offered, discarded, and forwarded statistics. This mode does not inhibit the function of CIR on the ingress policer and it does not prevent explicit in-profile and out-of-profile classification for packets offered to the policer.
Offered Profile Capped CIR	The system creates three offered-output counters in the network processor and three discard counters in the QChip. Use the Offered Profile Capped CIR option when ingress packets are being marked as either in-profile, or out-of-profile when the traffic is in excess of CIR. Packet priority is ignored, and is not separated in the accumulated statistics.
Offered Profile CIR	The system creates four offered-output counters in the network processor and four discard counters in the QChip. The Offered Profile CIR option is similar to the Offered Limited Profile CIR option, except that it includes the in-profile packet classification along with the out-of-profile and undefined-profile classifications. As with the Offered Limited Profile CIR option, packet priority is ignored and not separated in the statistics.

Table 50-2 Stats Mode options (continued)

Option	Description
Offered Profile No CIR	The system creates two offered-output counters in the network processor and two discard counters in the QChip. Use the Offered Profile No CIR option when all ingress packets are either in-profile or out-of-profile. Undefined-profile packets are treated as offered out from a statistics perspective. Undefined-profile packets are affected by the current state of the policer's CIR and are output as either in-profile or out-of-profile, depending on the CIR output state. The offered, discarded, and forwarded statistics do not reflect this behavior because they are based on the initial profile of the packets.
Offered Total CIR	The system creates two offered-output counters in the network processor and two discard counters in the QChip. Use the Offered Total CIR option when ingress priority and initial profile visibility is not required, and CIR profiling is in use. In many cases, all packets offered to a policer are of one priority level and all have the same initial profile (in-profile, out-of-profile, or undefined-profile). This option is different from the Minimal option because it provides visibility into the policer's CIR output.

50.2.7 Traffic mapping

You can specify the mapping between the ingress traffic and the ingress queue. Mapping is optional and can be based on combinations of customer QoS marking (dot1p, DSCP, LSP-EXP, and Precedence), and IP criteria, Ipv6 criteria, or MAC criteria. See the table below.

Adding an LspExp rule to a policy forces packets that match the specified MPLS LSP EXP criteria to override the existing forwarding class and enqueuing priority, based on the parameters specified in the LspExp rule. This functionality allows geographically distributed ISP sites to establish site-to-site interconnection service through a backbone network using VPLS/VLL. Each ISP site PE router connects to a 7x50 Ethernet L2 SAP in the backbone network, and traffic is encapsulated in a VPLS/VLL service tunnel. A maximum of eight LspExp rules are allowed on a single access ingress policy.

Table 50-3 SAP access ingress policy traffic mapping

Tab	Function
Dot1p	Maps the dot1p value of the ingress traffic to the ingress queue ID
DSCP	Maps the DSCP value of the ingress traffic to the ingress queue ID
LspExp	Maps the EXP value of the ingress traffic to the ingress queue ID
Precedence	Maps the precedence value of the ingress traffic to the ingress queue ID
IP Match Criteria	Maps the IP Match Criteria of the ingress traffic to the ingress queue ID
IPv6 Match Criteria	Maps the IPv6 Match Criteria of the ingress traffic to the ingress queue ID

Table 50-3 SAP access ingress policy traffic mapping (continued)

Tab	Function
MAC Match Criteria	Maps the MAC Match Criteria of the ingress traffic to the ingress queue ID

50.3 SAP access egress policies

50.3.1 General Information

SAP access egress policies are applied to access egress interfaces and specify QoS on egress.

SAP access egress policies define egress service queues and map forwarding class flows to queues. In the simplest access egress policy, all forwarding classes are treated like a single flow and mapped to a single queue.

The required SAP access egress policy elements include:

- a unique access egress policy ID
- at least one defined default queue
- an exclusive scope for one-time use, or a template scope for use with multiple SAPs and interfaces

The optional SAP access egress policy elements include:

- additional queues up to a total of 8 separate queues for each of the 8 supported forwarding classes
- IEEE 802.1p priority value remarking based on forwarding class
- a policer to control traffic flow rate
- a dynamic policer to specify traffic flow rate for dynamically-spawned policers
- a policer to specify the high drop tail and high plus drop tail burst size (by percentage) from the specified MBS

Each queue in a policy is associated with one or more of the supported forwarding classes. Each queue can have its individual queue parameters allowing individual rate shaping of the forwarding classes mapped to the queue. More complex service queuing models are supported, where each forwarding class is associated with a dedicated queue.

The forwarding class determination per service egress packet is determined at ingress. If the packet ingresses the service on the same managed device, the service ingress classification rules determine the forwarding class of the packet. If the packet was received over a service tunnel, the forwarding class is marked in the tunnel transport encapsulation.

There is one default SAP access egress policy. The default policy gives all traffic equal priority with the same chance of being sent or dropped during periods of congestion.

50.3.2 Policy override

You can override some or all settings associated with an access egress policy on an L2 or L3 access interface, SLA profile, or subscriber profile. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) for more information.

Note: You can override SAP access egress policies that have the Scope parameter set to template. See 50.30 “To configure a SAP access egress policy” (p. 1550) for more information.

50.3.3 Policers

You can add a policer to an access egress policy to provide traffic flow limiting. Policers are associated with the forwarding classes defined in the access ingress policy. Policers can also be linked to a policer control policy, which maintains a hierarchy of multiple policer objects in the NFM-P system. Policers are linked to policer control policies by means of an arbiter. For more information, see 50.16 “Policer control policies” (p. 1524).

When policers are configured in an access egress policy, they can be used to generate statistics. The type of statistics generated depends on the Stats Mode. Table 50-2, “Stats Mode options” (p. 1510) lists the options for counter allocation when generating statistics.

50.3.4 Traffic mapping

You can specify the mapping between the egress traffic and the egress queue. Mapping is optional and can be based on combinations of customer QoS marking (dot1p, DSCP and precedence), IP/IPv6 criteria, or MAC criteria, as shown in the following table.

Table 50-4 SAP access egress policy traffic mapping

Tab	Function
Dot1p	Maps the dot1p value of the egress traffic to the egress queue ID
DSCP	Maps the DSCP value of the egress traffic to the egress queue ID
Precedence	Maps the precedence value of the egress traffic to the egress queue ID
IP Match Criteria	Maps the IP Match Criteria of the egress traffic to the egress queue ID
IPv6 Match Criteria	Maps the IP Match Criteria of the egress traffic to the egress queue ID

Adding an LspExp rule to a policy forces packets that match the specified MPLS LSP EXP criteria to override the existing forwarding class and enqueuing priority, based on the parameters specified in the LspExp rule. This functionality allows geographically distributed ISP sites to establish site-to-site interconnection service through a backbone network using VPLS/VLL. Each ISP site PE router connects to a 7x50 Ethernet L2 SAP in the backbone network, and traffic is encapsulated in a VPLS/VLL service tunnel. A maximum of eight LspExp rules are allowed on a single access ingress policy.

50.4 ATM QoS policies

50.4.1 General Information

ATM QoS policies are used to specify how ATM traffic is managed, by using ATM traffic descriptors such as service category and shaping. You can create up to 2000 ATM QoS policies per router. ATM QoS policies are assigned to access interfaces.

The following table describes ATM service categories and typical applications:

Table 50-5 Service categories

Service category	Application
CBR	For applications such as voice and video that require high priority and have a known peak transmission rate. CBR guarantees bandwidth for constant bit rate traffic, very low cell loss, and very low delay. For circuit-switched data paths, the Service Category parameter is set to the CBR option and cannot be changed.
RT-VBR	For time-sensitive applications, such as voice and video, that have unpredictable, bursty traffic characteristics. Guarantees very low cell loss and very low delay. The NFM-P handles RT-VBR and NRT-VBR identically for routing and rerouting purposes, and uses both to calculate the total VBR bandwidth usage. You cannot exceed the physical speed of the ATM interface when setting the traffic descriptors.
NRT-VBR	For applications such as video and frame relay that have known or predictable traffic characteristics. Guarantees low cell loss and low delay. You cannot exceed the physical speed of the ATM interface when setting the traffic descriptors.
UBR	For applications that do not require guarantees of low cell loss or low delay. UBR paths emulate the connectionless services provided by conventional bridged and routed data networks.

50.5 Post Policer Mapping policies

50.5.1 General information

You can create a Post Policer Mapping policy to:

- change the access egress packet's forwarding class to redirect the packet to an alternate queue than the ingress forwarding class determination would have used
- change the access egress packet's profile (for example, in or out) to modify the congestion behavior within the egress queue

In both cases, egress marking decisions will be based on the new forwarding class and profile as opposed to the egress forwarding class or profile. The exception is when ingress remarking is configured. An ingress remark decision will not be affected by egress forwarding class or egress profile overrides.

See 50.36 "To configure a Post Policer Mapping policy" (p. 1563) for more information. Once configured, you can apply the Post Policer Mapping policy to a SAP Access Egress Policy; see 50.30 "To configure a SAP access egress policy" (p. 1550).

50.6 MC MLPPP ingress and egress QoS profiles

50.6.1 General Information

MC MLPPP ingress QoS profiles are used to configure the reassembly timeout for each of the four MLPPP classes. You can create up to 128 ingress QoS profiles per NE.

MC MLPPP egress QoS profiles are used to specify queue and queue scheduling parameters for each of the four MLPPP classes.

50.7 MCFR ingress and egress QoS profiles

50.7.1 General Information

MCFR ingress QoS profiles are used to configure the reassembly timeout for each of the four MLFR classes. You can create up to 128 ingress QoS profiles per NE.

MCFR egress QoS profiles are used to specify queue and queue scheduling parameters for each of the four MLFR classes.

50.8 Network policies

50.8.1 General information

Network policies are applied to network interfaces or access uplink ports and specify QoS on egress and ingress.

On ingress, a network policy maps incoming DSCP and EXP values to forwarding class and profile state for traffic received from the core network. On egress, the policy maps forwarding class and profile state to DSCP and EXP values for traffic to be transmitted into the core network.

50.9 Network queue policies

50.9.1 General information

Network queue policies are applied to network ports, uplink ports, or daughter cards.

Network queue policies determine:

- the default burst allocations for queues based on the queue’s forwarding class
- the CIR, PIR, FIR, and burst size parameters for the queue

Note: You cannot use the same network queue policy on devices at different releases.

For network egress, a network burst policy is associated with the network port buffer pool. For network ingress, the network burst policy is associated the network ingress buffer pool of the daughter card.

Network queue policies support multipoint queues and a variable number of forwarding classes. When you want to deploy a network queue policy to devices of different types, you may need to create an instance of the policy for each NE type. Policies that are configured for multipoint queues

and a variable number of forwarding classes are deployed only to devices that support the functionality. Modification of an existing policy results in automatic deployment to participating NEs, so must be done with consideration of the policy features and devices involved.

CIR, PIR, and FIR

CIR defines the rate that packets are considered in-profile by the system. In-profile packets are preferentially queued by the system at egress and at subsequent next-hop devices on which the packet can traverse. To be properly handled as in- or out-of-profile throughout the network, the packets must be marked accordingly for profiling at each hop.

PIR specifies the peak information rate for an ingress or egress daughter card queue to transmit packets through the switch fabric or out an egress interface. On ingress, the PIR defines the maximum rate at which the queue can transmit packets through the switch fabric. On egress, the PIR defines the maximum rate at which the queue can transmit packets out an egress interface. The specified PIR value does not guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or the available bandwidth.

FIR configuration allows the scheduling priority to be modified independently from the marking/drop precedence of the packets being scheduled from the queue. Specifically, a higher queue scheduling priority can be assigned without modifying the queue's CIR.

Burst size

The burst size for a queue determines whether the queue has exhausted its reserved buffer pool space while queuing packets. When the queue has exceeded the amount of buffer pool space considered in reserve for the queue, it must contend with other queues for the available shared buffer space in the buffer pool. Access to this shared pool space is controlled by the RED slope.

Two RED slopes are maintained in each buffer pool. A high-priority slope is used by in-profile packets. A low-priority slope is used by out-of-profile packets. All nc forwarding class packets are considered in-profile. Assured packets are handled by their in-profile and out-of-profile markings. All be packets are considered out-of-profile.

Premium queues should be configured such that the committed burst size is sufficient to prevent shared buffering of packets. This is handled by the CIR scheduling of premium queues and the overall small amount of traffic on the class. Premium queues in a properly designed system drain before all others, limiting their buffer utilization. The RED slopes detect congestion conditions and work to discard packets and slow down random TCP session flows through the queue.

The resultant committed burst size can be larger than the maximum burst size. This results in a portion of the committed burst size for the queue unused and should be avoided.

50.10 Shared-queue policies

50.10.1 General Information

A shared-queue policy can be applied to daughter cards for optional use by SAPs. The NFM-P provides one shared-queue policy for all eight default queues. Each queue is associated with a default forwarding class.

Shared-queue QoS policies can be implemented to facilitate queue consumption on a daughter card; for example, when VPLS, IES, and VPRN services are scaled on one card. Instead of allocating multiple hardware queues for each unicast queue that is defined in an access ingress policy, SAPs with the shared-queuing feature enabled only allocate one hardware queue for each unicast queue.

However, the total amount of traffic throughput at ingress is reduced because ingress packets that are serviced by a shared-queuing SAP are recirculated for additional processing. This can reduce the available bandwidth by half. Shared queuing can also add latency. Network planners should consider these restrictions when they try to scale services on one daughter card.

The following table lists the queue IDs used by the NFM-P to identify the shared-queue types.

Table 50-6 Shared queue types

Shared-queue ID	Shared-queue type
1 to 8	Unicast
9 to 16	Multicast
17 to 25	Broadcast
26 to 32	Unknown

Note: Queue IDs 9 to 32 are also known as multipoint shared queues.

50.10.2 Shared policer output queue

To support hierarchical policing, a default policer output queue policy is applied automatically to each IOM that supports ingress policing. In the NFM-P, this shared policer output queue is modeled in the same manner as the existing default shared queue, except that it has only 16 queues.

50.10.3 Multipoint shared-queue policies

Multipoint shared queues minimize the number of multipoint queues that are created for the following:

- ingress VPLS, IES, or VPRN SAPs
- ingress subscriber SLA profiles

Typically, ingress multipoint packets are handled by multipoint queues that are created for each SAP or subscriber SLA profile instance. In some cases, the number of SAPs or SLA profile instances are sufficient for the in-use multipoint queues to represent many thousands of queues on an ingress forwarding plane. If multipoint shared queuing is enabled for the SAPs or SLA profile instances on the forwarding plane, the multipoint queues are not created. Instead, the ingress multipoint packets are handled by the unicast queue that is mapped to the forwarding class of the multipoint packet.

Multipoint shared queues are a superset of shared queuing. With shared queuing on a SAP or SLA profile instance, only unicast packets are processed twice; once for the initial service-level queuing and a second time for switch fabric destination queuing. Shared queuing does not affect multipoint

packet handling. Multipoint packet handling in normal service queuing is the same as shared queuing. Shared queuing for unicast packets is automatically enabled when you enable multipoint shared queuing.

The NFM-P supports multipoint shared-queue policies for the following services on the 7750 SR, 7450 ESS:

- Layer 2: VPLS and MVPLS
- Layer 3: IES and VPRN

See the following chapters for more information about enabling multipoint shared queues for a service:

- Chapter 77, “VPLS management ” for VPLS and MVPLS
- Chapter 78, “IES management” for IES
- Chapter 79, “VPRN service management” for VPRN

You can also enable multipoint shared queues in an existing shared-queue policy. See the procedures in this chapter for more information about editing shared-queue policies.

50.11 Slope policies

50.11.1 General Information

Slope policies are applied to access ports, network ports, and network daughter cards.

Slope policies define weighted RED (WRED) slope characteristics for buffer pools. Low-priority and high-priority slopes are specified when you configure a slope policy. If a slope policy is not explicitly specified, a default policy is applied.

Each buffer pool supports a high-priority, a low-priority RED slope, an Exceed Slope, and a HighPlus Slope. The high-priority RED slope manages access to the shared portion of the buffer pool for the high-priority or in-profile packets. The low-priority RED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets. The Exceed slope handles the exceeded profile traffic congestion control at the pool and megapool level. The Exceed slope is only used for exceed profile traffic when it is not shutdown. The defaults for the Exceed slope parameters should be set so that they are lower than the low slope in the slope hierarchy. Finally, the HighPlus Slope handles the inplus profile traffic congestion control. A HighPlus Slope policy can be applied to pool-per-queue queues in an egress queue group template, which in turn can be applied at a network egress port. The HighPlus Slope is also supported in a slope policy applied to the egress WRED megapool. The values for the HighPlus Slope parameters should be set so that they are higher than the High Slope in the slope hierarchy. By default, the Low, High, Exceed, and HighPlus slopes are initially disabled.

A RED slope is a graph with an X (horizontal) and Y (vertical) axis. The X axis plots the percentage of shared buffer utilization, from 0 to 100%. The Y axis plots the probability of packet discard marked from 0 to 1. The slope is defined as four sections, as shown in the following figure.

Figure 50-1 RED slope characteristics

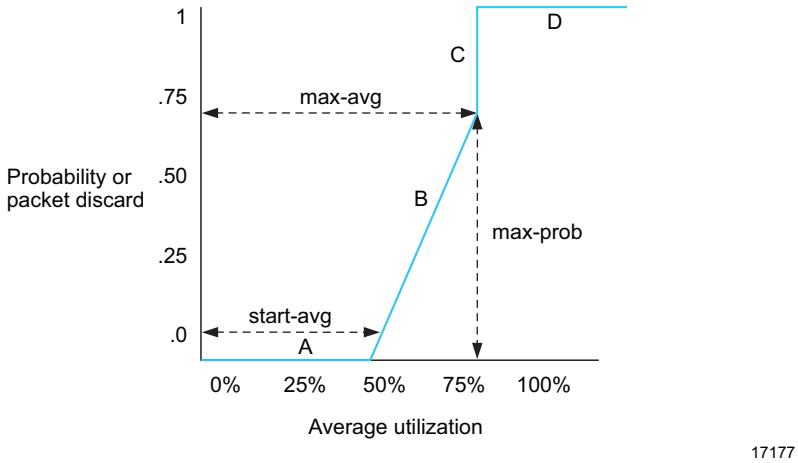

Section A is (0, 0) to (start-avg, 0). For this part of the slope, the packet discard value is always zero, which prevents the RED function from discarding packets when the shared buffer average utilization falls between 1 and start-avg.

Section B is (start-avg, 0) to (max-avg, max-prob). This part of the slope is linear where packet discard probability increases from zero to max-prob.

Section C is (max-avg, max-prob) to (max-avg, 1). This part of the slope shows the instantaneous increase of packet discard probability from max-prob to one. A packet discard probability of one results in an automatic discard of the packet.

Section D is (max-avg, 1) to (100%, 1). On this part of the slope, the shared buffer average utilization value of max-avg to 100% results in a packet discard probability of one.

50.12 HSMDA WRED slope policies

50.12.1 General Information

HSMDA WRED slope policies controls the management of the HSMDA queue depth. The polices are applied to queues defined in the SAP ingress and SAP egress QoS to provide congestion control. Congestion control includes a defined maximum depth that the queue can reach when packets and RED of congestion and slope-based discards based on queue depth are accepted.

Note: HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier. The NFM-P does not support HSMDA WRED slope policies on other NEs.

50.13 Scheduler policies

50.13.1 General Information

Scheduler policies determine the order in which queues are serviced. All ingress and egress queues operate within the context of a scheduler. Multiple queues share the same scheduler. Schedulers control the data transfer between the following queues and destinations:

- service ingress queues to switch fabric destinations
- service egress queues to access egress ports
- network ingress queues to switch fabric destinations
- network egress queues to network egress interfaces

The are two types of scheduler policies:

- single-tier, in which queues are scheduled based on the forwarding class of the queue and the operation state of the queue relative to the queue CIR and PIR
- hierarchical or multi-tier, which allow the creation of a hierarchy of schedules where queues or other schedulers are scheduled by superior schedulers

Scheduler policies are applied to access ingress and access egress interfaces.

50.13.2 Single tier schedulers

Single-tier scheduling is the default method of scheduling queues. Queues are scheduled with single-tier scheduler policies when no explicit hierarchical scheduler policy is defined or applied. In single-tier scheduling, queues are scheduled based on the forwarding class of the queue and the operational state of the queue relative to the queue CIR and PIR.

50.13.3 Hierarchical schedulers

Hierarchical scheduler policies are used for access ingress and access egress queues. Hierarchical scheduler policies allow you to create a hierarchy of schedulers where queues and other schedulers are scheduled by superior schedulers.

Virtual schedulers are created within the context of a hierarchical scheduler policy. A hierarchical scheduler policy defines the hierarchy and parameters for each scheduler. A scheduler is defined in the context of a tier. The tier level determines the scheduler's position in the hierarchy. Three tiers of virtual schedulers are supported, as shown in the following figure:

Figure 50-2 Hierarchical scheduler and queue association

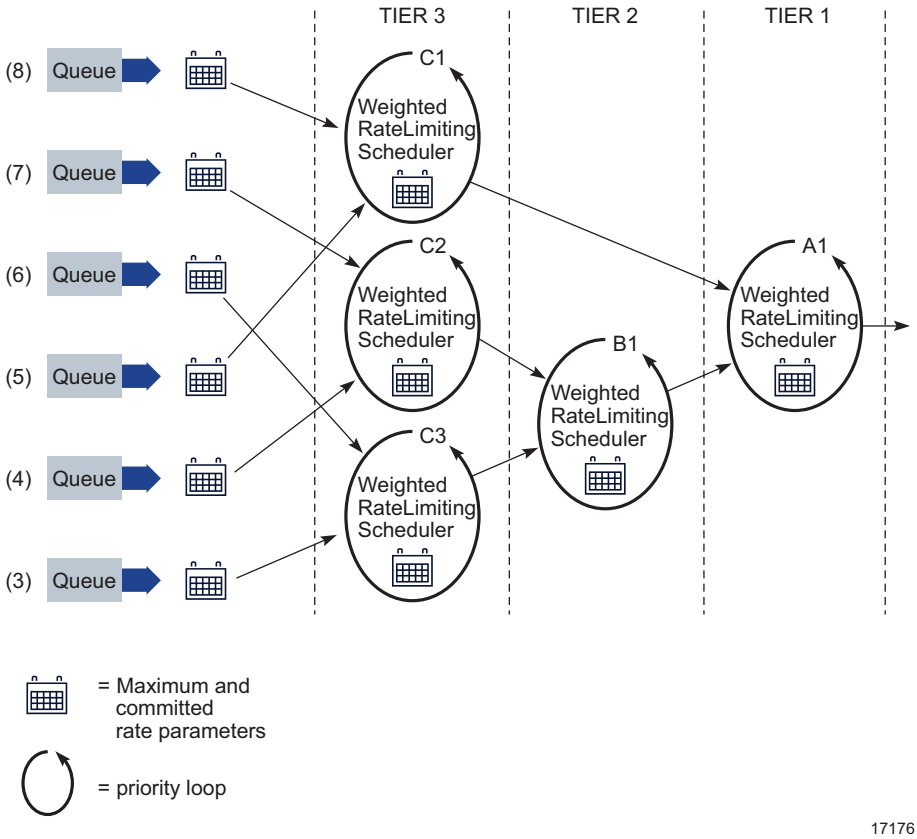

Tier 1 schedulers are defined without a parent scheduler. A scheduler can enforce a maximum rate of operation for all child queues and associated schedulers.

You can create a tier 2 scheduler without a parent tier 1 scheduler or a tier 3 scheduler without a parent tier 2 scheduler.

During configuration of the hierarchical scheduler, you can assign numerical level and weight values to the tier 2 or tier 3 child scheduler, to specify a priority level in comparison with other child schedulers of the same parent scheduler. The level value determines relative importance when tier 2 or tier 3 schedulers are contending for bandwidth. The higher the number, the higher the priority level of the child scheduler bandwidth request.

Child schedulers with a level value lower than other child schedulers do not receive bandwidth until all child schedulers with a higher level have reached their maximum bandwidth allocation, or have no packets to pass.

When two child schedulers have the same level value, the weight value determines which scheduler first receives bandwidth.

50.13.4 Policy override

You can override some or all settings associated with an ingress or egress scheduler policy on an L2 or L3 access interface or subscriber profile. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) for more information.

50.14 Port scheduler policies

50.14.1 General Information

Port scheduler policies determine the virtual scheduling of egress ports by allocating HQoS bandwidth based on the available bandwidth at the egress port level. A port scheduler is defined in the context of a tier. The tier level determines the position of the port scheduler in the hierarchy.

The first tier of the scheduling hierarchy manages the total frame bandwidth that the port scheduler allocates to the eight priority levels. The second tier receives bandwidth from the first tier in two priorities—a within-CIR distribution, and an above-CIR distribution. The within-CIR distribution of the second tier provides bandwidth to the third tier within-CIR distributions for each of the eight priority levels. The above-CIR distribution of the second tier provides bandwidth to the above-CIR distribution of the third tier for each of the eight priority levels.

Up to eight groups can be defined within each port scheduler policy. A group has a rate, an optional cir-rate, and inherits the highest scheduling priority of its member levels. In essence, a group receives bandwidth from the port and distributes it within the member levels of the group according to the weight of each level. Each level will compete for bandwidth within the group based on its weight under congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth. One or more levels can map to the same group.

Note: When the agg-rate-limit is applied to queues of a subscriber which are mapped to different priority levels in the same weighted scheduler group, the bandwidth distribution to the queues will be based on the priority of the level.

Orphan queues or schedulers that are not explicitly associated with the port scheduler receive bandwidth after all parented queues and schedulers are allocated bandwidth.

Port scheduler policies are configured on ports and channels.

50.14.2 Policy override

You can override some or all port scheduler settings associated with an access egress queue policy on an L2 or L3 access interface or SLA profile. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) for more information.

50.15 HSMDA scheduler policies

50.15.1 General Information

HSMDA scheduler policies controls the scheduling of a set of HSMDA scheduler classes. The policies are assigned to egress HSMDA ports, and the ingress control scheduler between the HSMDA and ingress forwarding plane. The policies assigned to HSMDA egress ports define how all queues associated with the egress port are scheduled. Scheduler policies assigned to the ingress

control scheduler between the HSMDA and ingress forwarding plane define how all ingress queues on the HSMDA, regardless of the ingress port, are scheduled.

Note: HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier. The NFM-P does not support HSMDA scheduler policies on other NEs.

50.16 Policer control policies

50.16.1 General Information

Policer control policies allow you to create a control framework, under which policer objects associated with SAPs or subscriber contexts are configured with traffic control parameters. The policer object is applied to the SAP or subscriber context as part of an access ingress policy or access egress policy. The policer control policy is also applied to the SAP or subscriber context. The policer control policy provides the control framework, and policers are associated with the framework by means of arbiters.

When a policer control policy is applied to a SAP or subscriber context, the system creates a parent policer that is bandwidth limited by the policy's maximum bandwidth rate, as defined under the root arbiter. In addition to the root arbiter, the policy may also contain user defined child arbiters that provide bandwidth control for subsets of child policers.

Figure 50-3 Policer association with SAPs and user contexts under policer control hierarchy

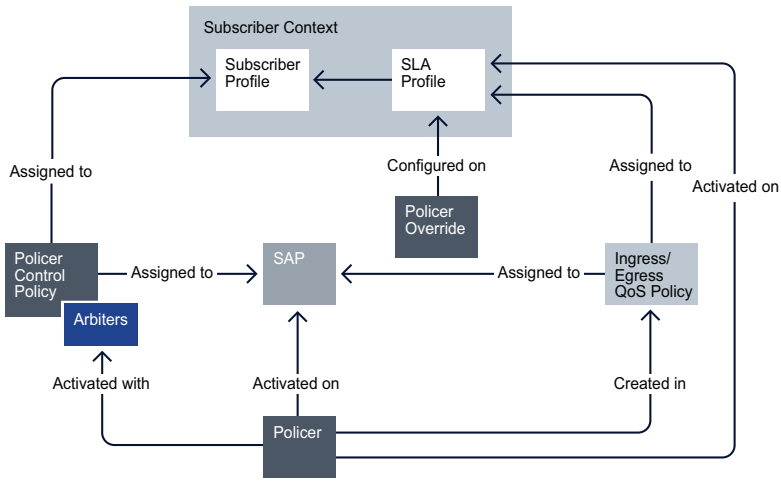

A policer is created as part of an access ingress or access egress policy. The policer is configured with a parent arbiter name. When a policy containing a policer is applied to a SAP, the system scans the available arbiters on the SAP. If an arbiter exists with a name that matches the Parent Arbiter parameter on the policer, a policer to arbiter association is established and the policer becomes part of the policer control hierarchy.

In the case of subscriber contexts, the policer control policy is applied to the sub profile used by the subscriber. The system creates a unique policer control hierarchy for each subscriber associated

with the sub profile. An access ingress or access egress policy containing a policer is applied to the subscriber SLA profile. The combination of the sub profile and SLA profile at the subscriber level provides the system with the information required to create the policer control hierarchy instance for the subscriber context.

50.26 “Workflow to configure policer control hierarchy” (p. 1535) describes the high-level steps to configure a policer control hierarchy.

50.17 Buffer pool policies

50.17.1 General Information

There are two types of buffer pool policies:

- **HSMDA policies:** The high scale Ethernet MDA (HSMDA) pool policy can be assigned to an ingress or egress HSMDA to control how buffers are distributed between HSMDA queues. HSMDA pool policies do not apply to ports.

  Note: HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier. The NFM-P does not support HSMDA pool policies on other NEs.

- **Named buffer pool policies:** Named buffer pool policies allows you to create named pools to override the default buffer pool behavior by creating and allocating ingress and egress queues. Named pools are configured in the named buffer pool policy and are applied at the MDA and/or port ingress and egress level.

  Named pools can be further defined in Q1 pools configuration. You can configure and assign queue pools to the following policies:

  - SAP Access Ingress
  - SAP Access Egress
  - Network Queue
  - Shared Queue

  When a policy is associated at the MDA level, named pools defined in the policy allow queues from any port to be associated. When a policy is associated at the port level, named pools are only available to queues associated with that port. A named pool policy that is currently applied to a MDA or port can be deleted after all association between the policy and the MDA or port have been removed.

  When buffer pools are created, renamed or deleted, queues mapped to the pools are moved to default pools. When a queue is moved, traffic that was destined for the queue is temporarily moved to a fail over queue. After the old queue is drained, and the new queue is created and associated with the buffer pool, the saved statistics are loaded to the new queue and traffic is moved from the fail over queue to the new queue.

  50.27 “Workflow to configure a named buffer pool” (p. 1537) describes the high-level steps to configure Named Buffer Pools and Q1 pools.

50.18 HS QoS policies

50.18.1 General Information

HS QoS policies allow you to configure and manage ESM and services using the HSQ IOM4-e-HSMDA on the 7750 SR, Release 20.10 and earlier. There are four types of HS QoS policies:

- **HS Attachment policy:** The high scale (HS) Attachment policy defines how queues map to the scheduling classes associated with the HS port scheduler policy. The HS Attachment policy also defines how multiple queues collapse into WRR (weighted round robin) groups allowing multiple queues to map to a single scheduling class. See 50.69 “To configure a HS Attachment policy” (p. 1612).
- **HS Pool policy:** The high scale (HS) Pool policy is used to control how buffers are distributed to root-tier and mid-tier buffer pools and how mid-tier buffer pools attach to root-tier buffer pools. The HS pool policy is applied to the Forwarding Plane at the IOM level on a 7750 SR node. See 50.70 “To configure a HS Pool policy” (p. 1613).
- **HS Port Pool policy:** The high scale (HS) Port Pool policy is used to control how buffers are distributed to port-class buffer pools and how port-class buffer pools attach to mid-tier buffer pools. The HS port pool policy is applied to an HSQ (IOM4-e-HSMDA) egress port on a 7750 SR, Release 20.10 and earlier. See 50.71 “To configure a HS Port Pool policy” (p. 1614).
- **HS Scheduler policy:** The high scale (HS) Scheduler policy defines the scheduling behavior of each HS scheduler class for all queues associated to an HSQ (IOM4-e-HSMDA) egress port on a 7750 SR, Release 20.10 and earlier, including specifying the Max. rate, scheduling class type (rate/group), bandwidth rate, and weight-in-group. See 50.73 “To configure a HS Scheduler policy” (p. 1616).

50.19 Queue Group policies

50.19.1 General Information

There are three types of Queue Group policies:

- **Queue Group Ingress Template policies:** are used for Access Ingress Queue Group creation on Ethernet access ports; see 50.74 “To configure a queue group ingress template policy” (p. 1617).
- **Queue Group Egress Template policies:** are used for Access Egress Queue Group and Network Egress Queue Group creations on Ethernet ports; see 50.75 “To configure a queue group egress template policy” (p. 1619).

  Queue Group template policies allow you to define the queuing and parenting structure for queue groups on Ethernet ports. The policy defines the number and types of queues within the port queue group, and provides the default queue parameters. Queue Group Template policies are not applicable to L3 interfaces associated with HSMDA ports.

  Queue Group Template policies are used in the following network applications:
  - Access SAP queue group applications
  - Network port queue groups for network interfaces
- **Queue Group Redirect List policies:** are used to provide QoS control to map IPv4 and IPv6 traffic using a single VXLAN or VXLAN GPE Virtual Network Identifier (VNI) to an ingress access

queue group instance or a port access egress queue group instance. Each entry in the redirect list policy matches to a specific VNI which then maps it to a specific queue group instance. A maximum of 16 match statements can be configured in a queue group redirect list. Queue Group Redirect List policies are supported on both egress/ingress physical interfaces (SAPs) that support IES or VPRN services. See 50.76 “To configure a queue group redirect list policy” (p. 1624).

See Chapter 71, “Queue groups” for more information about queue groups, the associated network components, and typical applications.

50.19.2 Default policer queue group

The system maintains a special default egress queue group template policy that is applied automatically to all Ethernet ports. The default policer-output-queues policy is configured with two queues:

- Queue 1: configured with a forwarding class value of Best Effort.
- Queue 2: configured with a forwarding class value of Expedite.

All other parameters are default. You cannot delete the default policer-output-queues policy.

50.20 FP Resource policies

50.20.1 General information

Forwarding Plane Resource policies allow custom allocation of card slot resources to queues. By default, ingress and egress queues are each allocated 50% of the queue resources on a forwarding plane. FP Resource policies allow you to configure custom ingress queue resources in a range of 4% to 97%. The corresponding percentage remaining is allocated to egress queues. See 50.72 “To configure an FP Resource policy” (p. 1615).

FP Resource policies are assigned to forwarding planes on the card slot properties form; see 15.59 “To assign an FP Resource policy to a forwarding plane” (p. 519).

50.21 Queue depth monitoring

50.21.1 General information

Queue depth monitoring provides insight into traffic flows on queues, which can help to fine-tune port service allocations and QoS policies, for capacity planning and SLA compliance.

The NFM-P supports the configuration of queue depth monitoring using overrides for access ingress, access egress, and network egress queue groups, and for L2 and L3 access egress interfaces.

- For egress queue group overrides on Ethernet ports, see 16.37 “To add a queue group to an Ethernet port” (p. 627).
- For QoS policy overrides for L2 and L3 access egress interfaces, see 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654).

Configuration includes options for fast polling, which can provide more accurate results, and for violation thresholds, which generate notifications when exceeded. Depth monitoring results and notifications are not available in the NFM-P, but can be obtained from NEs using CLI or telemetry, depending on NE version and release support. See the NE documentation for more information.

50.22 7705 SAR fabric profiles

50.22.1 General Information

Each daughter card in a 7705 SAR can have two assigned fabric profiles, one for access ingress and one for network ingress. These policies are assigned from the Daughter Card tab of a daughter card slot properties form.

A 7705 SAR fabric profile includes a mode, which cannot be changed after profile creation, and one or more shaping rates. The mode types are Aggregate and Destination. In Aggregate mode, one rate defines the maximum fabric shaping rate that is distributed to each daughter card slot. In Destination mode, there is one fabric shaping rate for each daughter card slot.

In Destination mode, there is also a multi-point shaping rate that is distributed to each daughter card slot.

50.23 7210 SAS QoS policies

50.23.1 General Information

The NFM-P provides QoS policies that are specific to 7210 SAS NEs. Some of these policies also apply to 7250 IXR devices, as indicated by the menu name of the policy. QoS policy information that is specific to the 7210 SAS may not necessarily apply for other devices.

QoS policies for the 7210 SAS accommodate multiple different 7210 SAS chassis types. See 6.5 “7210 SAS” (p. 216) in Chapter 6, “Device support” for more information about 7210 SAS chassis types. Not all chassis types support all policies, or all configurations within a policy. For information about chassis support for a specific policy, see Table 50-7, “7210 SAS QoS policies” (p. 1531).

For some 7210 SAS QoS policies and functions, device system resources must be appropriately allocated. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” and the NE documentation for more information.

Some 7210 SAS QoS functions and settings are not configured within the policy framework. Instead, they are configured as device or port properties, or during SAP configuration. For an overview of the tasks required to configure QoS on 7210 SAS NEs using the NFM-P, see 70.15 “Sample QoS configuration on the 7210 SAS” (p. 1954) .

You can override the settings for meters and queues in access ingress policies that are assigned to access interfaces. For information about policy overrides for access ingress meters on 7210 SAS NEs, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657). For information about policy overrides for access ingress queues on 7210 SAS-X NEs, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659).

On supporting 7210 SAS NEs, you can override the settings for queues in port access egress policies that are assigned to access ports; see 50.100 “To configure QoS policy overrides on port access egress queues for a 7210 SAS” (p. 1660).

A GQP can be used to assign 7210 and 1830 SAP access ingress policies and 7210 SAP access egress policies to access interfaces; see 50.96 “To configure a Generic QoS Profile” (p. 1649). GQPs are assigned using the XML API.

When a GQP is used to assign ingress and egress policies to 7210 SAS NEs that do not support QoS policy overrides, an API may be used to provide GQP overrides that change QoS values (CIR, PIR, CBS, MBS). The API creates a new “dynamic” local policy that is outside the NFM-P policy management. Local policy forms remain visible in the NFM-P GUI, and display the Policy Mode parameter, set to dynamic. This indicates that the QoS values are set by the GQP override API, and not from the GUI, CLI, or OSS.

See the 7210 SAS documentation for more information about 7210 SAS QoS and QoS policies.

50.23.2 Table-based ingress classification on the 7210 SAS

The NFM-P supports the configuration of table-based ingress classification on supporting 7210 SAS NEs. When table classification is enabled, ingress DSCP bits are mapped to FCs and profiles using hardware tables instead of CAM resources. The DSCP mappings are configured in a 7210/7250 DSCP classification policy.

You can enable table-based ingress classification for the following:

- L2 interfaces on VPLS and VLL Epipe services. The 7210/7250 DSCP classification policy is assigned to the SAP Access Ingress policy on the interface.
- L3 interfaces on IES and VPRN services. The 7210/7250 DSCP classification policy is assigned to the SAP Access Ingress policy on the interface.
- RVPLS on L3 interfaces. The 7210/7250 DSCP classification policy is assigned to the interface using the Routed VPLS tab of the properties form for the interface. On the form, the 7210/7250 DSCP classification policy is called the Routed Override QoS Policy.
- Ethernet ports. The 7210/7250 DSCP classification policy is assigned to the port using the Policies tab of the port properties form. When table-based classification is enabled on a port, all SAPs on the port use the 7210/7250 DSCP classification policy assigned to that port. You can configure a default FC for ingress packets that do not have a DSCP value.

See the NE documentation for more information about table-based classification.

50.23.3 Workflow to configure table-based DSCP ingress classification

The following workflow outlines the basic steps required to configure table-based ingress DSCP classification.

1

Configure a 7210/7250 DSCP classification policy; see 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638). Distribute the policy to the required NEs.

2

If you are configuring table-based ingress classification on interfaces, assign the 7210/7250 DSCP classification policy to a 7210, 7250 and 1830 SAP Access Ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544). Distribute the policy to the required NEs.

3

Configure table-based ingress classification on interfaces, and assign the 7210, 7250 and 1830 SAP Access Ingress policy from Stage 2 to those interfaces.

See the relevant procedure for the service type:

- VLL: 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184)
- VPLS: 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343)
- IES: 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479)
- VPRN: 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665)

4

Configure table-based ingress classification for RVPLS on L3 interfaces, and assign the 7210/7250 DSCP classification policy from Stage 1 to those interfaces.

See the relevant procedure for the service type:

- IES: 78.30 “To bind an IES L3 access interface to a VPLS site or VPLS I-site” (p. 2475)
- VPRN: 79.98 “To bind a VPRN L3 access interface to a VPLS site or VPLS I-site” (p. 2673)

5

Configure table-based ingress classification for Ethernet ports, and assign the 7210/7250 DSCP classification policy from Stage 1 to those ports. See 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636).

50.23.4 7210 SAS QoS policies

The following table describes the QoS policies specific to the 7210 SAS.

Table 50-7 7210 SAS QoS policies

QoS policy	Assigned to (in Procedure #)	Supported 7210 SAS chassis types	For policy configuration, see
7210, 7250 and 1830 SAP Access Ingress	L2 access interfaces in: • VLL (76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184)) • VPLS (77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343)) L3 access interfaces in: • IES (78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479)) • VPRN (79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665))	All 7210 SAS chassis types	50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
7210 and 1830 Port Access Egress	Access ports (16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636))	7210 SAS-D 7210 SAS-Dxp 7210 SAS-E 7210 SAS-M 7210 SAS-Mxp 7210 SAS-S 7210 SAS-Sx 7210 SAS-T 7210 SAS-R 7210 SAS-X	50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)
7210 SAP Access Egress	L2 access interfaces in: • VLL (76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184)) • VPLS (77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343)) • Mirror Service destinations (93.10 “To create an L2 access interface on a destination site” (p. 3174)) L3 access interfaces in: • IES (78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479)) • VPRN (79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665))	7210 SAS-K 7210 SAS-Mxp 7210 SAS-R 7210 SAS-X	50.32 “To configure a 7210 SAP access egress policy” (p. 1558)

Table 50-7 7210 SAS QoS policies (continued)

QoS policy	Assigned to (in Procedure #)	Supported 7210 SAS chassis types	For policy configuration, see
7210 and 1830 Network	Network ports (16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636)) Network interfaces[1] (27.17 “To create an L3 network interface on a routing instance” (p. 856) , 27.18 “To configure L3 network interfaces” (p. 863))	All 7210 SAS chassis types	50.42 “To configure a 7210 and 1830 network policy” (p. 1571)
7210 and 1830 Network Queue	Network ports (16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636))	All 7210 SAS chassis types	50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)
7210 and 1830 Slope	Egress ports in any mode (16.24 “To configure Ethernet ports” (p. 599)) Queues, during queue configuration in: • 7210 and 1830 SAP Access Ingress policies (50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)) • 7210 SAP Access Egress policies (50.32 “To configure a 7210 SAP access egress policy” (p. 1558)) • 7210 and 1830 Network Queue policies (50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)) • 7210 and 1830 Network policies (50.42 “To configure a 7210 and 1830 network policy” (p. 1571))	7210 SAS-D 7210 SAS-Dxp 7210 SAS-E 7210 SAS-K 7210 SAS-M 7210 SAS-S 7210 SAS-Sx 7210 SAS-T	50.50 “To configure a 7210 and 1830 slope policy” (p. 1589)
7210 and 7250 Queue Management	Queues, during queue configuration in: • 7210 and 1830 SAP Access Ingress policies (50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)) • 7210 and 1830 Port Access Egress policies (50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)) • 7210 SAP Access Egress policies (50.32 “To configure a 7210 SAP access egress policy” (p. 1558)) • 7210 and 1830 Network Queue policies (50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583))	7210 SAS-Mxp 7210 SAS-R 7210 SAS-X	50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592)
7210, 7250, and 1830 Port Scheduler[2]	Egress ports in any mode (16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636))	7210 SAS-D 7210 SAS-Dxp 7210 SAS-E 7210 SAS-M 7210 SAS-S 7210 SAS-Sx 7210 SAS-T	50.60 “To configure a 7210, 7250 and 1830 Port Scheduler policy” (p. 1601)

Table 50-7 7210 SAS QoS policies (continued)

QoS policy	Assigned to (in Procedure #)	Supported 7210 SAS chassis types	For policy configuration, see
7210/7250 Dot1p	Ingress policies: • 7210 and 1830 SAP Access Ingress policies (50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)) • 7210 and 1830 Network policies (50.42 “To configure a 7210 and 1830 network policy” (p. 1571))	7210 SAS-K	50.84 “To configure a 7210/7250 Dot1p classification policy” (p. 1637)
7210/7250 DSCP	Ingress policies: • 7210 and 1830 SAP Access Ingress policies (50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)) • 7210 and 1830 Network policies (50.42 “To configure a 7210 and 1830 network policy” (p. 1571)) Ethernet ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) RVPLS on L3 access interfaces: • 78.30 “To bind an IES L3 access interface to a VPLS site or VPLS I-site” (p. 2475) • 79.98 “To bind a VPRN L3 access interface to a VPLS site or VPLS I-site” (p. 2673)	7210 SAS-K 7210 SAS-Mxp 7210 SAS-R	50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638)
7210/7250 MPLS LSP-EXP	7210 and 1830 Network policies (50.42 “To configure a 7210 and 1830 network policy” (p. 1571))	7210 SAS-K12 7210 SAS-K30 ETR	50.86 “To configure a 7210/7250 MPLS LSP-EXP classification policy” (p. 1639)
7210 FC Meter Map	7210,7250 and 1830 SAP Access Ingress policies (50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544))	7210 SAS-Mxp	50.87 “To configure a 7210 FC Meter Map policy” (p. 1640)
7210 Remarking	Egress policies: • 7210 and 1830 Port Access Egress policies (50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)) • 7210 SAP Access Egress policies (50.32 “To configure a 7210 SAP access egress policy” (p. 1558)) • 7210 and 1830 Network policies (50.42 “To configure a 7210 and 1830 network policy” (p. 1571))	7210 SAS-K 7210 SAS-Mxp 7210 SAS-R 7210 SAS-S 7210 SAS-Sx 7210 SAS-T 7210 SAS-X	50.80 “To configure a 7210 remarking policy” (p. 1630)

Table 50-7 7210 SAS QoS policies (continued)

QoS policy	Assigned to (in Procedure #)	Supported 7210 SAS chassis types	For policy configuration, see
7210 MPLS LSP-EXP Mapping	7210 and 1830 Network policies of network interface type (50.42 “To configure a 7210 and 1830 network policy” (p. 1571)) NEs, using the NE properties form	7210 SAS-M 7210 SAS-Mxp 7210 SAS-R 7210 SAS-S 7210 SAS-Sx 7210 SAS-T 7210 SAS-X	50.83 “To configure a 7210 MPLS LSP-EXP Mapping policy” (p. 1635)
7210 Multipoint Bandwidth management	7210 SAS-X NEs, using the NE global properties (52.16 “To configure a 7210 multipoint bandwidth management policy” (p. 1725))	7210 SAS-X	52.16 “To configure a 7210 multipoint bandwidth management policy” (p. 1725)

Notes:

1. The 7210 SAS-M (in network mode), 7210 SAS-Mxp, 7210 SAS-R,7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T (in network mode), and 7210 SAS-X support 7210 and 1830 network policies on network interfaces. The policy must be of network interface type.
2. The 7210 SAS-K, 7210 SAS-Mxp, 7210 SAS-R and 7210 SAS-X do not use scheduler policies. Egress scheduling for these devices is defined by the port bandwidth and by user-configured settings in access egress and network queue policies, port properties, and SAP configuration. See 70.15 “Sample QoS configuration on the 7210 SAS” (p. 1954) in Chapter 70, “Service management and QoS”.

50.24 7250 IXR QoS policies

50.24.1 7250 IXR QoS policy support

If you are creating or synchronizing QoS policies for the 7250 IXR with SR OS, the following NFM-P QoS policies apply:

- SAP Access Ingress; see 50.28 “To configure a SAP access ingress policy” (p. 1538)
- Shared policer; see 50.101 “To configure a shared policer policy” (p. 1661)
- 7250 SROS Queue Management; see 50.53 “To configure a 7250 SROS Queue Management policy” (p. 1593)
- 7250 SROS VLAN QoS; see 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594)
- 7250 SROS Port QoS; see 50.61 “To configure a 7250 SROS Port QoS policy” (p. 1603)
- 7250 SROS Remarking; see 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634)
- 7250 SROS Network Ingress; see 50.43 “To configure a 7250 SROS Network Ingress policy” (p. 1575)

- 7250 SROS Ingress Classification; see 50.45 “To configure a 7250 SROS Ingress Classification policy” (p. 1579)

For supporting chassis variants and releases of the 7250 IXR, the NFM-P supports the following SROS FC mapping policies for ingress classification:

- 7250 SROS Dot1p FC Mapping
- 7250 SROS DSCP FC Mapping
- 7250 SROS LSP-EXP FC Mapping

See 50.44 “To configure FC mapping policies for ingress classification on the 7250 IXR” (p. 1577).

For supporting chassis variants and releases of the 7250 IXR, the NFM-P supports the following SROS FC mapping policies for egress remarking:

- 7250 SROS FC Dot1p Mapping
- 7250 SROS FC DSCP Mapping
- 7250 SROS FC LSP-EXP Mapping

See 50.81 “To configure FC mapping policies for egress remarking on the 7250 IXR” (p. 1632).

For more information about FC mapping policy support on the 7250 IXR, see the NE documentation.

50.25 OmniSwitch QoS policies

50.25.1 General Information

OmniSwitch QoS provides a way to control traffic flows through the switch based on configured policies. The control may be simple such as allowing or denying traffic, or complicated such as remapping 802.1p bits from a Layer 2 network to ToS values in a Layer 3 network.

The OmniSwitch supports the following types of QoS policies:

- basic QoS—to control traffic prioritization and bandwidth shaping
- ACLs—for Layer 2, 3, and 4 filtering

A QoS policy contains a condition and an action. The condition specifies parameters that the switch checks in incoming traffic flows, such as the destination address or ToS bits. The action specifies what the switch does with traffic that matches the condition; for example, the switch may queue the traffic flow with a higher priority or reset the ToS bits.

Ethernet service policies are used to create UNI and SAP profiles. The profiles are applied to stacked VLAN SAPs and UNIs to control traffic and specify options to manage certain traffic types.

50.26 Workflow to configure policer control hierarchy

50.26.1 Requirements

This workflow assumes that the following components already exist:

- a policer control policy

- an access ingress or access egress policy
- a subscriber
- an access interface or SAP

50.26.2 Stages

1

Configure arbiters in a policer control policy. See 50.62 “To configure a policer control policy” (p. 1604) for more information.

1. Create arbiters.
2. Configure parent/child relationships between arbiters to build a policer control hierarchy.

2

Configure a policer in a SAP access ingress or access egress policy. See 50.28 “To configure a SAP access ingress policy” (p. 1538) and 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) for more information.

1. Create a policer to configure traffic control parameters.
2. Assign a parent arbiter to the policer to associate it with the hierarchy configured in the policer control policy.

3

Assign the policer control policy to a subscriber that is configured with a SAP access ingress or access egress policy that includes a policer. See 64.4 “To configure a subscriber profile” (p. 1840) for more information.

4

As required, assign the policer control policy to an access interface or SAP that is configured with a SAP access ingress or access egress policy that includes a policer.

a. For an IES L3 access interface, see 78.28 “To configure an L3 access interface on an IES site” (p. 2472) for more information.

b. For an IES SAP, see 78.19 “To configure a group interface on an IES” (p. 2449) for more information.

c. For a VPLS L2 access interface, see 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) for more information.

d. For a VPLS B-L2 access interface, see 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) for more information.

e. For a VPLS I-L2 access interface, see 77.88 “To create a VPLS I-L2 access interface” (p. 2372) for more information.

f. For a VPRN L3 access interface, see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) for more information.

g. For a VPRN SAP, see 79.37 “To configure a group interface on a VPRN” (p. 2586) for more information.

50.27 Workflow to configure a named buffer pool

50.27.1 Stages

1

Configure a named pool buffer Policy. See 50.67 “To configure a named buffer pool policy” (p. 1610) for more information.

2

Enable named pool to be configured for a port or MDA. See 15.67 “To enable named pool mode” (p. 526) for more information.

3

Apply a named pool policy to an MDA. See 15.78 “To configure an MDA” (p. 536) for more information.

4

Apply a named pool policy to a port. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

5

Configure Q1 pools. See 50.68 “To configure Q1 pools” (p. 1611) for more information.

6

Assign Q1 pools to SAP access ingress policies. See 50.28 “To configure a SAP access ingress policy” (p. 1538) for more information.

7

Assign Q1 pools to SAP access egress policies. See 50.30 “To configure a SAP access egress policy” (p. 1550) for more information.

8

Assign Q1 pools to network queue policies. See 50.46 “To configure a network queue policy” (p. 1581) for more information.

9

Assign Q1 pools to shared queue policies. See 50.48 “To modify a shared-queue policy” (p. 1585) for more information.

QoS policies procedures

50.28 To configure a SAP access ingress policy

50.28.1 Before you begin

A SAP access ingress policy defines forwarding for traffic entering from a service. The policy is assigned to L2 access interfaces on VPLS and VLL services, and to L3 access interfaces on IES and VPRN services. See the following procedures, according to the service type:

- for VLL: 76.42 “To assign ingress and egress QoS policies to a VLL L2 access interface” (p. 2181)
- for VPLS: 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340)
- for IES: 78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477)
- for VPRN: 79.89 “To assign ingress and egress QoS policies to a VPRN L3 access interface” (p. 2662)

50.28.2 Steps

1

Choose Policies→QoS→SROS QoS→Access Ingress→SAP Access Ingress from the NFM-P main menu. The SAP Access Ingress Policies form opens.

2

Click Create or choose a policy and click Properties. The SAP Access Ingress Policy (Create|Edit) form opens.

3

Configure the required general parameters.

If you are configuring a SAP access ingress policy for tagged IP/IPv6 match entry overrides, set the IP/IPv6 Criteria Type parameters to Tagged-Entries. Tag IDs are configured in Step 13. See 50.2.2 “Policy override” (p. 1508).

If you are configuring a SAP access ingress policy for the 7250 IXR, select a 7250 SROS Ingress Classification policy in the Ingress Classification Policy panel; see50.45 “To configure a 7250 SROS Ingress Classification policy” (p. 1579). If you do not select a user-configured 7250 SROS Ingress Classification policy, a default policy is applied.

Note: For some of the general parameters, additional fields will be displayed when certain selections are made. For example, when setting the IP Criteria Type and IPv6 Criteria Type to VXLAN-VNI and then clicking Apply, the Match Criteria field is displayed, and shows a value of None.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

Note: For some of the following steps, the tabs are not available until after the policy is created. Click Apply if required to create the policy and make all tabs available.

4

Enable the Policers HQoS Manageable parameter if the policers in this policy are to be managed by the Hierarchical QoS process.

5

Configure access ingress queues.

1. Click on the Queues tab.

   The NFM-P provides two default queues — a default unicast queue (ID 1) and a default multipoint queue (ID 11). You cannot delete the default queues.

   You can add queues up to a total of 32.
2. Click Create or choose a queue and click Properties. The Queue (Create|Edit) form opens.
3. Configure the required general parameters.

   Before you can configure the Pool Name parameter, you must create a Q1 pool; see 50.68 “To configure Q1 pools” (p. 1611).
4. Select a slope policy.
5. Select an advanced configuration policy.
6. Click on the CIR/PIR/FIR tab and configure the required parameters.

   You can configure the CIR, PIR, and FIR as a rate or as a percentage, depending on the option selected for the Rate Type parameter.

   Ensure that the CIR value is lower than the PIR value.
7. Click on the Burst Size tab and configure the required parameters.

   Deselect the Default check box for each parameter that you need to configure.

   Ensure that the Committed Burst Size (kb) value is lower than the Maximum Burst Size (bytes) value.
8. Save your changes and close the form.

6

Configure a dynamic policer.

1. Click on the Dynamic Policer tab and configure the required general parameters.
2. Select a Parent Arbiter.

3. Click on the Burst Size tab and configure the Committed Burst Size and Maximum Burst Size parameters.

   Deselect the Default check box for each parameter that you need to configure.

   Ensure that the Committed Burst Size value is lower than the Maximum Burst Size value.

7

Configure ingress policers.

1. Click on the Policer tab.
2. Click Create or choose a policer and click Properties. The Access Ingress Policer (Create|Edit) form opens.
3. Configure the required general parameters.
4. Select a Parent Arbiter and an Advanced Configuration Policy.
5. Select a Scheduler Parent if required, and then configure the parameters in the Scheduler Association panel.
6. Click on the CIR/PIR tab and configure the required parameters.

   You can configure the CIR and PIR as a rate or as a percentage, depending on the option selected for the Rate Type parameter.

   Ensure that the CIR value is lower than the PIR value.
7. Click on the Burst Size tab and configure the required parameters.

   Deselect the Default check box for each parameter that you need to configure.
8. Save your changes and close the form.

8

Configure forwarding classes.

1. Click on the Forwarding Classes tab.
2. Click Create or choose a forwarding class and click Properties. The Forwarding Class (Create|Edit) form opens.
3. Configure the Forwarding Class parameter.
4. Perform any of the following:
   - If you need to configure a queue group for one or more traffic types, complete substeps 5 and 6 .
   - If you need to configure queues for one or more traffic types, complete substep 6 .
   - If you need to configure policers for one or more traffic types, complete substeps 7,8, and 9.
5. In the Traffic Control panel, select the Use Queue Group option for any of the Unicast, Broadcast, Multicast, and Unknown traffic type parameters.

   In the Queues panel, select a queue group template policy.

   When the Use Queue Group option is selected in the global policy, the Queue ID, Multipoint Queue ID, Broadcast Queue ID, and Unknown Queue ID parameters, if selected, are

validated against the specified Queue Group Template Policy. The Queue ID is validated against the local queues for each local policy instance.

6. In the Queues panel, configure the parameters:
   - Queue ID

     The parameter must be set to a unicast ID. Set the parameter to 0 if you want the default queue ID to be used.
   - Multipoint Queue ID

     The parameter must be set to a multicast ID. Set the parameter to 0 if you want the default queue ID to be used.
   - Broadcast Queue ID

     The parameter must be set to a multicast ID. Set the parameter to 0 if you want the default queue ID to be used.
   - Unknown Queue ID

     The parameter must be set to a multicast ID. Set the parameter to 0 if you want the default queue ID to be used.
7. In the Traffic Control panel, select the Use Policer option for any of the Unicast, Broadcast, Multicast, and Unknown traffic type parameters.
8. In the Policers panel, for each traffic type configured with the Use Policer option, configure the corresponding Policer ID parameter.

   The FP Redirect Group, Multipoint FP Redirect Group, Broadcast FP Redirect Group, and Unknown FP Redirect Group parameters are displayed only when the associated traffic types in the Traffic Control panel are configured to use policers.
9. Select a policer for each Policer ID parameter. The policer must exist in the SAP access ingress policy from Step 7 .
10. Configure the required remaining parameters.

    The In Precedence parameter is configurable when the In Remark parameter is set to precedence.

    The Out Precedence parameter is configurable when the Out Remark parameter is set to precedence.

    The In DSCP parameter is configurable when the In Remark parameter is set to dscp.

    The Out DSCP parameter is configurable when the Out Remark parameter is set to dscp.
11. Click on the Sub Classes tab to configure subclasses for the forwarding class.
12. Click Create or choose a sub class and click Properties. The Forwarding Sub Class (Create|Edit) form opens.
13. Configure the required parameters.

    The In Precedence parameter is configurable when the In Remark parameter is set to precedence.

    The Out Precedence parameter is configurable when the Out Remark parameter is set to precedence.

    The In DSCP parameter is configurable when the In Remark parameter is set to dscp.

    The Out DSCP parameter is configurable when the Out Remark parameter is set to dscp.
14. Save your changes and close the forms.

9

Map ingress dot1p bits.

1. Click on the Dot1p tab.
2. Click Create or choose a Dot1p entry and click Properties. The Dot1p (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

10

Map ingress DSCP bits.

1. Click on the DSCP tab.
2. Click Create or choose a DSCP entry and click Properties. The Dscp (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

11

Map ingress Precedence bits.

1. Click on the Precedence tab.
2. Click Create or choose a Precedence entry and click Properties. The Precedence (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

12

Map ingress LSP EXP bits.

1. Click on the LspExp tab.
2. Click Create or select an LspExp entry and click Properties. The Lsp Exp (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

13

Configure match criteria. Perform any the following:

Note: If you configure MAC match criteria entries, you cannot configure IPv4 or IPv6 match criteria entries. Conversely, when no MAC match criteria entries are configured, you can configure both IPv4 and IPv6 entries.

a. Configure IPv4 match criteria.

1. Click on the IP Match Criteria tab.
2. Click Create or choose an entry and click Properties. The IP Match (Create|Edit) form opens.
3. Configure the required parameters.

   You can configure the Tag ID parameter only when the IP/IPv6 Criteria Type is set to Tagged-Entries in Step 3.

   If you set the IP/IPv6 Criteria Type to Tagged-Entries, you cannot enable the Destination Port parameter in the Port Properties panel.

   The Source Port and Destination Port parameters are available only when the Protocol parameter value is TCP, UDP, or UDPTCP (*).

   The VXLAN-VNI Properties panel is only displayed when the Protocol parameter value is set to UDP (17). You can then configure the required Operator parameter.

   To configure some parameters, you must select the check box for the parameter.

   The Src Mask and Src Net Mask parameters are mutually exclusive.

   The Source IP and Source IP Prefix parameters are mutually exclusive.

   The Dst Mask and Dst Net Mask parameters are mutually exclusive.

   The Destination IP and Destination IP Prefix parameters are mutually exclusive.
4. In the IP Properties panel, select QoS IP Prefix List policies, as required.

   - For the Source IP Prefix and Destination IP Prefix parameters, enable the checkbox and click Select. The Select [Source|Destination] IP Prefix form opens.

   - Click Create, or click Search and choose an existing QoS IP Prefix list policy.

   See 50.33 “To configure a QoS prefix list policy” (p. 1561).
5. Select a bypass policer, if required.

   The bypass policer must be local to the SAP (that is, it cannot be part of a queue group), and it must already be created within the SAP ingress policy, as described in Step 7 .
6. Save your changes and close the form.

b. Configure MAC match criteria.

1. Click on the MAC Match Criteria tab.
2. Click Create or choose an entry and click Properties. The MAC Match (Create|Edit) form opens.
3. Configure the required parameters.

   The parameters that are available vary depending on the option selected for the Frame Type parameter.

   To configure some parameters, you must select the check box for the parameter.
4. Select a bypass policer, if required.

   The bypass policer must be local to the SAP (that is, it cannot be part of a queue group), and it must already be created within the SAP ingress policy, as described in Step 7 .
5. Save your changes and close the form.

c. Configure IPv6 match criteria.

1. Click on the IPv6 Match Criteria tab.
2. Click Create or choose an entry and click Properties. The IPv6 Match form opens.
3. Configure the required parameters.

   The Source Port and Destination Port parameters appear only when the Protocol parameter value is TCP, UDP, or UDPTCP (*).

   The VXLAN-VNI Properties panel is only displayed when the Protocol parameter value is set to UDP (17). You can then configure the required Operator parameter.

   To configure some parameters, you must select the check box for the parameter.

   The Src Mask and Src Net Mask parameters are mutually exclusive.

   The Dst Mask and Dst Net Mask parameters are mutually exclusive.
4. In the IP Properties panel, select QoS IP Prefix List policies, as required.

   - For the Source IPv6 Prefix and Destination IPv6 Prefix parameters, enable the checkbox and click Select. The Select [Source|Destination] IPv6 Prefix form opens.

   - Click Create, or click Search and choose an existing QoS IP Prefix list policy.

   See 50.33 “To configure a QoS prefix list policy” (p. 1561).
5. Select a bypass policer, if required.

   The bypass policer must be local to the SAP (that is, it cannot be part of a queue group), and it must already be created within the SAP ingress policy, as described in Step 7 .
6. Save your changes and close the form.

14

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.29 To configure a 7210, 7250, and 1830 SAP Access Ingress policy

50.29.1 Purpose

A 7210, 7250 and 1830 SAP Access Ingress policy defines forwarding for traffic entering from a service. The policy specifies the following:

- an FC and profile for packets, based on MAC and IP match criteria, or dot1p or DSCP classification
- a default FC and profile for packets that do not meet any of the match criteria
- mapping of FCs to ingress meters or queues
- ingress meters, with rate values (CIR, PIR), burst sizes, and algorithm modes
- ingress queues, with rate values (CIR, PIR), scheduler priority values, and assigned queue management or slope policies

Meters are defined for unicast and multipoint traffic, and are mapped to FCs according to the user-

defined configuration in the policy. The maximum number of meters supported on a local policy varies depending on the chassis type.

The 7210 SAS-K and 7210 SAS-X support ingress queues. The 7210, 7250 and 1830 SAP Access Ingress policy allows the creation of up to eight queues per SAP, and maps the queues to FCs according to the user-defined configuration in the policy. For the 7210 SAS-K, unicast and multicast ingress queues are configured in the policy. For the 7210 SAS-X, ingress multipoint traffic is handled separately; see 52.16 “To configure a 7210 multipoint bandwidth management policy” (p. 1725) .

The 7210, 7250 and 1830 SAP Access Ingress policy is assigned to L2 access interfaces on VPLS and VLL services, and to L3 access interfaces on IES and VPRN services. See the following procedures, according to the service type:

- 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
- 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
- 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
- 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN

The NFM-P provides one default 7210, 7250 and 1830 SAP Access Ingress policy (policy ID 1). The default policy cannot be deleted or modified, and is assigned to an interface unless a user-defined policy is explicitly assigned. The default policy defines two default meters, one for unicast traffic (Meter 1) and one for multipoint traffic (Meter 11). The default meters map all packets to the *be* (best effort) FC.

MAC and IP match criteria entries consume NE system resources. The system resource profile for a 7210 SAS NE that uses the policy must be configured appropriately. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” for more information.

You can override the settings for meters and queues in access ingress policies that are assigned to access interfaces. For information about policy overrides for access ingress meters on 7210 SAS NEs, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) . For information about policy overrides for access ingress queues on 7210 SAS-X NEs, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

The 7210, 7250 and 1830 SAP Access Ingress policy is supported on all 7210 SAS chassis types. If you are configuring a SAP access ingress policy for a 7250 IXR, see 50.28 “To configure a SAP access ingress policy” (p. 1538). For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

Support for specific parameters and features depends on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

50.29.2 Steps

1

Choose Policies→QoS→SROS QoS→Access Ingress→7210, 7250 and 1830 Access Ingress Policies from the NFM-P main menu. The 7210, 7250 and 1830 Access Ingress Policies form opens.

2

Click Create, or choose a policy and click Properties. The 7210, 7250 and 1830 SAP Access Ingress Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

To avoid deployment errors, configure the Number of QoS Classifiers parameter to coordinate with the system resource profile settings for the NE. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” . The combined SAP-based allocations must not exceed the system allocations. See the NE documentation for more information.

For supporting 7210 SAS-Mxp NEs, table-based ingress classification is available using an assigned FC Meter Map policy. You must set the Use SVC Meter Pool to True, and assign an FC Meter Map policy in the Classification Type panel.

You can configure the IPV6 Match Criteria type parameter only when the Allow any IPV6 Match parameter is set to true.

You cannot change the setting of the MAC Match Criteria type, IP Match Criteria type, or IPV6 Match Criteria type parameters when entries exist for that match type. You must delete the entries first, and release the global policy.

You cannot set the IP MAC Match Criteria type parameter to None, or change it from None, when entries exist. You must delete the entries first, and release the global policy.

You cannot change the setting of the IP MAC Match Criteria type parameter when the policy is assigned to an access interface.

When the IP MAC Match Criteria type parameter is set to IP First or MAC first, the Allow any IPV6 Match parameter must be set to false.

4

If you are configuring a policy for the 7210 SAS-K, select a Dot1p classification policy in the Classification Type panel. See 50.84 “To configure a 7210/7250 Dot1p classification policy” (p. 1637) for information about 7210 Dot1p classification policies.

The classification policies selected in the Classification Type panel apply only to the 7210 SAS-K. For 7250 IXR classification configuration, you must select an Ingress CoS classification policy in the Ingress CoS Classification panel; see Step 6.

5

If you are configuring a policy for the 7210 SAS-K, or for table-based classification on supporting chassis types, select a DSCP classification policy in the Classification Type panel.

See 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638) for information about 7210 DSCP classification policies.

The classification policies selected in the Classification Type panel apply only to the 7210 SAS-K. For 7250 IXR classification configuration, you must select an Ingress CoS classification policy in the Ingress CoS Classification panel; see Step 6.

6

If you are configuring a policy for the 7250 IXR, or for table-based classification on supporting chassis types, select a 7250 Ingress CoS classification policy in the Ingress CoS Classification panel. See 50.89 “To configure a 7250 Ingress CoS policy” (p. 1642) for information about 7250 Ingress CoS policies.

You cannot select a 7250 Ingress CoS policy that contains a 7210/7250 MPLS LSP-EXP policy.

7

If you are configuring a policy for a 7210 SAS-Mxp that supports 7210 Meter Map policies, perform the following:

1. Set the User SVC Meter Pool parameter to True. FC-Meter-Map parameters appear on the form.

   The Number of QoS Classifiers parameter supported value is 2.
2. Select a 7210 FC Meter Map policy in the Classification Type panel. See 50.87 “To configure a 7210 FC Meter Map policy” (p. 1640);
3. Go to Step 16.

8

Configure ingress meters.

The 7250 IXR NEs use the term “policer” internally, but the NFM-P policy configuration uses the term “meter”. When you configure a 7210, 7250 and 1830 SAP Access Ingress policy for 7250 IXR NEs, meters configured in the policy correspond to policers in 7250 IXR NEs.

The Number of QoS Classifiers parameter, configured in Step 3 , affects the number of meters available.

Perform the following:

1. Click on the Meter tab.
2. Select a meter and click Properties, or click Create. The Meter form opens.
3. Configure the required parameters on the General tab.

   You cannot choose ID values 1 or 11 for user-configured meters; these ID values are used by default meters.

   Note:

   To enable hierarchical ingress policing (H-metering), set the Rate Mode parameter to trTCM (RFC 4115).
4. Click on the CIR/PIR and Burst Size tabs and configure the required parameters.
5. Save your changes and close the form.

9

Configure ingress queues.

The Number of QoS Classifiers parameter, configured in Step 3 , affects the number of queues available.

1. Click on the Queues tab.
2. Select a queue and click Properties, or click Create. The Queue form opens.
3. Configure the required queue identification parameters.
4. Select a queue management policy in the Queue Management Policy panel. See 50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592) for information about queue management policies.
5. Select a slope policy in the Slope Policy panel. See 50.50 “To configure a 7210 and 1830 slope policy” (p. 1589) for information about slope policies.
6. Configure the scheduling parameters in the Port Parent panel.
7. Configure the rate parameters in the CIR/PIR panel.
8. Click on the Burst Size tab and configure the required parameters.
9. Save your changes and close the form.

   Note:

   You can configure queues in the policy anytime, but when the policy is assigned to an interface, the system resource profile for the 7210 SAS NE must allocate sufficient resources to the Queues function. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

10

Click on the Forwarding Classes tab.

11

Select a forwarding class and click Properties, or click Create. The Forwarding Class form opens.

12

Configure the Forwarding Class parameter, if required.

13

Map the FC to a queue or meter. Perform one of the following:

a. Assign a queue or queues to the forwarding class.

   1. Select a queue for each of the required queue types in the Queue panel. To create queues, see Step 9 .
   2. Save your changes and close the form.

b. Assign a meter or meters to the forwarding class.

1. Select a meter for each of the required meter types in the Meter panel. To create meters, see Step 4 .
2. Save your changes and close the form.

You can also select meters or queues by directly entering the meter or queue ID.

14

Repeat Step 11 to Step 13 for each forwarding class that you need to configure.

15

Configure match criteria entries.

Note: The Number of QoS Classifiers parameter, configured in Step 3 , affects the maximum number of entries available.
MAC and IP match criteria entries consume NE system resources. The system resource profile for a 7210 SAS NE must be configured appropriately. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” .

Perform any of the following:

a. Associate forwarding classes with MAC match criteria.

1. Click on the MAC Match Criteria tab.
2. Select an entry in the list and click Properties, or click Create. The MAC Match form opens.
3. Configure the required parameters.

   You can configure the Source MAC, Destination MAC, and Ether Type parameters only when the MAC Match Criteria type parameter is set to Any in Step 3 .
4. Save your changes and close the form.

b. Associate forwarding classes with IP or IPv6 match criteria.

1. Click on the IP Match Criteria tab or IPv6 Match Criteria tab.
2. Select an entry in the list and click Properties, or click Create. The IP Match or IPv6 Match form opens.
3. Configure the required parameters.

   You can configure the Fragment parameter only when the IP Match Criteria type parameter is set to Any in Step 3 .

   You can configure the Protocol, Source IP, Src Mask, Destination IP, and Dst Mask parameters only when the IP Match Criteria type parameter is set to Any in Step 3 .

   You can configure the Source Port, Port Src, Destination Port, and Port Dst parameters only when the IP Match Criteria type parameter is set to Any in Step 3 and the Protocol parameter is set to TCP or UDP.

   You can configure parameters in only one of the DSCP and IP Precedence panels. When you configure one panel, the other is not available.
4. Save your changes and close the form.

16

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.30 To configure a SAP access egress policy

50.30.1 Steps

1

Choose Policies→QoS→SROS QoS→Access Egress→SAP Access Egress from the NFM-P main menu. The SAP Access Egress Policies form opens.

2

Click Create, or choose a policy in the list and click Properties. The SAP Access Egress Policy (Create|Edit) form opens.

3

Configure the required general parameters.

Note: For some of the following steps, the tabs are not available until after the policy is created. Click Apply if required to create the policy and make all tabs available.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

4

Enable the Policers HQoS Manageable parameter if the policers in this policy are to be managed by the Hierarchical QoS process.

5

If the SAP access egress policy requires the forwarding class to be changed to an alternate queue than the ingress forwarding class determination would have used, or you need to modify the forwarding profile to modify the congestion behavior of the egress queue, select a policy in the Post Policer Mapping Policy panel.

6

Click Select in the HS panel to associate an alternate HS Attachment Policy other than the default policy. The Select HS Attachment Policy – Egress Queue Group Template Policy form opens.

7

Select the required HS Attachment Policy and click OK.

8

If you are creating a 7705 SAR MC MLPPP access egress policy, configure the Use As Multiclass MLPPP Policy For 7705 SAR parameter and select a WRR Policy in the Egress HSMDA panel.

9

Configure HSMDA queues.

1. Click on the HSMDA Queues tab. Eight default queues are displayed.
2. Choose a queue and click Properties. The Access Egress HSMDA Queue (Create) form opens.
3. Configure the required general parameters.
4. Select a slope policy in the Slope Policy panel.
5. Click on the PIR/Burst Size tab and configure the required parameters.

   To configure the parameters, deselect the MAX or Default check boxes as required.
6. Save your changes and close the forms.

10

Configure access egress queues.

1. Click on the Queues tab.

   The NFM-P provides a default egress queue (ID 1). You can modify the default queue but you cannot delete queue ID 1.
2. Click Create, or choose a queue and click Properties. The Queue (Create|Edit) form opens.

   You can create up to eight queues.
3. Configure the required general parameters.

   Note:

   You must enable the Use WRED Queue parameter to display and configure additional parameters associated with a WRED queue.
4. Select a Slope Policy, Pool Name, Advanced Configuration Policy, and Scheduler.

   Before you can select a Pool Name, you must create a Q1 pool; see 50.68 “To configure Q1 pools” (p. 1611) .

   Scheduler selection is available when the Port Parent parameter is set to false.

   The hardware aggregate shaper scheduler policy can be supported using Shaper Weight and Scheduler Class parameters. See 50.64 “To configure a Hardware Aggregate Shaper Scheduler policy” (p. 1608).
5. Click on the CIR/PIR tab and configure the required parameters.

You can configure the CIR and PIR as a rate or as a percentage, depending on the option selected for the Rate Type parameter.

6. Click on the Burst Size tab and configure the required parameters.

   Deselect the Default check box for each parameter that you need to configure.

   Ensure that the Committed Burst Size (KB) value is lower than the Maximum Burst Size (bytes) value.
7. Click on the HSQ tab if required and configure the required parameters. You can only configure HSQ on default queue entries 1 to 8.
8. Select a HS WRED Slope Policy if required and click OK.

   Note: To configure the parameters on the HS WRED Slope Policy Properties form, you must first create a HS WRED Slope policy. See 50.49 “To configure a WRED slope policy” (p. 1587).
9. Save your changes and close the form.

11

Configure a dynamic policer.

1. Click on the Dynamic Policer tab and configure the required general parameters.
2. Select a Parent Arbiter.
3. Click on the Burst Size tab and configure the Committed Burst Size and Maximum Burst Size parameters.

   Deselect the Default check box for each parameter that you need to configure.

   Ensure that the Committed Burst Size value is lower than the Maximum Burst Size value.

12

Configure egress policers.

1. Click on the Policer tab.
2. Click Create, or choose a policer and click Properties. The Access Egress Policer (Create|Edit) form opens.
3. Configure the general parameters.
4. Select a Parent Arbiter and an Advanced Configuration Policy.
5. Select a Scheduler Parent if required, and then configure the parameters in the Scheduler Association panel.

   Note:

   You cannot configure a Scheduler Parent and a Port Parent. They are mutually exclusive.
6. Set the Port Parent parameter to True if you need to create a port parent association for the policer, as opposed to a scheduler association. You can then configure the parameters in the Port Parent Association panel.
7. Click on the CIR/PIR tab and configure the required parameters.

You can configure the CIR and PIR as a rate or as a percentage, depending on the option selected for the Rate Type parameter.

8. Click on the Burst Size tab and configure the required parameters.

   Deselect the Default check box for each parameter that you need to configure.

9. Save your changes and close the form.

13

Configure forwarding classes. You can create up to eight forwarding classes.

1. Click on the Forwarding Classes tab.
2. Click Create, or choose a forwarding class and click Properties. The Forwarding Class (Create|Edit) form opens.
3. Configure the Forwarding Class parameter.
4. Complete one of the following for each Forwarding Class you configure:
   - If you need to configure a policer, set the Traffic Control parameter for the Use Policer option and go to 5 .
   - If you need to configure a local queue, set the Traffic Control parameter for the Use Queue option and configure the Port Redirect Queue Group parameter as required, then go to 9 .
   - If you need to configure a queue from a queue group, set the Traffic Control parameter to the Use Queue Group option and go to 7 .
5. In the Policers panel, select a policer.
6. Configure the Port Redirect Queue Group parameter as required. If you disable this parameter, then go to 7 . Otherwise go to 9 .
7. In the Queue Group panel, select a Queue Group Template Policy.
8. Configure the Instance ID parameter.
9. In the Queue panel, select a queue.

   Note:

   You cannot map a Forwarding Class to a non-existent or deleted Queue ID. Use the Select button to choose a Queue ID, as opposed to manually entering it in the Queue ID field.

   The queue you can select here is dependent on the Traffic Control parameter setting you chose in 4 , as follows:

   - Use Queue: If you specified this setting, then you can only select a locally-defined queue. The local queue in this context is either the default queue or one you created for this policy in Step 10 . The local queue you select here must be set to a unicast ID. The default ID for unicast queues is 1. Alternatively, you can manually enter 0 if you want the default Queue ID to be used.
   - Use Queue Group: If you specified this setting, then you can only select one of the queues defined in the Queue Group Template Policy you chose in 7 .

     Note:

     If the Use Queue Group option is enabled for the global policy, then the Queue ID you select is validated against the specified Queue Group Template Policy. The Queue ID is validated against the local queues for each local policy instance.

You cannot delete a Queue Group Template Policy that is referenced by a queue group.

- Use Policer: If you specified this setting and disabled the Port Redirect Queue Group parameter in 6 , then you can select one of the queues defined in the Queue Group Template Policy you chose in 7 . If you enabled the Port Redirect Queue Group parameter in 6 , then select the required local queue here.

10. In the Hsmda Queue panel, configure the Queue ID and Port Redirect Queue Group parameters, as required.
11. Configure the required parameters in the Dot1p and Properties panels.
12. Save your changes and close the form.

14

Configure IPv4 match criteria.

1. Click on the IP Match Criteria tab.
2. Click Create, or choose an entry and click Properties. The IP Match (Create|Edit) form opens.
3. Configure the required parameters.

   To configure some parameters, you must select the check box for the parameter.

   The Source Port and Destination Port parameters appear only when the Protocol parameter value is TCP, UDP, or UDPTCP (*).

   The Src Mask and Src Net Mask parameters are mutually exclusive.

   The Source IP and Source IP Prefix parameters are mutually exclusive.

   The Dst Mask and Dst Net Mask parameters are mutually exclusive.

   The Destination IP and Destination IP Prefix parameters are mutually exclusive.
4. In the IP Properties panel, select QoS IP Prefix List policies, as required.

   - For the Source IP Prefix and Destination IP Prefix parameters, enable the checkbox and click Select. The Select [Source|Destination] IP Prefix form opens.

   - Click Create, or click Search and choose an existing QoS IP Prefix list policy.

   See 50.33 “To configure a QoS prefix list policy” (p. 1561).
5. In the Bypass Policer panel, select a Policer and Queue and configure the Port Redirect Queue Group parameter, as required.

   Note:

   The bypass policer must be local to the SAP (that is, it cannot be part of a queue group), and it must already be created within the SAP egress policy, as described in Step 12 .
6. Save the changes and close the form.

15

Configure IPv6 match criteria.

1. Click on the IP Match Criteria tab.

2. Click Create, or choose an entry and click Properties. The IP Match (Create|Edit) form opens.
3. Configure the required parameters.

   To configure some parameters, you must select the check box for the parameter.

   The Source Port and Destination Port parameters appear only when the Protocol parameter value is TCP or UDP.

   The Src Mask and Src Net Mask parameters are mutually exclusive.

   The Dst Mask and Dst Net Mask parameters are mutually exclusive.
4. In the IP Properties panel, select QoS IP Prefix List policies, as required.

   - For the Source IPv6 Prefix and Destination IPv6 Prefix parameters, enable the checkbox and click Select. The Select [Source|Destination] IPv6 Prefix form opens.

   - Click Create, or click Search and choose an existing QoS IP Prefix list policy.

   See 50.33 “To configure a QoS prefix list policy” (p. 1561).
5. In the Bypass Policer panel, select a Policer and Queue and configure the Port Redirect Queue Group parameter, as required.

   Note:

   The bypass policer must be local to the SAP (that is, it cannot be part of a queue group), and it must already be created within the SAP egress policy, as described in Step 12 .
6. Save your changes and close the form.

16

Map DSCP values to FCs and profiles.

1. Click on the DSCP tab.
2. Click Create, or choose an entry and click Properties. The DSCP form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

17

Map Precedence values to FCs and profiles.

1. Click on the Precedence tab.
2. Click Create, or choose an entry and click Properties. The Precedence form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

18

Map dot1p values to FCs and profiles.

1. Click on the Dot1p tab.
2. Click Create, or choose an entry and click Properties. The Dot1p form opens.

3. Configure the required parameters.
4. Save your changes and close the form.

19

Configure the HSQ parameters if required to configure HS WRR Groups..

1. Click on the HS WRR Groups tab.
2. Choose an existing HS WRR group and click Properties. The HS WRR Group form opens.

   Note: You can only modify the existing WRR Group parameters; you cannot create new WRR Groups.
3. As required, modify the parameters.
4. Save your changes and close the form.

20

Configure Scheduler Class Elevation.

1. Click on the Scheduler Class Elevation tab.
2. Choose an existing Scheduler Class Elevation entry and click Properties. The Scheduler Class Elevation form opens.

   Note: You can only modify the existing Scheduler Class Elevation parameters; you cannot create new Scheduler Class Elevation entries.
3. As required, modify the parameters.
4. Save your changes and close the form.

21

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.31 To configure a 7210 and 1830 port access egress policy

50.31.1 Purpose

A 7210 and 1830 port access egress policy defines remarking values and queue rate and scheduling parameters for traffic that egresses the switch toward a service. The policy is assigned to access ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) . Policy definitions are applied to all SAPs on the port.

You can enable or disable remarking in the policy. Remarking is configured per FC, or a 7210 remarking policy is assigned.

You can configure CIR and PIR for up to eight queues, and you can assign queue management policies to define slope and burst size parameters for queues. Queues are mapped to forwarding classes by default; the mapping is not user-configurable.

The NFM-P provides a default 7210 and 1830 port access egress policy (policy ID 1). The default policy cannot be modified or deleted, and policy ID 1 cannot be used by another policy. The default policy specifies dot1p remarking, a default remarking policy, and queue rates set to the port line rate. The default policy is assigned to access ports if no user-configured policy is assigned.

All 7210 SAS chassis types (except the 7210 SAS-K) support the 7210 and 1830 port access egress policy. Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy.

50.31.2 Steps

1

Choose Policies→QoS→SROS QoS→Access Egress→7210 and 1830 Port Access Egress from the NFM-P main menu. The 7210 and 1830 Port Access Egress Policies form opens.

2

Click Create or choose a policy and click Properties. The Port Access Egress Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

The Remarking Type parameter is configurable when the Egress Remark parameter is set to true.

4

Select a remarking policy, if required. See 50.80 “To configure a 7210 remarking policy” (p. 1630) for information about 7210 remarking policies.

Note: A 7210 remarking policy that is assigned to a 7210 and 1830 port access egress policy must be of one of the following types:

- DOT1P
- DSCP
- DOT1P - DSCP
- DOT1P - LSP-EXP SHARED

5

Configure port access egress queues.

1. Click on the Queues tab. A list of eight queues is displayed.
2. Choose a queue and click Properties. The Port Access Egress Queue form opens.
3. Configure the required parameters on the General tab.
4. Select a queue management policy, if required. See 50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592) for information about 7210 queue management policies.

5. Click on the CIR/PIR tab and configure the required parameters.
6. Save the changes and close the form.

6

Assign remarking values to forwarding classes.

1. Click on the Forwarding Classes tab.
2. Choose an FC from the list and click Properties, or click Create. The Port Access Egress Forwarding Class form opens.
3. Configure the required parameters.

   For dot1p and DE remarking, settings for the Dot1p and Force DE value parameters take effect when the Mark DE bits parameter is set to true.
4. Save the changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.32 To configure a 7210 SAP access egress policy

50.32.1 Purpose

A 7210 SAP access egress policy defines remarking and SAP egress queue parameters for traffic egressing the switch into a service.

For packets, the policy defines SAP-based remarking values using an assigned 7210 remarking policy.

For SAP-based egress queues, the policy defines:

- rate parameters
- scheduling priority values and weights
- slope parameters and burst size values, using an assigned queue management policy or slope policy

Egress queues are mapped to forwarding classes. For 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-X devices, the mapping is static and not user-configurable. The mapping for these devices is not displayed in the NFM-P. The following table shows the system-defined mapping.

Table 50-8 7210 SAS system-defined FC-to-queue mapping

Forwarding class	Queue ID
nc	8
h1	7

Table 50-8 7210 SAS system-defined FC-to-queue mapping (continued)

Forwarding class	Queue ID
ef	6
h2	5
l1	4
af	3
l2	2
be	1

For the 7210 SAS-K, the FC-to-queue mapping is user-configurable. You must explicitly enable queues in policies distributed to the 7210 SAS-K, except for Queue 1 which is enabled by default. Only queues that are applicable are available for mapping to FCs.

The 7210 SAP access egress policy is assigned to L2 access interfaces on VPLS and VLL services, and to L3 access interfaces on IES and VPRN services. You can also assign the policy to destination L2 interfaces on mirror services. The 7210 SAP access egress policy is SAP-based, not port-based, and applies policy definitions only on the interfaces to which it is assigned. The policy is assigned to interfaces in the following procedures, according to the service type:

- 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
- 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
- 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
- 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN
- 93.10 “To create an L2 access interface on a destination site” (p. 3174) for mirror services

The NFM-P assigns a default 7210 SAP access egress policy to an interface unless a user-configured policy is explicitly assigned. The default policy (policy ID 1) cannot be altered or deleted.

The 7210 SAP access egress policy is supported on the 7210 SAS-K, 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-X. Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

Note: For the 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-X, access egress remarking is either SAP-based or port-based. L3 access egress remarking is always port-based. For the 7210 SAS-Mxp and 7210 SAS-R, L2 access egress remarking is always SAP-based.

For the 7210 SAS-X, to use the values configured in a 7210 SAP access egress policy for SAP-based remarking, you must enable the SAP QoS Marking parameter on the port properties form for access and hybrid ports; see 16.24 “To configure Ethernet ports” (p. 599) .

See the NE documentation for more information.

50.32.2 Steps

1

Choose Policies→QoS→SROS QoS→Access Egress→7210 SAP Access Egress from the NFM-P main menu. The 7210 SAP Access Egress Policies form opens.

2

Click Create or choose a policy and click Properties. The 7210 SAP Access Egress Policy (Create|Edit) form opens.

3

Configure the required parameters.

For 7210 SAS-K NEs, you must enable the Remarking Policy Applicable parameter to apply the remarking policy selected in Step 4 , or a default remarking policy is applied.

4

Select a remarking policy in the Remarking Policy panel. See 50.80 “To configure a 7210 remarking policy” (p. 1630) for more information about 7210 remarking policies.

5

Configure SAP egress queues.

1. Click on the Queues tab. Eight default queues are displayed.
2. Choose a queue and click Properties. The Queue, 7210 SAP Access Egress Policy form opens.
3. Configure the required parameters on the General tab.

   The Port Parent panel contains scheduling parameters.
4. Select a Queue Management Policy in the Queue Management Policy panel.
5. Select a Slope Policy in the Slope Policy panel.
6. Click on the CIR/PIR and Burst Size tabs and configure the required parameters.
7. Save your changes and close the form.

6

Click on the Forwarding Classes tab.

7

Map FCs to queues.

1. Click Create, or choose an FC from the list and click Properties. The Forwarding Class form opens.
2. Configure the Forwarding Class parameter and select queues as required.

   Queues are available only when the Queue Applicable parameter is enabled for the queue.

3. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.33 To configure a QoS prefix list policy

50.33.1 Purpose

QoS IPv4 and IPv6 prefix lists are associated with the following QoS policies:

- SAP access ingress policies ; see 50.28 “To configure a SAP access ingress policy” (p. 1538).
- SAP access egress policies; see 50.30 “To configure a SAP access egress policy” (p. 1550).
- Network policies ; see 50.41 “To configure a QoS network policy” (p. 1568).

Various tabs on the QoS IP Prefix List form show the associations for lists in distributed policies.

50.33.2 Steps

1

Choose Policies→QoS→SROS QoS→QoS Match→ Prefix List from the NFM-P main menu. The QoS – Prefix Lists form opens.

2

Click Create, or choose a policy in the list and click Properties. The QoS – IP Prefix List (Create|Edit) form opens.

3

If you are creating a new policy, configure the IP Prefix List Name and IP Prefix List Type parameters, and click Apply.

4

Configure IP Prefix List members.

1. Click on the IP Prefix List Members tab.
2. Click Create, or choose an entry and click Properties. The QoS IP Prefix List Member form opens.
3. Configure the Prefix and Mask parameters.
4. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To

release and distribute a policy" (p. 1568) to distribute the policy to NEs.

End of steps

50.34 To configure a QoS port list policy

50.34.1 Purpose

A QoS port list policy defines TCP/UDP ports or ranges of port values for use as IPv4 and IPv6 match criteria.

QoS port list policies are associated with QoS network policies; see 50.41 "To configure a QoS network policy" (p. 1568).

Tabs on the QoS - Port List form show the associations for lists in distributed policies.

50.34.2 Steps

1

Choose Policies→QoS→SROS QoS→QoS Match→QoS Port List from the NFM-P main menu. The QoS - Port List form opens.

2

Click Create, or choose a policy in the list and click Properties. The QoS - Port List (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Configure Port List members.

Port list policies can contain a maximum of 16 members.

1. Click on the Port List Members tab.
2. Click Create, or choose a member entry and click Properties. The Port List Member form opens.
3. For the Port parameter, choose EQUAL or RANGE in the drop-down, and configure the port value or values.
4. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 "To release and distribute a policy" (p. 1476) to distribute the policy to NEs.

End of steps

50.35 To configure an ATM QoS policy

50.35.1 Steps

1

Choose Policies→QoS→SROS QoS→ATM QoS from the NFM-P main menu. The Manage ATM QoS Policies form opens.

2

Click Create or choose an ATM QoS policy and click Properties. The ATM QoS Policy (Create | Edit) form opens.

3

Configure the parameters on the General tab.

4

Click on the QoS tab and configure the required parameters.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.36 To configure a Post Policer Mapping policy

50.36.1 Steps

1

Choose Policies→QoS→SROS QoS→Post Policer Mapping from the NFM-P main menu. The Post Policer Mapping Policy form opens.

2

Click Create or choose an Post Policer Mapping Policy and click Properties. The Post Policer Mapping Policy (Create | Edit) form opens.

3

Configure the parameters on the General tab.

4

Click on the Post Policer Mapping Entry tab and configure the required parameters.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.37 To configure a MC MLPPP ingress QoS profile

50.37.1 Steps

1

Choose Policies→QoS→SROS QoS→MLPPP→Ingress QoS Profile from the NFM-P main menu. The Manage MLPPP Ingress QoS Profile form opens.

2

Click Create or choose an existing profile and click Properties. The MLPPP Ingress QoS Profile (Create | Edit) form opens.

Note: You can edit the properties of the default ingress MC profile, but you cannot delete the profile.

3

Configure the required parameters on the General tab and click Apply.

4

Click on the Classes tab.

5

Select a class in the list and click Properties. The MLPPP Ingress QoS Profile Class (Edit) form opens.

6

Configure the Reassembly Timeout (msec) parameter and click OK.

7

Repeat Step 5 and Step 6 for each class that you need to configure.

8

Save your changes and close the form.

END OF STEPS

50.38 To configure a MC MLPPP egress QoS profile

50.38.1 Steps

1

Choose Policies→QoS→SROS QoS→MLPPP→Egress QoS Profile from the NFM-P main menu. The Manage MLPPP Egress QoS Profile form opens.

2

Click Create or choose an existing profile and click Properties. The MLPPP Ingress QoS Profile (Create | Edit) form opens.

Note: You can edit the properties of the default egress MC profiles, but you cannot delete them.

3

Configure the required parameters on the General tab and click Apply.

4

Click on the Classes tab.

5

Select a class in the list and click Properties. The MlpppEgressQosProfileClass (Edit) form opens.

6

Configure the required parameters and click OK.

7

Repeat Step 5 and Step 6 for each class that you need to configure.

8

Click on the Forwarding Classes tab.

9

Select a forwarding class in the list and click Properties. The MlpppEgressQosProfileForwardingClass (Edit) form opens.

10

Configure the MLPPP Class parameter and click OK.

11

Repeat Step 9 and Step 10 for each forwarding class that you need to configure.

12

Save your changes and close the form.

END OF STEPS

50.39 To configure a MCFR ingress QoS profile

50.39.1 Steps

1

Choose Policies→QoS→SROS QoS→MCFR→Ingress QoS Profile from the NFM-P main menu. The Manage MCFR Ingress QoS Profile form opens.

2

Click Create or choose an existing profile and click Properties. The MCFR Ingress QoS Profile (Create | Edit) form opens.

Note: You can edit the properties of the default ingress MC profile, but you cannot delete the profile.

3

Configure the required parameters on the General tab and click Apply.

4

Click on the Classes tab.

5

Select a class in the list and click Properties. The MCFR Ingress QoS Profile Class (Edit) form opens.

6

Configure the Reassembly Timeout (msec) parameter and click OK.

7

Repeat Step 5 and Step 6 for each class that you need to configure.

8

Save your changes and close the form.

END OF STEPS

50.40 To configure a MCFR egress QoS profile

50.40.1 Steps

1

Choose Policies→QoS→SROS QoS→MCFR→Egress QoS Profile from the NFM-P main menu. The Manage MCFR Egress QoS Profile form opens.

2

Click Create or choose an existing profile and click Properties. The MCFR Egress QoS Profile (Create | Edit) form opens.

Note: You can edit the properties of the default egress MC profile, but you cannot delete the profile.

3

Configure the required parameters on the General tab and click Apply.

4

Click on the Classes tab.

5

Select a class in the list and click Properties. The MCFR Egress QoS Profile Class (Edit) form opens.

6

Configure the required parameters and click OK.

7

Repeat Step 5 and Step 6 for each class that you need to configure.

8

Save your changes and close the form.

END OF STEPS

50.41 To configure a QoS network policy

50.41.1 Steps

1

Choose Policies→QoS→SROS QoS→Network→Network from the NFM-P main menu. The Network Policies form opens.

2

Click Create, or choose a policy and click Properties. The Network Policy (Create|Edit) form opens.

3

Configure the required general parameters.

Note: To view or configure the NE policy name parameter, the Show Display Name of Form system preference must be enabled. See the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

4

Configure the parameters in the Ingress panel and Egress panel.

5

If you are configuring a network policy for a 2-port ring adapter MDA on a 7705 SAR NE, perform the following:

1. Set the Network Policy Type for 7705 parameter to Ring. The Ring Dot1p tab becomes available, and other tabs are unavailable.
2. Click on the Ring Dot1p tab and choose an entry, or click Create. The Ring Dot 1 p (Create|Edit) form opens.
3. Configure the required parameters to map dot1p values to queues and profiles.
4. When you finish mapping the required queues, go to Step 11 .

6

Configure egress forwarding classes.

1. Click on the Egress Forwarding Classes tab. Eight default FCs are displayed.
2. Click Create, or choose an FC and click Properties. The Egress Forwarding Class (Create|Edit) form opens.
3. Configure the required parameters.

   Select the Use check box to configure the Queue ID and Policer ID parameters. A valid entry for the Queue ID parameter must be specified in the queue group template policy.
4. Save your changes and close the form.

7

Configure ingress forwarding classes.

1. Click on the Ingress Forwarding Classes tab. Eight default FCs are displayed.
2. Click Create, or choose an FC and click Properties. The Ingress Forwarding Class (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

8

To configure the mapping of ingress or egress traffic markings to FCs and profiles, click on the appropriate tab. Mapping is optional and is based on combinations of customer QoS marking for LSP EXP, DSCP, Dot1p, and Precedence. Table 50-9, “Network policy traffic-mapping options” (p. 1568) describes the options.

Table 50-9 Network policy traffic-mapping options

Tab	Mapping
Ingress LSP EXP Bits	LSP EXP bits of the ingress traffic to an FC and profile
Ingress DSCP	DSCP of the ingress traffic to an FC and profile
Ingress Dot1p	Dot1p tag of the ingress traffic to an FC and profile
Egress DSCP	DSCP of the egress traffic to an FC and profile
Egress Precedence	Precedence value of the egress traffic to an FC and profile

Perform the following steps for each mapping that you want to configure.

1. Click on the required tab.
2. Click Create, or choose an entry and click Properties. A (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

9

Configure ingress match criteria.

1. Click on the Ingress IP Match Criteria or the Ingress IPv6 Match Criteria tab, as required.
2. Click Create. The [IP|IPv6] Match (Create) form opens.
3. Configure the required parameters.

   Match entries are evaluated in numerical order based on the entry ID from the lowest to highest ID value. The first entry that matches all match criteria has its action performed.
4. In the IP Properties panel, select a QoS IP Prefix List policy, as required.

   - For the Source [IP|IPv6] Prefix or Destination [IP|IPv6] Prefix parameters, enable the check box and click Select. The Select [Source|Destination] IPv6 Prefix form opens.

The Source IP and Source IP Prefix parameters are mutually exclusive. The Destination IP and Destination IP Prefix parameters are mutually exclusive.

- Click Create, or click Search and choose an existing QoS IP Prefix list policy.

See 50.33 “To configure a QoS prefix list policy” (p. 1561).

You can assign only one QoS IP Prefix list policy per match entry, either as a Source IP Prefix or a Destination IP Prefix.

5. In the Port Properties panel, select a QoS port list policy, as required.

   - For the Source Port List or Destination Port List parameters, enable the check box and click Select. The Select [Source|Destination] Port List form opens.

   The Source Port and Source Port List parameters are mutually exclusive. The Destination Port and Destination Port List parameters are mutually exclusive.

   - Click Create, or click Search and choose an existing QoS port list policy.

   See 50.34 “To configure a QoS port list policy” (p. 1562).

   You can assign only one QoS port list policy per match entry, either as a Source Port List or a Destination Port List.

6. Save your changes and close the form. The new match criteria entry is displayed in the list.

Note: If required, you can renumber the IPv4 or IPv6 match criteria entries in the global policy or in a local definition. Select the entry on the Ingress IP Match Criteria or the Ingress IPv6 Match Criteria tab and click Renumber. The Renumber Entry ID form opens and displays the existing (Old) Entry ID. Configure a New Entry ID and click OK.

Note: The use of the Source IP Prefix parameter in the IP Properties panel within a single entry is mutually exclusive with the [Source|Destination] Port List parameters in the Port Properties panel.

10

Configure egress match criteria.

1. Click on the Egress IP Match Criteria or the Egress IPv6 Match Criteria tab, as required.
2. Click Create. The Egress IP Match or Egress IPv6 Match (Create) form opens, as appropriate.
3. Configure the required parameters.

   Match entries are evaluated in numerical order based on the entry ID from the lowest to highest ID value. The first entry that matches all match criteria has its action performed.

4. In the Port Properties panel, select a QoS port list policy, as required.

   - For the Source Port List or Destination Port List parameters, enable the check box and click Select. The Select [Source|Destination] Port List form opens.

   - Click Create, or click Search and choose an existing QoS port list policy.

   See 50.34 “To configure a QoS port list policy” (p. 1562).

   You can assign only one QoS port list policy per match entry, either as a Source Port List or a Destination Port List.

5. Save your changes and close the form. The new match criteria entry is displayed in the list.

Note: If required, you can renumber the IPv4 or IPv6 match criteria entries in the global policy or in a local definition. Select the entry on the Egress IP Match Criteria or the Egress IPv6 Match Criteria tab and click Renumber. The Renumber Entry ID form opens and displays the existing (Old) Entry ID. Configure a New Entry ID and click OK.

11

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.42 To configure a 7210 and 1830 network policy

50.42.1 Purpose

A 7210 and 1830 network policy defines forwarding for network traffic on both ingress and egress.

For ingress traffic, the policy associates dot1p, DSCP, or LSP-EXP bits to forwarding classes and profile states. The associations are either configured directly in the network policy, or by using an assigned Dot1p, DSCP, or MPLS LSP-EXP classification policy. For packets with no dot1p, DSCP, or LSP-EXP values, a user-defined default FC and profile are assigned.

FCs are mapped to ingress meters or queues. You can configure unicast and multipoint ingress meters. One default unicast meter and one default multipoint meter are provided within the policy. The FC-to-meter mapping is user-configurable. The network policy defines rate, burst size, and algorithm mode for ingress meters.

Ingress queues (supported on the 7210 SAS-K) are configured directly within the policy. Eight queues are provided by default. You can configure the following for network ingress queues:

- CIR and PIR as a percentage of port bandwidth
- scheduling and burst size parameters
- WRED slope parameters using an assigned slope policy; see 50.50 “To configure a 7210 and 1830 slope policy” (p. 1589)

Ingress queue-to-FC mapping is user-configurable.

For egress traffic, a network policy defines remarking behavior. Remarking is configured per FC, or a 7210 remarking policy is assigned to the network policy.

On egress, FCs are statically mapped to egress queues by default; the FC-to-queue mapping is not configurable. Rates and other parameters for egress queues are defined in a separate network queue policy; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583) .

There are two types of 7210 and 1830 network policies: port type and network interface type. Port-type network policies are assigned to network, hybrid, or L2 uplink ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) . Network interface-type policies are assigned to L3 network interfaces; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863) . For network interface-type policies, you can

select an MPLS LSP-EXP Profile. See 50.83 “To configure a 7210 MPLS LSP-EXP Mapping policy” (p. 1635) for information about creating an MPLS LSP-EXP Profile.

7210 and 1830 network policies are identified by both the Nw Mgr ID value and the Policy ID value. The Nw Mgr ID value is unique for all 7210 and 1830 network policies. The Policy ID value is unique for a policy type (port or network interface). It is possible to have two 7210 and 1830 network policies with the same Policy ID if one is a port-type policy and one is a network interface-type policy. However, each of these policies will have a different Nw Mgr ID.

The NFM-P provides two default 7210 and 1830 network policies. Ports use default network policy ID 1 and network interfaces use default network policy ID 2. You cannot modify or delete a default policy. A default network policy is assigned to a port or network interface unless a user-configured policy is explicitly assigned. For the 7210 SAS-K12 and 7210 SAS-K30 ETR, a default port-type policy for network ports is automatically created as Policy ID 2 when the device is discovered. An auto-assigned Nw Mgr ID value distinguishes it from the network interface default policy.

The 7210 and 1830 network policy is supported on all 7210 SAS chassis types. Support for specific parameters and features depends on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy.

50.42.2 Steps

1

Choose Policies→QoS→SROS QoS→Network→7210 and 1830 Network from the NFM-P main menu. The 7210 and 1830 Network Policies form opens.

2

Perform one of the following:

a. Modify an existing policy. Select a policy and click Properties. The Network Policy (Edit) form opens.

b. Create a new policy.

   1. Click Create. The Network Policy (Create) form opens.
   2. Configure the Nw Mgr ID, Policy ID, and Type parameters.

      The Nw Mgr ID value is unique for all 7210 and 1830 network policies regardless of type. The Policy ID value is unique for policies within each type (port or network interface).
   3. Click Apply to create the policy. The Network Policy (Edit) form is displayed.

      Note:

      You cannot modify the Nw Mgr ID, Policy ID, or Type parameter after the policy is created.

3

Configure the required parameters on the General tab.

The Remarking Type parameter is configurable only when the Type parameter is set to Port and the Remarking parameter is set to true.

4

In the Classification Type panel, select a Dot1p, DSCP, or MPLS LSP-EXP classification policy, as required. See 50.84 “To configure a 7210/7250 Dot1p classification policy” (p. 1637), 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638) or 50.83 “To configure a 7210 MPLS LSP-EXP Mapping policy” (p. 1635).

You can select Dot1p, DSCP, or MPLS LSP-EXP classification policies only for a port-type network policy.

5

Select an MPLS LSP-EXP Profile in the Ingress panel, if required.

You can select an MPLS LSP-EXP Profile only for a network interface-type policy.

Note: The MPLS LSP-EXP Profile defines how LSP-EXP bits are mapped to profile states on ingress. The MPLS LSP-EXP Profile is also referred to as an MPLS LSP-EXP Map policy. See 50.83 “To configure a 7210 MPLS LSP-EXP Mapping policy” (p. 1635) for more information about MPLS LSP-EXP Map policies.

6

Select a remarking policy in the Egress panel, if required. See 50.80 “To configure a 7210 remarking policy” (p. 1630) for information about creating a remarking policy.

7

Associate egress forwarding classes to dot1p, DE, DSCP, or LSP-EXP values.

1. Click on the Egress Forwarding Classes tab. Eight forwarding classes are listed with associated dot1p, DE, DSCP, or LSP-EXP values.

   Port-type policies display dot1p, DE, and DSCP values. Network interface-type policies display LSP-EXP values.
2. Choose a forwarding class and click Properties. The Egress Forwarding Class form opens.
3. Configure the required parameters.

   For dot1p and DE remarking for port-type network policies, settings for the Dot1p and Force DE value parameters take effect when the Mark DE bits parameter is set to true.
4. Save your changes and close the form.

8

If you are configuring a network interface-type policy, associate ingress LSP EXP values with forwarding classes and profiles. Otherwise, go to Step 9 .

1. Click on the Ingress LSP EXP Bits tab.
2. Choose an LSP EXP value from the list and click Properties, or click Create. The Network Ingress Lsp Exp, 7210 Network Policy form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

5. Go to Step 11 .

9

If you are configuring a port-type policy, perform one of the following:

a. Associate ingress dot1p values with forwarding classes and profiles.
   1. Click on the Ingress Dot1p tab.
   2. Choose a dot1p value from the list and click Properties, or click Create. The Network Ingress Dot1p, Network Policy form opens.
   3. Configure the required parameters.
   4. Save your changes and close the form.

b. Associate ingress DSCP values with forwarding classes and profiles.
   1. Click on the Ingress DSCP tab.
   2. Choose a DSCP value from the list and click Properties, or click Create. The Network Ingress DSCP, Network Policy form opens.
   3. Configure the required parameters.
   4. Save your changes and close the form.

10

Configure ingress queues.

1. Click on the Queues tab.
2. Choose an ingress queue from the list and click Properties. The Network Ingress Queue, Network Policy form opens.
3. Configure the required parameters on the General tab.
4. Select a slope policy in the Slope Policy panel.
5. Click on the CIR/PIR and Burst Size tabs and configure the required parameters.
6. Save your changes and close the form.

11

Configure ingress meters.

1. Click on the Ingress Meter tab.
2. Choose an ingress meter from the list and click Properties, or click Create. The Network Ingress Meter, Network Policy form opens.
3. Configure the required parameters on the General tab.
4. Click on the CIR/PIR and Burst Size tabs and configure the required parameters.
5. Save your changes and close the form.

12

Map forwarding classes to queues or meters.

1. Click on the Ingress FC tab.
2. Choose an entry from the list and click Properties, or click Create. The 7210 Network Ingress Forwarding Class, Network Policy form opens.
3. Configure parameters and select queues, as required.
4. Save your changes and close the form.

13

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

End of steps

50.43 To configure a 7250 SROS Network Ingress policy

50.43.1 Purpose

A 7250 SROS Network Ingress policy defines forwarding for ingress traffic on a network interface.

The policy associates ingress Dot1p, DSCP, or LSP-EXP bits with forwarding classes and profiles, using an assigned 7250 SROS Ingress Classification policy.

The 7250 SROS Network Ingress policy also defines network ingress policers. Eight unicast policers, with IDs 1 to 8, exist by default and cannot be deleted. The CIR, PIR, and burst size values for these policers are user-configurable; if no user-configured values are defined, default values are used. You can configure up to eight additional policers for multicast, using IDs 9 to 16.

Ingress policers are mapped to FCs. The default mapping for unicast policers 1 to 8 is fixed, and not user-configurable. The mapping of FCs to multicast policers is user-configurable.

7250 SROS Network Ingress policies are assigned to L3 network interfaces; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863).

7250 SROS Network Ingress policies are identified by the policy name; policy ID numbers are not used.

A default 7250 SROS Network Ingress policy is applied if no user-configured policy is explicitly assigned. The default policy cannot be modified or deleted.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.43.2 Steps

1

Choose Policies→QoS→SROS QoS→Network→7250 SROS Network Ingress from the NFM-P main menu. The 7250 SROS Network Ingress form opens.

2

Click Create, or choose a policy and click Properties. The 7250 SROS Network Ingress Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

In the Ingress Classification Policy panel, select a 7250 SROS Ingress Classification policy; see 50.45 “To configure a 7250 SROS Ingress Classification policy” (p. 1579).

5

Configure policers.

Policers 1 through 8 exist in the policy by default. Unicast traffic uses policers 1 to 8.

You can create additional policers 9 through 16. These additional policers can be mapped to FCs for multicast traffic only.

Click on the Policer tab and perform the following:

1. Choose a policer from the list and click Properties, or click Create. The Policer, 7250 SROS Network Ingress (Create|Edit) form opens.
2. Configure the required parameters on the General tab.
3. Click on the CIR/PIR and Burst Size tabs and configure the required parameters.
4. Save your changes and close the form.

6

To view the mapping of policers to forwarding classes, click on the Forwarding Classes tab.

The mapping of FCs to unicast policers 1 through 8 is fixed, and is not user-configurable.

The mapping of FCs to multicast policers is optional; see Step 7.

7

To configure the mapping of multicast policers to forwarding classes, click on the Forwarding Classes tab. For each required FC, perform the following:

1. Choose the forwarding class from the list and click Properties. The Network Ingress Forwarding Class (Create) form opens.
2. Configure the Multicast Policer ID parameter.

   You can enter a policer ID from 1 to 16. You can assign one of the policers that exist in the policy by default (1 to 8), or a user-created policer (9 to 16).
3. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.44 To configure FC mapping policies for ingress classification on the 7250 IXR

50.44.1 Purpose

7250 SROS FC mapping policies for ingress classification associate Dot1p, DSCP, and LSP EXP CoS values with forwarding classes and profiles. A separate policy is configured for each classification type (Dot1p, DSCP, or LSP-EXP). For each policy, a default FC and profile action is configurable for ingress traffic that does not have a CoS value.

The policies are assigned to a 7250 SROS Ingress Classification policy; see 50.45 “To configure a 7250 SROS Ingress Classification policy” (p. 1579).

7250 SROS FC mapping policies are identified by the policy name; policy ID numbers are not used.

The NFM-P provides a default 7250 SROS Dot1p FC mapping policy. The default policy is not configurable, and cannot be deleted. The default policy is assigned to a 7250 SROS Ingress Classification policy when no user-configured policy is explicitly assigned.

7250 SROS FC mapping policies are not supported on all 7250 IXR chassis types and releases. For information about FC mapping policy support on the 7250 IXR, see the NE documentation.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.44.2 Steps

1

If required, configure a 7250 SROS Dot1p FC Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Ingress Classification→7250 SROS Dot1p FC Mapping from the NFM-P main menu. The 7250 SROS Dot1p FC Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS Dot1p FC Mapping (Create|Edit) form opens.
3. Configure the required parameters on the General tab.

   The default values are applied to ingress packets that do not have a CoS value.
4. Click on the Dot1p tab.
5. Click Create, or choose a mapping in the list. The Dot1p FC form opens.
6. Configure the mapping of Dot1p values, forwarding classes, and profiles.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional mappings, as required.

9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

2

If required, configure a 7250 SROS DSCP FC Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Ingress Classification→7250 SROS DSCP FC Mapping from the NFM-P main menu. The 7250 SROS DSCP FC Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS DSCP FC Mapping (Create|Edit) form opens.
3. Configure the required parameters on the General tab.

   The default values are applied to ingress packets that do not have a CoS value.
4. Click on the DSCP tab.
5. Click Create, or choose a mapping in the list. The DSCP FC form opens.
6. Configure the mapping of DSCP values, forwarding classes, and profiles.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional mappings, as required.
9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

3

If required, configure a 7250 SROS LSP EXP FC Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Ingress Classification→7250 SROS LSP EXP FC Mapping from the NFM-P main menu. The 7250 SROS LSP EXP FC Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS LSP EXP FC Mapping (Create|Edit) form opens.
3. Configure the required parameters on the General tab.

   The default values are applied to ingress packets that do not have a CoS value.
4. Click on the LSP-EXP tab.
5. Click Create, or choose a mapping in the list. The LSP-EXP FC form opens.
6. Configure the mapping of LSP-EXP values, forwarding classes, and profiles.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional mappings, as required.
9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.45 To configure a 7250 SROS Ingress Classification policy

50.45.1 Before you begin

A 7250 SROS Ingress Classification policy associates ingress Dot1p, DSCP, and LSP-EXP CoS values with forwarding classes and profiles. A default FC and profile action is configurable for ingress traffic that does not have a CoS value.

The policy is assigned to the following QoS policies:

- SAP Access Ingress policies configured for the 7250 IXR, see 50.28 “To configure a SAP access ingress policy” (p. 1538)
- 7250 SROS Network Ingress policies; see 50.43 “To configure a 7250 SROS Network Ingress policy” (p. 1575)

7250 SROS Ingress Classification policies are identified by the policy name; policy ID numbers are not used.

The NFM-P provides two default 7250 SROS Ingress Classification policies, one for SAP Access Ingress policies and the other for 7250 SROS Network Ingress policies. The default policies are not configurable, and cannot be deleted. A default policy is assigned to a SAP Access Ingress policy or 7250 SROS Network Ingress policy when no user-configured policy is explicitly assigned.

The 7250 IXR supports 7250 SROS Ingress Classification policies.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.45.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7250 SROS Ingress Classification from the NFM-P main menu. The 7250 SROS Ingress Classification form opens.

2

Click Create, or choose an existing policy and click Properties. The 7250 SROS Ingress Classification Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

If you are configuring a policy for a 7250 IXR NE that supports SROS FC mapping policies for ingress classification, assign the FC mapping policies in the Properties panel, then go to Step 7.

For more information about SROS FC mapping policies, see 50.44 “To configure FC mapping policies for ingress classification on the 7250 IXR” (p. 1577).

4

Click on the Dot1p tab to associate ingress Dot1p values.

1. Click Create, or choose a Dot1p value from the list and click Properties. The Dot1p Forwarding Class form opens.
2. Configure the parameters.
3. Save your changes and close the form.

5

Click on the DSCP tab to associate ingress DSCP values.

1. Click Create, or choose a DSCP value from the list and click Properties. The DSCP Forwarding Class form opens.
2. Configure the parameters.
3. Save your changes and close the form.

6

Click on the LSP-EXP tab to associate ingress LSP-EXP values.

1. Click Create, or choose an LSP-EXP value from the list and click Properties. The LSP-EXP Forwarding Class form opens.
2. Configure the parameters.
3. Save your changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.46 To configure a network queue policy

50.46.1 Purpose

Use this procedure to configure a network queue policy.

Note: You cannot use the same network queue policy on devices at different releases.

Note: HSMDA is only supported on 7750 SR-7/12/12e and 7450 ESS-7/12, Release 20.10 and earlier.

50.46.2 Steps

1

Choose Policies→QoS→SROS QoS→Network Queue→Network Queue from the NFM-P main menu. The Network Queue Policies form opens.

2

Click Create or choose a policy and click Properties. The Network Queue Policy (Create|Edit) form opens.

3

Configure the general parameters.

4

Click Select in the HS panel to associate an alternate HS Attachment Policy other than the default policy. The Select HS Attachment Policy – Egress Queue Group Template Policy form opens.

5

Select the required HS Attachment Policy and click OK.

6

Configure the Packet Byte Offset parameter in the Egress HSMDA panel.

7

Select a WRR policy in the WRR Policy panel and click Apply.

8

Click on the Egress HSMDA Queues tab. Eight default queues are displayed on the form.

9

Double-click on a queue in the list to configure a default Egress HSMDA queue. The Network Queue Egress HSMDA Queue (Edit) form opens with the General tab displayed.

10

Configure the WRR Weight parameter.

Note: The WRR Weight parameter is only displayed and configurable for Egress HSMDA Queues 1, 2, and 3.

11

Select a slope policy in the Slope Policy panel.

12

Click on the PIR/Burst Size tab and configure the parameters.

13

Configure the parameters and click OK.

You must deselect the associated Default check box to configure the Maximum Burst Size (bytes) or Burst Limit (bytes) parameter.

14

Click on the Queues tab. Two default queues are displayed on the form.

15

Perform one of the following:

a. To configure a default queue, double-click on a queue in the list. The Entry, Network Queue Policy (Edit) form opens with the General tab displayed.

b. To create a new queue, click Create. The Entry, Network Queue Policy (Create) form opens with the General tab displayed.

16

Configure the required parameters and select a slope policy in the Slope Policy panel.

17

Click on the CIR/PIR/FIR and Burst Size tabs and configure the required parameters.

18

Configure the HSQ parameters if required. You can only configure HSQ on default queue entries 1 to 8.

1. Click on the HSQ tab.
2. Configure the required parameters.
3. Click Select adjacent to the HS WRED Slope Policy parameter. The Select HS WRED slope Policy form opens.

   Note: To configure the HS WRED Slope Policy parameter, you must first create a HS WRED Slope policy. See 50.49 “To configure a WRED slope policy” (p. 1587).
4. Select the required policy and click OK.

19

Click OK. The Entry, Network Queue Policy configuration form closes and the Network Queue Policy (Create) form reappears with a list of the network queues displayed.

20

Click on the Forwarding Classes tab.

21

Double-click on an existing forwarding class or click Create to create a new forwarding class. The Forwarding Class, Network Queue Policy (Edit) form opens.

22

Configure the required parameters.

The Forwarding Class parameter can be configured only when you are creating a forwarding class network queue policy.

23

Configure the HSQ parameters if required to configure HS WRR Groups.

1. Click on the HS WRR Groups tab.
2. Choose an existing HS WRR group and click Properties. The HS WRR Group form opens.

   Note: You can only modify the existing WRR Group parameters; you cannot create new WRR Groups.
3. As required, modify the parameters.
4. Save your changes and close the form.

24

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.47 To configure a 7210 and 1830 network queue policy

50.47.1 Purpose

A 7210 and 1830 network queue policy defines the following for egress queues on a network-facing port:

- CIR and PIR as a percentage of port bandwidth
- scheduling priority mode and weight
- WRED slope parameters and burst size values, using an assigned queue management policy or slope policy
- burst size values; these are either system-defined, configured directly in the policy, or defined in an assigned queue management policy, depending on the NE

7210 and 1830 network queue policies are assigned to network-facing ports (network, L2 uplink, or hybrid ports); see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) . The supported port types vary, depending on the NE chassis type; see the NE documentation for more information.

7210 and 1830 network queue policies are identified by the policy name and description; policy ID numbers are not used.

The NFM-P provides default 7210 and 1830 network queue policies. A default policy is applied unless another network queue policy is explicitly assigned. The local definition of the default policy varies depending on the chassis type. See the 7210 SAS Quality of Service Guides for more information.

All 7210 SAS chassis types support the 7210 and 1830 network queue policy. Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

50.47.2 Steps

1

Choose Policies→QoS→SROS QoS→Network Queue→7210 and 1830 Network Queue from the NFM-P main menu. The 7210 and 1830 Network Queue Policies form opens.

2

Click Create or choose a policy and click Properties. The Network Queue, Global Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Configure the required queues.

1. Click on the Queues tab. Eight default queues are displayed.
2. Select a queue and click Properties. The NQueue Entry, Network Queue form opens.
3. Configure the required parameters.
4. Select a queue management policy in the Queue Management Policy panel. See 50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592) for more information about queue management policies.
5. Select a slope policy in the Slope Policy panel. See 50.50 “To configure a 7210 and 1830 slope policy” (p. 1589) for more information about slope policies.
6. Configure the CIR and PIR for the queue. Click on the CIR/PIR tab and configure the parameters.
7. Configure burst size parameters for the queue. Click on the Burst Size tab and configure the parameters.
8. Save your changes and close the form.

5

Configure FC-to-queue mapping.

1. Click on the Forwarding Classes tab.
2. Click Create, or choose an FC and click Properties. The Forwarding Class form opens.

   The Create button is not available when the Displayed Name parameter on the General tab is set to default; see Step 3 .
3. Configure the required parameters.
4. Save your changes and close the form.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.48 To modify a shared-queue policy

50.48.1 Purpose

You cannot create shared-queue policies. Instead, you modify the existing default policy (default or policer-output-queues). Modification of shared-queue policies is supported on the 7750 SR and 7450 ESS.

The default behavior automatically chooses the expedited or non-expedited nature of the queue based on the forwarding classes mapped to it. As long as all forwarding classes mapped to the queue are expedited (nc, ef, h1 or h2), the queue is treated as an expedited queue by the hardware schedulers. When any non-expedited forwarding classes are mapped to the queue (be, af, l1 or l2), the queue is treated as best effort (be) by the hardware schedulers.

50.48.2 Steps

1

Choose Policies→QoS→SROS QoS→Shared Queue from the NFM-P main menu. The Shared Queue Policies form opens.

2

Choose Local or Global from the Policy scope drop-down menu. When you choose Local, you can specify the Local Node IP Address parameter and choose a device by clicking Select and selecting a device in the list.

3

Choose a policy in the list and click Properties. The Shared Queue Policy (Edit) form opens with the General tab displayed.

4

Configure the Description parameter.

5

Click on the Queues tab.

The following table lists the queue IDs that are used by the NFM-P to identify the shared queue types.

Table 50-10 Shared queue types

Shared queue ID	Shared queue type
1 to 8	Unicast
9 to 16	Multicast
17 to 25	Broadcast
26 to 32	Unknown

There are 32 default queues displayed on the form. Configure the queues as required. Only 16 queues are available for the policer-output-queues policy.

1. Double-click on a queue in the list. The Queue configuration form opens with the General tab displayed.
2. Configure the Expedite and Pool Name parameters.

   Note:

   Before you can configure the Pool Name parameter, you must create a Q1 pool; see 50.68 “To configure Q1 pools” (p. 1611) .
3. Click on the CIR/PIR/FIR and Burst Size tables and configure the parameters.
4. Save your changes.

6

Click on the Forwarding Classes tab.

Default objects based on the eight forwarding classes are displayed on the form. You cannot associate different forwarding classes with the shared queues.

7

Select a forwarding class and click Properties. The Forwarding Class, Local Policy (Edit) form opens.

You can only configure a forwarding class if you chose Local in Step 2 .

8

Configure the required parameters and click OK.

9

Click on the L2 Interfaces or L3 Interfaces tab. (Applies to default shared queue policy only.)

10

Select an interface in the list and click Properties. The L2 or L3 Access Interface form opens with the General tab displayed.

11

Click on the QoS tab and configure the Use Multipoint Shared Queue parameter. (Applies to default shared queue policy only.)

12

Click OK. The Shared Queue Policy form reappears.

13

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.49 To configure a WRED slope policy

50.49.1 Steps

1

Choose Policies→QoS→SROS QoS→Slope→WRED Slope from the NFM-P main menu. The WRED Slope Policies form opens.

2

Click Create, or choose a policy in the list and click Properties. The WRED Slope Policy (Create|Edit) form opens.

3

Configure the required general parameters.

When packets are queued, shared buffer average utilization is calculated using the Time Average Factor for the buffer pool. The time average factor specifies the weighting between the previous shared buffer average utilization result and the new shared buffer utilization to determine the new shared buffer average utilization.

Note: The Time Average Factor parameter must be set to 3 on the 7705 SAR.

4

Click on the Burst Size tab. The Burst Size form opens, as shown in Figure 50-4, “Slope policy form - Burst Size” (p. 1588) .

Figure 50-4 Slope policy form - Burst Size

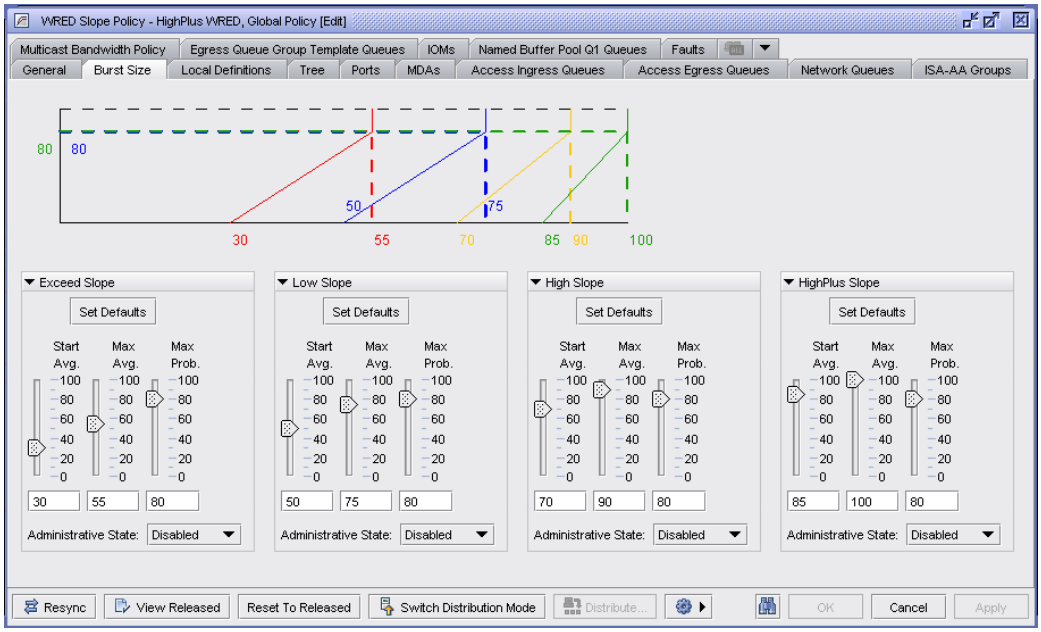

Each buffer pool supports a high-priority RED High Slope, a low-priority RED Low Slope, an Exceed Slope, and a HighPlus Slope. The High Slope manages access to the shared portion of the buffer pool for high-priority or in-profile packets. The Low Slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets. The Exceed Slope handles the exceeded profile traffic congestion control at the pool and megapool level. The values for the Exceed Slope parameters should be set so that they are lower than the Low Slope in the slope hierarchy, as shown in the example above, specifically, Start Avg=30, Max Avg=55, and Max Prob=80. Finally, the HighPlus Slope handles the inplus profile traffic congestion control. A HighPlus Slope policy can be applied to pool-per-queue queues in an egress queue group template, which in turn can be applied at a network egress port. The HighPlus Slope is also supported in a slope policy applied to the egress WRED megapool. The values for the HighPlus Slope parameters should be set so that they are higher than the High Slope in the slope hierarchy, as shown in the example above, specifically, Start Avg=85, Max Avg=100, and Max Prob=80. See 50.11 “Slope policies” (p. 1519) in this chapter for more information.

5

Configure the parameters for the Exceed Slope, Low Slope, High Slope, and the HighPlus Slope.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.50 To configure a 7210 and 1830 slope policy

50.50.1 Purpose

A 7210 and 1830 slope policy defines the values used for calculating drop precedence for egress queues in port buffer pools. Values are configured as drop rate percentages for each queue or as parameters for WRED slope calculations.

Note: The allocation of device memory for buffer pools and queues varies, depending on the chassis type. The slope calculation algorithm also varies by chassis type. See the NE documentation for more information.

The policy allows the configuration of the following slopes for a queue:

- High Slope (high priority TCP slope for in-profile packets)
- Low Slope (low priority TCP slope for out-of-profile packets)
- Non-TCP Slope

You can configure some 7210 SAS chassis types to use only TCP slopes (two slopes per queue). See 12.48 “To configure two WRED slopes on a 7210 SAS” (p. 379) .

You can also configure a High Ring Slope and a Low Ring Slope.

7210 and 1830 slope policies are identified by the policy name and description; policy ID numbers are not used.

The policy is assigned to egress buffer pools during port configuration; see 16.24 “To configure Ethernet ports” (p. 599) . The policy can also be assigned to queues in the following policies:

- 7210 and1830 SAP access ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
- 7210 SAP access egress policy; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)
- 7210 and 1830 network queue policy; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)
- 7210 and 1830 network policy (port type); see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571)

A default slope policy is applied if no policy is explicitly assigned.

The policy is supported on the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-S, 7210 SAS-Sx, and 7210 SAS-T. Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

The 7210 SAS-R and 7210 SAS-X do not support 7210 and 1830 slope policies; these devices use queue management policies instead. See 50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592) .

50.50.2 Steps

1

Choose Policies→QoS→SROS QoS→Slope→7210 and 1830 Slope from the NFM-P main menu. The 7210 and 1830 Slope Policies form opens.

2

Click Create or choose a policy and click Properties. The WRED Slope Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Configure drop rates for queues. Click on the Burst Size tab and configure the parameters in the Low Slope and High Slope panels.

5

Configure slope calculation values for queues.

1. Click on the Queue Slope tab. Eight default queues are displayed.
2. Choose a queue and click Properties. The Queue Slope, WRED Slope Policy form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

6

Click on the Slope tab and configure the required parameters.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.51 To configure an HSMDA WRED slope policy

50.51.1 Steps

1

Choose Policies→QoS→SROS QoS→Slope→HSMDA WRED Slope from the NFM-P main menu. The HSMDA Slope Policies form opens.

2

Click Create or choose a policy and click Properties. The HSMDA Slope Policy (Create|Edit) form opens.

3

Configure the general parameters.

4

Click on the HSMDA Burst Size tab. The Burst Size form opens, as shown in Figure 50-5, “HSMDA Slope Policy form - Burst Size” (p. 1590) .

Figure 50-5 HSMDA Slope Policy form - Burst Size

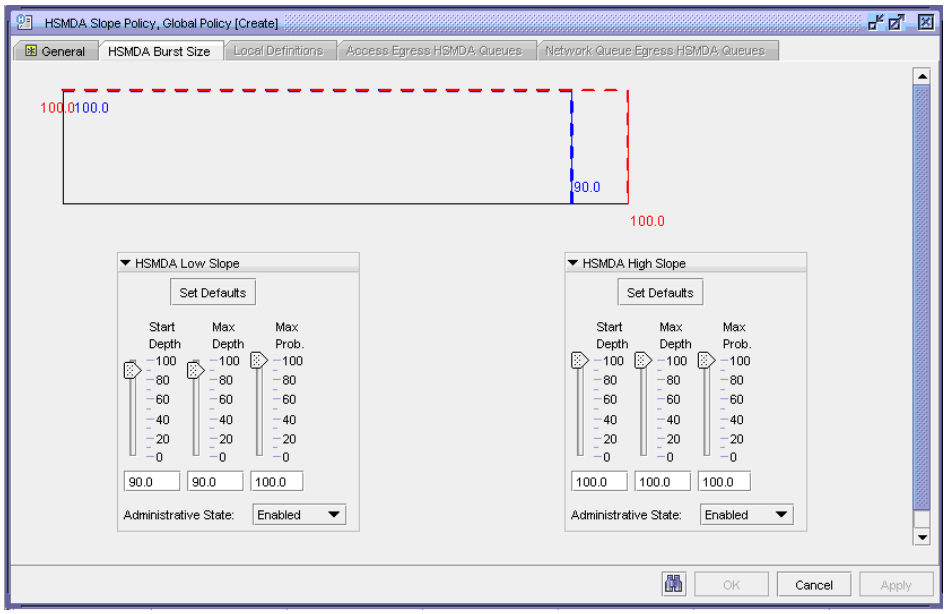

Each buffer pool supports a high-priority RED slope and a low-priority RED slope. The high-priority RED slope manages access to the shared portion of the buffer pool for high-priority or in-profile packets. The low-priority RED slope manages access to the shared portion of the buffer pool for low-priority or out-of-profile packets. See 50.11 “Slope policies” (p. 1519) in this chapter for more information.

5

Configure the parameters for the Low Slope and the High Slope and click Apply. The HSMDA Slope Policy form is refreshed with additional buttons.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.52 To configure a 7210 and 7250 Queue Management policy

50.52.1 Purpose

A 7210 and 7250 Queue Management policy defines parameters for calculating drop precedence for ingress or egress queues. The policy defines values for CBS and MBS, and for high and low WRED slope calculations.

7210 and 7250 Queue Management policies are assigned to queues during queue configuration in the following policies:

- 7210 and 1830 Access Ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
- 7210 and 1830 Port Access Egress policy; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)
- 7210 SAP Access Egress policy; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)
- 7210 and 1830 Network Queue policy; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)
- 7210, 7250 and 1830 Port Scheduler policy; see50.60 “To configure a 7210, 7250 and 1830 Port Scheduler policy” (p. 1601)

The NFM-P assigns a default 7210 and 7250 Queue Management policy to a queue if no user-configured policy is explicitly assigned. The default policy cannot be deleted or modified. The default policy allocates the appropriate amount of CBS and MBS buffers based on whether the queue is associated with a SAP or a network port. The WRED slopes (high and low) are disabled in the default policy. If WRED is not configured in a 7210 and 7250 Queue Management policy, tail drop is used.

7210 and 7250 Queue Management policies are identified by the policy name and description; policy ID numbers are not used.

The 7210 and 7250 Queue Management policy is supported on the 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-X.

If you are configuring a queue management policy for a 7250 IXR, see 50.53 “To configure a 7250 SROS Queue Management policy” (p. 1593). For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy.

50.52.2 Steps

1

Choose Policies→QoS→SROS QoS→Slope→7210 and 7250 Queue Management from the NFM-P main menu. The 7210 and 7250 Queue Management Policies form opens.

2

Click Create, or choose a policy and click Properties. The 7210 and 7250 Queue Management Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.53 To configure a 7250 SROS Queue Management policy

50.53.1 Purpose

A 7250 SROS Queue Management policy defines parameters for calculating drop precedence for egress queues. The policy defines values for Time Average Factor, MBS, and WRED slope calculations. If the slopes are not configured in the policy, the queue tail drops any excess traffic. The policy also defines a port speed value that supporting NEs use to internally optimize the allocation of virtual output queues in NE system memory.

7250 SROS Queue Management policies are assigned to queues in the following policies:

- 7250 SROS VLAN QoS policy; see 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594)
- 7250 SROS Port QoS policy; see 50.61 “To configure a 7250 SROS Port QoS policy” (p. 1603)

You can view the associated 7250 SROS VLAN QoS and 7250 SROS Port QoS policies using the 7250 SROS Queue Management policy properties form.

7250 Queue Management policies are identified by the policy name; policy ID numbers are not used.

The NFM-P provides a default 7250 SROS Queue Management policy. The default policy cannot be deleted or modified. The default queue management policy is assigned to queues in 7250 SROS VLAN QoS policies and 7250 SROS Port QoS policies if no user-configured queue management policy is explicitly assigned. The NFM-P default 7250 SROS Queue Management policy does not

assign a port speed. When the default policy is distributed to an NE that supports port speed configuration, the NE automatically assigns a port speed.

The 7250 IXR supports 7250 SROS Queue Management policies.

Support for specific parameters and features varies, depending on the chassis type and release. Local definitions of the policy do not necessarily support all configurations in the global policy.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.53.2 Steps

1

Choose Policies→QoS→SROS QoS→Slope→7250 SROS Queue Management from the NFM-P main menu. The 7250 SROS Queue Management form opens.

2

Click Create, or choose a policy and click Properties. The 7250 SROS Queue Management Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.54 To configure a 7250 SROS VLAN QoS policy

50.54.1 Purpose

A 7250 SROS VLAN QoS policy defines queue parameters for VLAN egress queues. The policy allows configuration of CIR weight profiles, CIR/PIR values, and the assignment of 7250 SROS Queue Management policies.

The 7250 SROS VLAN QoS policy also allows viewing of the mapping of forwarding classes to queues. FCs are mapped to queues by default; the mapping is not user-configurable.

The 7250 SROS VLAN QoS policy allows the enabling or disabling of VOQ (Virtual Output Queue) statistics collection, for supporting 7250 IXR chassis types. Statistics collection is enabled using the Stat Mode parameter.

For supporting NEs, the policy allows configuration of a Packet Byte Offset value. See the NE documentation for more information about PBO on the 7250 IXR.

The 7250 SROS VLAN QoS policy is assigned to L2 access interfaces on VPLS and VLL services, to L3 access interfaces on IES and VPRN services, and to network interfaces. See the following procedures:

- for VLL: 76.42 “To assign ingress and egress QoS policies to a VLL L2 access interface” (p. 2181)
- for VPLS services: 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340)
- for IES: 78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477)
- for VPRN: 79.89 “To assign ingress and egress QoS policies to a VPRN L3 access interface” (p. 2662)
- for L3 network interfaces: 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863)

7250 SROS VLAN QoS policies are identified by the policy name; policy ID numbers are not used.

A default 7250 VLAN QoS policy is applied if no user-configured policy is explicitly used. The default policy cannot be deleted or modified.

The 7250 IXR supports 7250 SROS VLAN QoS policies. For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

Support for specific parameters and features varies, depending on the chassis type. Local definitions of the policy do not necessarily support all configurations in the global policy.

50.54.2 Steps

1

Choose Policies→QoS→SROS QoS→Slope→7250 SROS VLAN QoS from the NFM-P main menu. The 7250 SROS VLAN QoS form opens.

2

Click Create, or choose a policy and click Properties. The 7250 SROS VLAN QoS Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

For supporting NEs, you can configure the Packet Byte Offset parameter when QoS Packet Byte Offset is globally enabled in NE properties; see 12.67 “To globally enable or disable Packet Byte Offset on a 7250 IXR” (p. 396).

4

Configure CIR weight profiles.

CIR weight profiles configured in this step are used during CIR configuration in Step 6.

1. Click on the CIR Weight Profiles tab, choose a profile from the list, and click Properties. The VLAN QoS CIR Weight Profile, Site form opens.
2. Configure the parameters.
3. Save your changes and close the form.

5

Click on the Forwarding Classes tab to view the mapping of FCs to queues. The mapping is not user-configurable.

6

Configure queues.

1. Click on the Queues tab. Eight default queues are displayed.
2. Choose a queue and click Properties. The VLAN QoS Queue form opens.
3. Configure the required parameters on the General tab.
4. Select a 7250 SROS Queue Management policy in the Queue Management policy panel. For information about 7250 SROS Queue Management policies, see 50.53 “To configure a 7250 SROS Queue Management policy” (p. 1593).
5. Click on the CIR/PIR tab and configure the required parameters.

   CIR weight profiles selected in this step are configured in Step 4.
6. Save your changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.55 To configure a scheduler policy

50.55.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Scheduler from the NFM-P main menu. The Scheduler Policies form opens.

2

Click Create or choose a policy and click Properties. The Scheduler Policy (Create|Edit) form opens.

3

Configure the required parameters and click Apply.

4

Click on the Schedulers tab and click Create. The Entry, Scheduler Policy (Create) form opens.

The scheduler defines bandwidth control that limits each child (other schedulers and queues) that are associated with the scheduler.

5

Configure the required parameters.

The Limit Unused Bandwidth parameter appears only when the Tier parameter is set to 1.

The Parent Scheduler panel and its parameters appear only when the Port Parent parameter is set to false.

The Tier parameter identifies the level of hierarchy with which a group of schedulers is associated. A parent is tier 1. Children are tier 2. Grandchildren are tier 3. You can create a tier 2 child scheduler without creating a parent tier 1 scheduler, and you can create a grandchild tier 3 scheduler without creating a child tier 2 scheduler.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.56 To create an Aggregation Scheduler

Note: You can create an Aggregation Scheduler when you have created two or more scheduler policies. See 50.13.3 “Hierarchical schedulers” (p. 1521) and 50.55 “To configure a scheduler policy” (p. 1596) in this chapter for more information.

50.56.1 Steps

1

Choose Manage→Service→Customers from the NFM-P main menu. The Manage Customers form opens.

2

Select a customer in the list and click Properties. The Customer (Edit) form opens.

3

Click on the Sites tab.

4

Select a site in the list and click Properties. The Site (Edit) form opens.

5

Click on the Aggregation tab and click Create. The Create Aggregation Scheduler step form opens.

6

Select a site in the list and click Next.

7

Configure the parameters in the Define Name step and click Next.

8

Select the assignment scope and click Next.

9

Select a port, card, or FPE and click Next.

10

Select an ingress and egress scheduler policy and click Next.

11

Configure the parameters in the Select Egress Aggregate Rate Limit step and click Next.

12

If you are creating an Aggregation Scheduler for an IOM3 port, go to Step 13 . Otherwise go to Step 14 .

13

Select an ingress and egress scheduler policy and click Next.

14

Select a time of day suite and click Finish.

The Aggregation Scheduler is listed in the Aggregation tab of the Site (Edit) form. It is also listed in the Aggregation Schedulers tab of the properties form for the specific object you assigned in Step 12 , either a card, port (including LAGs), or FPE.

15

Save your changes and close the forms.

END OF STEPS

50.57 To configure a port scheduler policy

50.57.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Port Scheduler from the NFM-P main menu. The Port Scheduler Policies form opens.

2

Click Create or choose a policy and click Properties. The Port Scheduler (Create|Edit) form opens.

3

Configure the required parameters.

The Maximum Rate parameter is configurable when the MAX check box is deselected.

When the HQoS Algorithm parameter is set to the Above Offered Allowance Control option, the policy uses values configured in an Advanced Configuration policy; see 50.63 “To configure an Advanced Configuration policy” (p. 1606). The required Advanced Configuration policy must be distributed to the NE along with the Port Scheduler policy.

4

Click on the Group tab and click Create. The Port Scheduler Group (Create) form opens.

5

Configure the required parameters.

If you set the Rate Type parameter to kbps, the PIR (kbps) and CIR (kbps) parameters are configurable when the associated MAX check box is deselected.

6

Save your changes.

7

Click on the Level tab on the Port Scheduler (Create) form to configure the required levels.

Note: Contiguous mapping of levels to a group is enforced. Priority levels cannot be added to a group unless the resulting set of priority levels is contiguous. When a level is not explicitly mapped to any group, it maps directly to the root of the port scheduler at its own priority.

8

Associate a port scheduler group to a level:

1. Click Select in the Level panels to choose a port scheduler group from the list in the Select Group form. The port scheduler group must be preexisting on the port scheduler policy (as created in Step 4 to Step 6).
2. Configure the Weight in group parameter as required.

9

Configure the Rate Type, PIR, and CIR parameters for each level required.

The PIR and CIR parameters are configurable when the associated MAX check box is deselected.

10

Configure the required parameters in the Orphan panel.

11

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.58 To configure an HSMDA scheduler policy

50.58.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→HSMDA Scheduler from the NFM-P main menu. The HSMDA Scheduler Policies form opens.

2

Click Create or choose a policy and click Properties. The HSMDA Scheduler Policy (Create) form opens.

3

Configure the required general parameters.

4

Click on the HSMDA Group Rates tab and configure the required parameters for Group 1 and Group 2.

5

Click on the HSMDA Scheduler Classes tab and configure the parameters for classes 1 to 8.

The Group parameter must be configured with the same group (Group 1 or Group 2) for up to three sequential classes for each group. For example, when you configure Class 1 as Group 2, you can configure Class 2 and Class 3 as Group 2, but not as Group 1.

The Rate (Mbps) and Burst Limit (bytes) parameters are configurable when the Group parameter is set to None. The Weight parameter is configurable when the Group parameter is set to a value other than None.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.59 To configure an HSMDA WRR policy

50.59.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→HSMDA WRR Policy from the NFM-P main menu. The HSMDA WRR Policies form opens.

2

Click Create or choose a policy and click Properties. The HSMDA WRR Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.60 To configure a 7210, 7250 and 1830 Port Scheduler policy

50.60.1 Purpose

For 7210 SAS and 1830 PSS NEs, a 7210, 7250 and 1830 Port Scheduler policy defines the scheduling mode (algorithm type) used to determine the forwarding priority for egress queues. When a weighted algorithm is chosen, the policy allows the configuration of relative weights for up to eight queues.

The policy affects access egress and network (or L2 uplink) egress traffic, and is assigned to ports in any mode. For policy assignment on 7210 SAS NEs, see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636). For policy assignment on 7250 IXR NEs, see 16.24 “To configure Ethernet ports” (p. 599).

The NFM-P provides a default 7210, 7250 and 1830 Port Scheduler policy. The default policy is applied unless a user-configured port scheduler policy is explicitly assigned. The default policy specifies the Strict scheduler mode and cannot be modified. Strict priority means that queues are serviced based on their queue ID number, with higher numbers serviced first.

The 7210, 7250 and 1830 Port Scheduler policies are identified by the policy name and description; policy ID numbers are not used.

The 7210, 7250 and 1830 Port Scheduler policy is supported on the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-E, 7210 SAS-M, 7210 SAS-S, 7210 SAS-Sx, and 7210 SAS-T.

Support for specific parameters and features varies, depending on the chassis type. Local definitions of a policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

Note: For 7210 SAS devices that do not support 7210, 7250 and 1830 Port Scheduler policies, egress scheduling is defined by the port bandwidth, and by user-configured settings in access egress and network queue policies, port properties, and SAP configuration. See 70.15 “Sample QoS configuration on the 7210 SAS” (p. 1954) in Chapter 70, “Service management and QoS” .

The maximum bandwidth (line rate) for a port is either system-defined, or is configured using port level egress rate limiting.

50.60.2 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→ 7210, 7250 and 1830 Port Scheduler from the NFM-P main menu. The 7210, 7250 and 1830 Port Scheduler Policies form opens.

2

Click Create or choose a policy and click Properties. The Port Scheduler, Global Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

If you selected a value of Strict or RoundRobin for the Mode parameter, go to Step 6 .

5

If you selected a value of WeightedRoundRobin or WeightedDeficitRoundRobin for the Mode parameter, assign weight values to queues.

1. Click on the Queue Weights tab.
2. Configure the required parameters.
3. Go to Step 6.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.61 To configure a 7250 SROS Port QoS policy

50.61.1 Purpose

A 7250 SROS Port QoS policy defines scheduler parameters for egress queues. The policy allows configuration of WRR weights, CIR and PIR values, and the assignment of 7250 SROS Queue Management policies.

The 7250 SROS Port QoS policy also allows viewing of the mapping of forwarding classes to queues. FCs are mapped to queues by default; the mapping is not user-configurable.

For supporting NEs, the policy allows configuration of a Packet Byte Offset value. See the NE documentation for more information about PBO on the 7250 IXR.

7250 SROS Port QoS policies are assigned to ports, using the Policies tab of the port properties form; see 16.24 “To configure Ethernet ports” (p. 599).

7250 SROS Port QoS policies are identified by the policy name; policy ID numbers are not used.

A default 7250 Port QoS policy is applied if no user-configured policy is explicitly used. The default policy cannot be deleted or modified.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.61.2 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→7250 SROS Port QoS from the NFM-P main menu. The 7250 SROS Port QoS form opens.

2

Click Create, or choose a policy and click Properties. The 7250 SROS Port QoS Policy, (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

For supporting NEs, you can configure the Packet Byte Offset parameter when QoS Packet Byte Offset is globally enabled in NE properties; see 12.67 “To globally enable or disable Packet Byte Offset on a 7250 IXR” (p. 396).

4

Click on the Forwarding Classes tab to view the mapping of FCs to queues. The mapping is not user-configurable.

5

Click on the Queues tab to configure egress queues.

1. Choose a queue from the list and click Properties. The Port QoS Queue, 7250 SROS Port QoS (Create|Edit) form opens.
2. Configure the required parameters on the General tab.

   For the WRR Weight parameter, ID and weight values are configured in Step 6.
3. Select a 7250 SROS Queue Management policy. For information about 7250 SROS Queue Management policies, see 50.53 “To configure a 7250 SROS Queue Management policy” (p. 1593).
4. Click on the CIR/PIR tab and configure the required parameters.
5. Save your changes and close the form.

6

Click on the WRR Weights tab to configure multicast and unicast WRR weights for queues.

1. Choose a WRR Weight ID from the list and click Properties. The Port QoS WRR Weights, Site (Create|Edit) form opens.
2. Configure the required parameters.

   WRR IDs and weights configured in this step are used in Step 5 .
3. Save your changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.62 To configure a policer control policy

50.62.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Policer Control from the NFM-P main menu. The Policer Control Policies form opens.

2

Click Create or choose an existing policer control policy and click Properties. The Policer Control Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

Click on the Arbiter tab.

1. Click Create. The Arbiter (Create) form opens.
2. Configure the required parameters.

   The Parent Arbiter parameter is configurable when the Tier parameter is set to 2.
3. Click OK. The Arbiter (Create) form closes.

5

Optionally, click on the Graphical Arbiters tab to view a graphical display of the arbiter hierarchy within the policer control policy. You can perform any of the following steps:

a. Right-click on the grid and select New Arbiter from the drop-down menu. The Arbiter (Create) form opens. You can create a new arbiter, as described in Step 4 .

b. Right-click on an arbiter object and select Edit from the drop-down menu. The Arbiter (Edit) form opens. You can change the arbiter, as described in Step 4 .

c. Right-click on an arbiter and select Remove from the drop-down menu. The arbiter object is removed from the hierarchy.

d. Click on a Tier 2 arbiter object and draw a line to a Tier 1 arbiter object. The Tier 1 arbiter becomes the parent of the Tier 2 arbiter.

6

Click on the Tree tab to view a hierarchical display of the policer control policy properties and sub components.

7

Click on the Priority Level tab to configure the MBS contribution for eight priority levels. For each level, configure the Cumulative MBS Contribution and Fixed MBS contribution parameters.

8

Click Apply.

9

Click Distribute to distribute the policy to NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

10

Click OK to save the policy and close the forms.

END OF STEPS

50.63 To configure an Advanced Configuration policy

50.63.1 Purpose

The Advanced Configuration policy allows for additional configurations to queues and policers in the access ingress/egress policies and the ingress/egress queue group templates. This extra configuration is subsequently applied by assigning the Advanced Configuration policy to a queue or policer.

The Advanced Configuration policy also provides bandwidth distribution values used by Port Scheduler policies to control the HQoS algorithm.

50.63.2 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Advanced Configuration Policy from the NFM-P main menu. The Advanced Configuration Policy (Create) form opens.

2

Configure the required parameters.

3

Click on the Measured Offered Rate tab.

1. Configure the general parameters as required.
2. Configure the required parameters in the Measured Offered Rate Increase panel.

   The Value parameter will either be expressed as a kbps rate or a percentage, based on your choice of Rate Type.

The Default parameter causes the Value parameter to use whatever value is set for the entity in its own configuration form.

3. Configure the required parameters in the Granularity panel.

   The Value parameter will either be expressed as a kbps rate or a percentage, based on your choice of Rate Type.

   The Default parameter causes the Value parameter to use whatever value is set for the entity in its own configuration form.

4. Configure the required parameters in the Max Decrement panel.

   The Value parameter will either be expressed as a kbps rate or a percentage, based on your choice of Rate Type.

   The Default parameter causes the Value parameter to use whatever value is set for the entity in its own configuration form.

5. Click OK. The Arbiter (Create) form closes.

4

Click on the Bandwidth Distribution tab.

1. Configure the required general parameters.
2. Configure the required parameters in the Above Offered Cap panel.

   The Value parameter will either be expressed as a kbps rate or a percentage, based on your choice of Rate Type.

   The Default parameter causes the Value parameter to use whatever value is set for the entity in its own configuration form.

3. Configure the required parameters in the Granularity panel.

   The Value parameter will either be expressed as a kbps rate or a percentage, based on your choice of Rate Type.

   The Default parameter causes the Value parameter to use whatever value is set for the entity in its own configuration form.

4. Configure the required parameters in the Above Offered Allowance panel.

   The Above Offered Allowance settings are used by a Port Scheduler policy when its HQoS Algorithm parameter is set to the Above Offered Allowance Control option. See 50.57 “To configure a port scheduler policy” (p. 1599).

5

Click Apply.

A number of other tabs become active. These allow you to view where the policy is used, once it has been assigned to the various applicable entities.

6

Click on the General tab.

7

Click on the Switch Mode button beside the Configuration Mode parameter to distribute the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

8

Click OK to save the policy and close the form.

END OF STEPS

50.64 To configure a Hardware Aggregate Shaper Scheduler policy

50.64.1 Purpose

The Hardware Aggregate Shaper Scheduler policy improves reaction time and downstream bursts for aggregate-rate enforcement at subscriber level.

It improves on other deficiencies of traditional HQoS, such as potential for underrun during weight-based BW distribution between different queues.

50.64.2 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Hardware Aggregate Shaper Scheduler from the NFM-P main menu. The Hardware Aggregate Shaper Scheduler form opens.

2

Click Create or choose an existing Hardware Aggregate Shaper Scheduler policy and click Properties. The Hardware Aggregate Shaper Scheduler Policy (Create) form opens.

3

Configure the required parameters on the General tab.

4

In the Group tab, click Create to enter the Displayed Name parameter.

5

Click OK to save changes.

6

Click the Scheduler Class tab to configure the scheduler class entries.

1. Choose a scheduler class from the list and click Properties. The Scheduler Class, Hardware Aggregate Shaper Scheduler Policy (Create) form opens.
2. Click Select to associate group name to the selected scheduler class.

7

Click OK to save the policy and close the form, or click Apply to save the policy.

See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.65 To configure a port policy

50.65.1 Steps

1

Choose Policies→QoS→SROS QoS→Scheduler→Port Policy from the NFM-P main menu. The Port Policies form opens.

2

Click Create or choose an existing port policy and click Properties. The Port Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

Click Select in the Port Scheduler Policy panel and choose a Port Scheduler Policy, then click OK.

5

Click Apply.

6

Click Distribute to manually distribute the policy locally to devices. Policies are also automatically distributed to devices when they are used by resources on the device. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

7

Click OK to save the policy and close the forms.

END OF STEPS

50.66 To configure an HSMDA pool policy

Note: To associate an HSMDA buffer pool with an MDA, see 15.78 “To configure an MDA” (p. 536) .

50.66.1 Steps

1

Choose Policies→QoS→SROS QoS→Buffer Pool→HSMDA Pool from the NFM-P main menu. The Manage HSMDA Pool Policies form opens.

2

Click Create or choose an HSMDA pool policy and click Properties. The HSMDA Pool Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click on the Root Pools tab and configure the Allocation Weight parameter for root pools 1 to 8.

The Allocation Weight parameter is configurable when the Default check box is deselected.

5

Click on the Class Pools tab and configure the required parameters for root pools 1 to 8.

The Allocation Percent parameter is configurable when the Default check box is deselected.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.67 To configure a named buffer pool policy

Note: To associate a Named Buffer Pool to an MDA, see 15.78 “To configure an MDA” (p. 536) .

50.67.1 Steps

1

Choose Policies→QoS→SROS QoS→Buffer Pool→Named Buffer Pool from the NFM-P main menu. The Named Buffer Pool Policies form opens.

2

Click Create or choose an existing named buffer pool policy and click Properties. The Named Pool Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

To configure Q1 pools, perform Step 3 to Step 7 in 50.68 “To configure Q1 pools” (p. 1610) .

5

Click Apply.

6

Click Distribute to manually distribute the policy locally to devices. Policies are also automatically distributed to devices when they are used by resources on the device. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

7

Click OK to save the policy and close the forms.

END OF STEPS

50.68 To configure Q1 pools

50.68.1 Before you begin

You must configure a Q1 pool before you can assign the pool to an access ingress, access egress, network queue or shared queue policy.

50.68.2 Steps

1

Choose Policies→QoS→SROS QoS→Buffer Pool→Named Buffer Pool from the NFM-P main menu. The Manage Named Buffer Pool Policies form opens.

2

Click Create or choose a pool policy and click Properties. The Named Pool Policy (Create | Edit) form opens.

3

Click on Q1 Pool tab. A list of configured Q1 Pools is displayed.

4

Click Create or choose a Q1 pool and click Properties. The Q1 pool (Create | Edit) form opens.

5

Configure the required parameters.

6

Click Select beside the Displayed Name parameter and select a slope policy.

7

Save your changes and close the forms.

END OF STEPS

50.69 To configure a HS Attachment policy

50.69.1 Steps

1

Choose Policies→QoS→SROS QoS→HS QoS→HS Attachment policy from the NFM-P main menu. The HS QoS – Attachment Policy form opens.

2

Click Create or choose an existing HS attachment policy and click Properties. The HS Attachment Policy, Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click Apply to save the policy. Confirm the action.

5

As required, click on the Queue tab, select a default queue entry and click Properties to modify the parameters. By default, all queue entries have the same parameter settings applied to them. The HS Attachment Queue form opens.

6

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Attachment Policy, Global Policy (Create | Edit form reappears.

7

As required, click on the WRR Group tab, select an entry and click Properties to modify the parameters. By default, all entries have the same parameter settings applied to them. The HS Attachment WRR Group form appears.

8

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Attachment Policy, Global Policy (Create | Edit form reappears.

9

Click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

50.70 To configure a HS Pool policy

50.70.1 Steps

1

Choose Policies→QoS→SROS QoS→HS QoS→HS Pool policy from the NFM-P main menu. The HS QoS – Pool Policy form opens.

2

Click Create or choose an existing HS pool policy and click Properties. The HS Pool Policy, Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

As required, click on the Root-Tier Pool tab, select a default entry and click Properties to modify the parameters. By default, all entries have the same parameter settings applied to them. The HS Root Tier Pool form opens.

5

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Pool Policy, Global Policy (Create | Edit) form reappears.

6

As required, click on the Mid-Tier Pool tab, select a default entry and click Properties to modify the parameters. By default, all entries have the same parameter settings applied to them. The HS Mid Tier Pool form opens.

7

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Pool Policy, Global Policy (Create | Edit) form reappears.

8

Click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

9

As required, associate the HS Port policy to the Forwarding Plane at the IOM level on a 7750 SR node:

1. From the equipment view, navigate to a Forwarding Plane on a 7750 SR node (Node→Shelf→Card Slot→Forwarding Plane) and choose Properties. The Forwarding Plane [Edit] form opens.
2. On the General tab, choose a HS Pool policy in the HSQ panel to associate the policy with the Forwarding Plane and click OK.
3. Save your changes and close the form.

END OF STEPS

50.71 To configure a HS Port Pool policy

50.71.1 Steps

1

Choose Policies→QoS→SROS QoS→HS QoS→HS Port Pool policy from the NFM-P main menu. The HS QoS – Port Pool policy form opens.

2

Click Create or choose an existing HS port pool policy and click Properties. The HS Port Pool Policy, Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

As required, click on the Standard Port Class Pools tab, select a default entry and click Properties to modify the parameters. By default, all entries have the same parameter settings applied to them. The HS Port STD Class Port form opens.

5

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Port Pool Policy, Global Policy (Create | Edit) reappears.

6

As required, click on the Alternate Port Class Pools tab, select a Class pool entry and click Properties to modify the default parameters. By default, all entries have the same parameter settings applied to them. The HS Port ALT Class Port form opens.

7

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Port Pool Policy, Global Policy (Create | Edit) reappears.

8

Click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

9

As required, associate the HS Port Pool policy to a HSQ IOM4-e-HS MDA port on a 7750 SR node:

1. From the equipment view, navigate to a HSQ IOM4-e-HS MDA port on a 7750 SR node (Node→Shelf→Card Slot→Daughter Card→[Supported] Port) and choose Properties. The Physical Port [Edit] form opens.
2. On the General tab, choose a HS Port Pool policy in the HSQ panel to associate the policy with the port and click OK.
3. Save your changes and close the form.

END OF STEPS

50.72 To configure an FP Resource policy

50.72.1 Before you begin

Forwarding Plane Resource policies allow custom allocation of card slot resources to queues. By default, ingress and egress queues are each allocated 50% of the queue resources on a forwarding plane. FP Resource policies allow you to configure custom ingress queue resources in a range of 4% to 97%. The corresponding percentage remaining is allocated to egress queues.

FP Resource policies are assigned to forwarding planes on the card slot properties form of supporting NEs; see 15.59 “To assign an FP Resource policy to a forwarding plane” (p. 519)

50.72.2 Steps

1

Choose Policies→QoS→SROS QoS→FP Resource Policy from the NFM-P main menu. The FP Resource Policy form opens.

2

Click Create or choose an existing FP Resource policy and click Properties. The FP Resource Policy, Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

Save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

End of steps

50.73 To configure a HS Scheduler policy

50.73.1 Steps

1

Choose Policies→QoS→SROS QoS→HS QoS→HS Scheduler policy from the NFM-P main menu. The HS QoS – Scheduler Policy form opens

2

Click Create or choose an existing HS scheduler policy and click Properties. The HS Scheduler Policy, Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click Apply to save the policy. Confirm the action.

5

As required, click on the Scheduling Class tab, select a default entry and click Properties to modify the default parameters. By default, all entries have the same parameter settings applied to them. The HS Scheduling Classes form opens.

6

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Scheduler Policy, Global Policy (Create | Edit) form reappears..

7

As required, click on the Scheduling Group tab, select a default entry and click Properties to modify the default parameters. By default, all entries have the same parameter settings applied to them. The HS Scheduling Group form appears.

8

Modify the parameters on the form as required and click OK to save the changes and close the form. The HS Scheduler Policy, Global Policy (Create | Edit) form reappears.

9

Click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

10

As required, associate the HS Port Pool policy to a HSQ IOM4-e-HS MDA port on a 7750 SR node:

1. From the equipment view, navigate to a HSQ IOM4-e-HS MDA port on a 7750 SR node (Node→Shelf→Card Slot→Daughter Card→[Supported] Port) and choose Properties. The Physical Port [Edit] form opens.
2. On the General tab, choose a HS Scheduler policy in the HSQ panel to associate the policy with the port and click OK.
3. Save your changes and close the form.

END OF STEPS

50.74 To configure a queue group ingress template policy

50.74.1 Steps

1

Choose Policies→QoS→SROS QoS→Queue Group→Ingress Template from the NFM-P main menu. The Ingress Queue Group Template Policies form opens.

2

Click Create or choose an existing ingress queue group template policy and click Properties. The Ingress Queue Group Template (Create | Edit) form opens.

3

Configure the required parameters.

4

Click on the Queues tab.

5

Click Create or choose an existing queue and click Properties. The Ingress Queue (Create | Edit) form opens.

6

Configure the required parameters.

7

Configure the Named Buffer Pool parameter:

1. Click on Select adjacent to the Named Buffer Pool parameter. The Named Buffer Pool form opens.

   Note:

   To configure the Named Buffer Pool parameter, you must first create a Named Buffer Pool policy. See 50.67 “To configure a named buffer pool policy” (p. 1610) .
2. Select the required policy and click OK.

8

Configure the Scheduler parameter:

1. Click on Select adjacent to the Scheduler parameter. The Schedulers form opens.
2. Select the required scheduler and click OK.
3. Configure the required parameters in the Scheduler Association panel.

9

Configure the Advanced Configuration Policy Displayed Name parameter:

1. Click Select in the Advanced Configuration Policy panel. The Select Advanced Configuration Policy search form opens.
2. Select the required policy and click OK.

10

Configure the CIR/PIR/FIR parameters:

1. Click on the CIR/PIR/FIR tab.
2. Configure the required parameters.

   Ensure that the CIR value is lower than the PIR value.

11

Configure the Burst Size parameters:

1. Click on the Burst Size tab.
2. Configure the required parameters.

   The parameters are configurable when the Default check box above each is deselected.

   Ensure that the Committed Burst Size (kb) value is lower than the Maximum Burst Size (bytes) value.
3. Click OK.

12

Click on the Policers tab.

13

Click Create or choose an existing policer and click Properties. The Ingress Policer (Create | Edit) form opens.

14

Configure the required parameters.

15

Configure the Advanced Configuration Policy Displayed Name parameter:

1. Click Select in the Advanced Configuration Policy panel. The Select Advanced Configuration Policy search form opens.
2. Select the required policy and click OK.

16

Configure the CIR/PIR parameters:

1. Click on the CIR/PIR tab.
2. Configure the required parameters.

17

Configure the Burst Size parameters:

1. Click on the Burst Size tab.
2. Configure the required parameters.

   The parameters are configurable when the Default check box above each is deselected.
3. Click OK.

18

Click OK to save the policy and close the form.

END OF STEPS

50.75 To configure a queue group egress template policy

50.75.1 Steps

1

Choose Policies→QoS→SROS QoS→Queue Group→Egress Template from the NFM-P main menu. The Egress Queue Group Template Policies form opens.

2

Click Create or choose an existing egress queue group template policy and click Properties. The Egress Queue Group Template (Create | Edit) form opens.

3

Configure the required parameters.

4

Enable the Queues HQoS Manageable parameter if you intend to allow Hierarchical QoS management of the queues in this template.

5

To configure HSMDA queues, go to Step 6. To associate an HS Attachment Policy, go to Step Step 9. Otherwise, go to Step 14.

6

Configure the Packet Offset (bytes) and Low Burst Max Class parameters.

7

Click Select in the WRR Policy panel to associate a WRR Policy. The Select WRR Policy form opens.

8

Select the required WRR Policy and click OK.

9

Click Select in the HS panel to associate an alternate HS Attachment Policy other than the default policy. The Select HS Attachment Policy – Egress Queue Group Template Policy form opens

10

Select the required HS Attachment Policy and click OK.

11

Click on the HSMDA Queues tab.

12

Edit the HSMDA queues as required:

1. Select a queue from the list and click Properties.
2. Configure the required parameters.
3. Click Select and choose the required Slope Policy, then click OK.
4. Click on the PIR/Burst Size tab.

5. Configure the required parameters.

   The sliders can be configured when the checkboxes above them are disabled.

6. Click OK. The selected HSMDA queue is revised with your configuration changes.

13

Repeat Step 6 to Step 12 for any other required HSMDA queues.

14

Click on the Queues tab.

15

Click Create or choose an existing queue and click Properties. The Egress Queue (Create | Edit) form opens.

16

Configure the required parameters.

The Level, Weight, CIR Level, and CIR Weight parameters are configurable only if the Port Parent parameter is set to True.

17

Configure the Named Buffer Pool parameter:

1. Click on Select adjacent to the Named Buffer Pool parameter. The Named Buffer Pool form opens.

   Note:

   To configure the Named Buffer Pool parameter, you must first create a Named Buffer Pool policy. See 50.67 “To configure a named buffer pool policy” (p. 1610) .

2. Select the required policy and click OK.

18

Configure the Scheduler parameter:

1. Click on Select adjacent to the Scheduler parameter. The Schedulers form opens.

   Note:

   The Scheduler parameter is configurable only if you set the Port Parent parameter to False.

2. Select the required scheduler and click OK.
3. Configure the required parameters in the Scheduler Association panel.

19

Configure the Advanced Configuration Policy Displayed Name parameter:

1. Click Select in the Advanced Configuration Policy panel. The Select Advanced Configuration Policy search form opens.
2. Select the required policy and click OK.

20

Configure the WRED Queue parameters:

1. Select the Use WRED Queue parameter.
2. Click Select adjacent to the Displayed Name parameter. The Select WRED Queue Slope Policy form opens.

   Note:

   To configure a WRED Queue Slope policy, you must first create the policy. See 50.49 “To configure a WRED slope policy” (p. 1587).
3. Specify a filter to search for existing policy.
4. Select the required WRED Slope policy and click OK.

21

Configure the parameters in the Latency panel.

Enabling the Dynamic MBS parameter permits the dynamic modification of the MBS size of a queue in order to maintain the maximum latency for traffic in the queue, based on the queue’s admin PIR and configured MBS. The Queue Delay parameter allows you to configure the target queue delay for packets forwarded through the queue. It is used to determine the related queue parameters based on the PIR of the queue. If the No checkbox is selected, the determination of the queue parameters based on the queue delay is disabled.

22

Configure the CIR/PIR parameters:

1. Click on the CIR/PIR tab.
2. Configure the required parameters.

   The choice of CIP and PIR parameter definitions depends on the Rate Type parameter setting on the General tab.

23

Configure the Burst Size parameters:

1. Click on the Burst Size tab.
2. Configure the required parameters.

   The parameters are configurable when the Default check box above each is deselected.

24

Configure the HSQ parameters if required. You can only configure HSQ on default queue entries 1 to 8.

1. Click on the HSQ tab.
2. Configure the required parameters.
3. Click Select adjacent to the HS WRED Slope Policy parameter. The Select HS WRED slope Policy form opens.

   Note: To configure the HS WRED Slope Policy parameter, you must first create a HS WRED Slope policy. See 50.49 “To configure a WRED slope policy” (p. 1587).
4. Select the required policy and click OK.
5. Click OK. The Egress Queue form closes.

25

Click on the Forwarding Classes tab.

26

Click Create or choose an existing forwarding class and click Properties. The Forwarding Class (Create | Edit) form opens.

27

Configure the required parameters.

28

Click OK.

29

Click on the Policers tab.

30

Click Create or choose an existing policer and click Properties. The Egress Policer (Create | Edit) form opens.

31

Configure the required parameters.

32

Configure the Advanced Configuration Policy Displayed Name parameter:

1. Click Select in the Advanced Configuration Policy panel. The Select Advanced Configuration Policy search form opens.
2. Select the required policy and click OK.

33

Configure the CIR/PIR parameters:

1. Click on the CIR/PIR tab.
2. Configure the required parameters.

34

Configure the Burst Size parameters:

1. Click on the Burst Size tab.
2. Configure the required parameters.

   The parameters are configurable when the Default check box above each is deselected.
3. Click OK. The Egress Policer (Create | Edit) form closes.

35

Configure the HSQ parameters if required to configure HS WRR Groups.

1. Click on the HS WRR Groups tab.
2. Choose an existing HS WRR group and click Properties. The HS WRR Group form opens.

   Note: You can only modify the existing WRR Group parameters; you cannot create new WRR Groups.
3. As required, modify the parameters.
4. Save your changes and close the form.

36

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.76 To configure a queue group redirect list policy

50.76.1 Purpose

You can create a Queue Group Redirect List policy to provide QoS control to map IPv4 and IPv6 traffic using a single VXLAN or VXLAN GPE VNI to an ingress access queue group instance or a port access egress queue group instance. See 50.19 “Queue Group policies” (p. 1526) for more information.

After you create a queue group redirect list policy, you must release and distribute the policy, then associate the policy with an egress/ingress physical interface (SAP) that supports IES or VPRN services.

50.76.2 Steps

1

Choose Policies→QoS→SROS QoS→Queue Group→Redirect List from the NFM-P main menu. The Redirect List Queue Group Policies form opens.

2

Click Create or choose an existing queue group redirect list policy and click Properties. The Queue Group Redirect List Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

Click on the Match tab and click Create. The Redirect List Match, Queue Group Redirect List Policy form opens.

5

Configure the required parameters.

6

Click OK to save the changes and close the form. The Queue Group Redirect List Policy (Create | Edit) form reappears.

7

Click OK to save the changes and close the form. The Redirect List Queue Group Policies form reappears.

Release and distribute the queue group redirect list policy to the NEs

8

To release and distribute the policy, click Search, select the newly created policy and click Properties. The Queue Group Redirect List Policy (Edit) form opens.

9

Perform the required steps in 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy.

Associate the queue group redirect list policy to an IES or VPRN ingress SAP

10

Create a queue group ingress template policy as described in 50.74 “To configure a queue group ingress template policy” (p. 1617) .

11

Perform the required steps in 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy.

12

Attach the newly created queue group ingress template policy to the forwarding plane of a card slot on the Access Ingress Queue Group tab; see 15.65 “To configure an ingress queue group on a forwarding plane” (p. 523).

13

Perform the following actions on the QoS tab associated with an ingress L3 access interface (SAP) that supports an IES or VPRN service; see78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477).

1. Click Select to choose the queue group ingress template policy created in Step 10.
2. Click Select to choose the queue group redirect list policy created in Step 1 to Step 7.
3. Configure any other required parameters on the form and click OK to save your changes and close the form.

Associate the queue group redirect list policy to an IES or VPRN egress SAP

14

Create a queue group egress template policy as described in 50.75 “To configure a queue group egress template policy” (p. 1619) .

15

Perform the required steps in 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy.

16

Attach the newly created queue group egress template policy to the physical port on the Access Egress Queue Group tab; see 16.37 “To add a queue group to an Ethernet port” (p. 627).

17

Perform the following actions on the QoS tab associated with an egress L3 access interface (SAP) that supports an IES or VPRN service; see78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477).

1. Click Select to choose the queue group ingress template policy created in Step 14.
2. Click Select to choose the queue group redirect list policy created in Step 1 to Step 7.
3. Configure any other required parameters on the form and click OK to save your changes and close the form.

END OF STEPS

50.77 To configure a 7705 SAR fabric profile

50.77.1 Steps

1

Choose Policies→QoS→SROS QoS→Fabric from the NFM-P main menu. The Fabric Profiles form opens.

2

Click Create or choose an existing fabric policy and click Properties. The Fabric (Create | Edit) form opens.

3

Configure the required parameters.

4

Click on the Rates tab. Perform one of the following:

a. If you chose Aggregate Mode in Step 3 , configure the parameters:
   - Aggregate Rate (kbps)
   - Unshaped SAP CIR (kbps)

b. If you chose Destination Mode in Step 3 , configure the Rate To MDA in kbps for each daughter card slot to set the fabric shaping rate to each daughter card.

   You can also configure the MultiPoint Rate (kbps) parameter for distribution to all of the daughter cards.

5

Click Apply. The form refreshes to enable the other tabs.

6

Click Distribute to distribute the policy to NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

7

Click OK to save the policy and close the form.

END OF STEPS

50.78 To configure a 7705 SAR security queue policy

50.78.1 Purpose

The security queue policy defines the queue type (Best Effort or Expedited) applied to 7705 SAR ingress traffic to deter DoS-like attacks from overwhelming the control plane, while ensuring that

critical control traffic such as signaling is always serviced in a timely manner. These queues are either fixed use (each queue handles a certain type of traffic which is not user-configurable) or fixed configuration (each queue is user-configurable; you can configure the CIR/PIR rate, burst size, and buffering capacity of each queue).

50.78.2 Steps

1

Choose Policies→QoS→SROS QoS→Security Queue from the NFM-P main menu. The Security Queue Policies form opens.

2

Click Create. The Security Queue Policy, Global Policy (Create) form opens.

3

Configure the general parameters and click Apply. The form refreshes to enable the other tabs.

4

Click on the Queue tab. Two default queues appear - Best Effort and Expedited.

5

Choose a default queue and click Properties.

6

Configure the MIN and MAX for the CIR (kbps), PIR (kbps), Committed Burst Size (KB), Maximum Buffer Size (KB), and Buffer Size (%), using the sliders.

The sliders are operational when the checkboxes are disabled.

7

Click Apply to save the policy, or click OK to save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

End of steps

50.79 To configure a 7705 SAR shaper policy

50.79.1 Purpose

All 7705 SAR devices have a default shaper QoS policy and an associated shaper group that cannot be modified. If you create a new shaper QoS policy, you can modify the associated default shaper group.

50.79.2 Steps

1

Choose Policies→QoS→SROS QoS→Shaper from the NFM-P main menu. The Shaper Policies form opens.

2

Click Create or choose an existing Shaper policy and click Properties. The Shaper Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

Click OK and confirm.

5

Click Distribute to distribute the policy to NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Select the newly-created policy in the Shaper Policies form and click Distribute to distribute the policy to NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

7

To modify a default shaper group, choose a newly-created policy from the Shaper Policies form and click Properties. The Shaper Policy, Global Policy (Edit) form opens.

8

Click on the Shaper Group tab.

9

Choose the default Shaper group and click Properties. The Shaper Group (Edit) form opens.

10

Configure the Description parameter and the MIN and MAX values for the PIR (kbps) and CIR (kbps) using the sliders.

The sliders are operational when the checkboxes are disabled.

11

Click OK to save your changes and close the forms.

END OF STEPS

50.80 To configure a 7210 remarking policy

50.80.1 Purpose

A 7210 remarking policy maps forwarding classes and profile states to dot1p, DE, DSCP, and LSP-EXP bit values for access egress or network egress traffic. Policies are created as one of the following types:

- DOT1P
- DSCP
- LSP-EXP
- DOT1P-DSCP
- DOT1P-LSP-EXP SHARED
- DOT1P-LSP-EXP
- DSCP-LSP-EXP
- DOT1P-LSP-EXP-DSCP

A 7210 remarking policy can be assigned to:

- a 7210 and 1830 port access egress policy; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556) . The remarking policy must be of type DOT1P, DSCP, DOT1P-DSCP, or DOT1P-LSP-EXP SHARED.
- a 7210 SAP access egress policy; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558) . The remarking policy must be of type DOT1P, DSCP, DOT1P-DSCP, or DOT1P-LSP-EXP SHARED.
- a 7210 and 1830 network policy; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) . If the network policy is of port type, the remarking policy must be of type DOT1P, DSCP, DOT1P-DSCP, DOT1P-LSP-EXP, DOT1P-LSP-EXP-DSCP, or DSCP-LSP-EXP. If the network policy is of network interface type, the remarking policy must be of type LSP-EXP or DOT1P-LSP-EXP SHARED.

The NFM-P provides two default remarking policies:

- policy ID 1: DOT1P-DSCP type
- policy ID 2: DOT1P-LSP-EXP SHARED type

Policy ID 1 is the default for port-type network policies; policy ID 2 is the default for network interface-type network policies. Default policies cannot be deleted or modified. A default 7210 remarking policy is assigned unless a user-configured remarking policy is explicitly assigned. You can view the default mapping of FCs and profiles to dot1p, DSCP, and LSP-EXP bit values by clicking on the Forwarding Classes tab of the properties form for the default policy.

The 7210 remarking policy is supported on the 7210 SAS-K, 7210 SAS-Mxp, 7210 SAS-R, 7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T, and 7210 SAS-X. Support for specific parameters and features varies, depending on the chassis type. Local definitions of a policy do not necessarily support all configurations in the global policy. See the NE documentation for more information.

Note: To avoid synchronization and deployment errors when you create remarking policies using a CLI on the 7210 SAS-X, Nokia recommends that you create each remarking policy separately and allow at least three seconds between each policy creation.

50.80.2 Steps

1

Choose Policies→QoS→SROS QoS→Egress Remarking→7210 Remarking from the NFM-P main menu. The 7210 Remarking Policies form opens.

2

Click Create or choose an existing policy and click Properties. The 7210 Remarking Policy (Create| Edit) form opens.

3

Configure the required parameters on the General tab and click Apply.

Note: You cannot modify the Type parameter after the policy is created.

4

Configure the mapping of forwarding classes, profiles, and remark values:

1. Click on the Forwarding Classes tab.
2. Choose a forwarding class and click Properties. The 7210 Remark Forwarding Class form opens.
3. Configure the required parameters.

   The available parameters vary, depending on the remarking policy type and the settings for other parameters on the form.

   For DOT1P and DOT1P-DSCP remarking policy types, settings for the Dot1p and Force DE value parameters take effect when the Mark DE bits parameter is set to true.
4. Save your changes and close the form.

5

Click Distribute to distribute the policy to NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Click Apply to save the policy or click OK to save the policy and close the form.

END OF STEPS

50.81 To configure FC mapping policies for egress remarking on the 7250 IXR

50.81.1 Purpose

7250 SROS FC mapping policies for egress remarking associate Dot1p, DSCP, and LSP-EXP CoS values with forwarding classes and profiles. A separate policy is configured for each classification type (Dot1p, DSCP, or LSP EXP).

The policies are assigned to a 7250 SROS Egress Remarking policy; see 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634).

7250 SROS FC mapping policies are identified by the policy name; policy ID numbers are not used.

The NFM-P provides a default 7250 SROS FC Dot1p mapping policy. The default policy is not configurable, and cannot be deleted. The default policy is assigned to a 7250 SROS Egress Remarking policy when no user-configured policy is explicitly assigned.

7250 SROS FC mapping policies are not supported on all 7250 IXR chassis types and releases. For information about FC mapping policy support on the 7250 IXR, see the NE documentation.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.81.2 Steps

1

If required, configure a 7250 SROS FC Dot1p Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Egress Remarking→7250 SROS FC Dot1p Mapping from the NFM-P main menu. The 7250 SROS FC Dot1p Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS FC Dot1p Mapping (Create|Edit) form opens.
3. Configure the Displayed Name parameter on the General tab.
4. Click on the FC Classification tab.
5. Choose a forwarding class in the list and click Properties. The FC Dot1p form opens.
6. Configure the mapping of forwarding classes, profiles, Dot1p values, and DE values.

   Settings for the Force DE value parameter take effect when the Mark DE bit parameter is set to true.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional FCs, as required.
9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

2

If required, configure a 7250 SROS FC DSCP Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Egress Remarking→7250 SROS FC DSCP Mapping from the NFM-P main menu. The 7250 SROS FC DSCP Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS FC DSCP Mapping (Create|Edit) form opens.
3. Configure the Displayed Name parameter on the General tab.
4. Click on the FC Classification tab.
5. Choose a forwarding class in the list, or click Create. The FC DSCP form opens.
6. Configure the mapping of forwarding classes, profiles, and DSCP values.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional FCs, as required.
9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

3

If required, configure a 7250 SROS FC LSP EXP Mapping policy. Perform the following:

1. Choose Policies→QoS→SROS QoS→Egress Remarking→7250 SROS FC LSP EXP Mapping from the NFM-P main menu. The 7250 SROS FC LSP EXP Mapping form opens.
2. Click Create or choose an existing policy and click Properties. The 7250 SROS FC LSP EXP Mapping (Create|Edit) form opens.
3. Configure the Displayed Name parameter on the General tab.
4. Click on the FC Classification tab.
5. Choose a forwarding class in the list, or click Create. The FC LSP EXP form opens.
6. Configure the mapping of forwarding classes, profiles, and LSP-EXP values.
7. Save your changes and close the form.
8. Repeat substep 5 to substep 7 to configure additional FCs, as required.
9. Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

End of steps

50.82 To configure a 7250 SROS Remarking policy

50.82.1 Before you begin

A 7250 SROS Remarking policy maps forwarding classes and profile states to Dot1p, DSCP, DE, and LSP-EXP values for egress traffic.

The 7250 SROS Remarking policy is assigned to L2 access interfaces on VPLS and VLL services, to L3 access interfaces on IES and VPRN services, and to network interfaces. See the following procedures:

- for VLL: 76.42 “To assign ingress and egress QoS policies to a VLL L2 access interface” (p. 2181)
- for VPLS: 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340)
- for IES: 78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477)
- for VPRN: 79.89 “To assign ingress and egress QoS policies to a VPRN L3 access interface” (p. 2662)
- for L3 network interfaces: 27.17 “To create an L3 network interface on a routing instance” (p. 856)and 27.18 “To configure L3 network interfaces” (p. 863)

7250 SROS Remarking policies are identified by the policy name; policy ID numbers are not used.

The NFM-P provides three default 7250 SROS Remarking policies:

- access default, for L2 and L3 access interfaces
- network default, for network interfaces
- push-implicit-null-default, for PHP network interfaces

Note: The push-implicit-null-default policy is supported only for releases earlier than 22.2 R1 on the IXR nodes.

The default policies are not configurable. A default policy is assigned if no user-configured 7250 SROS Remarking policy is explicitly assigned.

For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.82.2 Steps

1

Choose Policies→QoS→SROS QoS→Egress Remarking→7250 SROS Remarking from the NFM-P main menu. The 7250 SROS Remarking form opens.

2

Click Create, or choose an existing policy and click Properties. The 7250 SROS Egress Remark Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

If you are configuring a policy for a 7250 IXR NE that supports SROS FC mapping policies for egress remarking, assign the FC mapping policies in the Properties panel, then go to Step 5.

For more information about SROS FC mapping policies, see 50.81 “To configure FC mapping policies for egress remarking on the 7250 IXR” (p. 1632).

4

Click on the Forwarding Classes tab to configure remarking values for FCs.

1. Choose a forwarding class and click Properties. The Egress Remark Forwarding Class form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.83 To configure a 7210 MPLS LSP-EXP Mapping policy

50.83.1 Purpose

A 7210 MPLS LSP-EXP Mapping policy defines a local mapping of forwarding classes to MPLS LSP-EXP bits for traffic ingressing the switch from an MPLS network. The policy is used primarily for an RSVP LSP with FRR one-to-one. For LDP LSPs, or when using FRR facility, the use of a single MPLS LSP-Exp Map policy for all IP interfaces is recommended.

The policy is assigned to a network policy of network interface type; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) . You can also assign the policy to an NE using the NE properties form. The 7210 MPLS LSP-EXP Mapping policy is called an MPLS LSP-EXP Profile when it is assigned to a network policy or NE.

The LDP Local FC parameter must be set to true in NE properties in order for 7210 MPLS LSP-EXP Mapping policy local definitions to take effect.

The NFM-P provides one default MPLS LSP-EXP Mapping policy (policy ID 1). The default policy cannot be modified or deleted. In the default policy, EXP bit values 0 and 2 are out-of-profile. The other six EXP bit values (1, 3, 4, 5, 6, 7) are in-profile.

The MPLS LSP-EXP Mapping policy is supported on the following MPLS-enabled 7210 SAS devices:

- 7210 SAS-M (in network mode)
- 7210 SAS-Mxp
- 7210 SAS-R

- 7210 SAS-S
- 7210 SAS-Sx
- 7210 SAS-T (in network mode)
- 7210 SAS-X

50.83.2 Steps

1

Set the LDP Local FC parameter to true.

Note: For 7210 SAS NEs, the LDP Local FC parameter is read-only and is set to true.

Perform the following:

1. Right-click on a 7210 SAS NE in the navigation tree and choose Properties. The Network Element (Edit) form opens.
2. Click on the LDP QoS Config tab and configure the LDP Local FC parameter, as required.
3. Save your changes and close the form.

2

Choose Policies→QoS→SROS QoS→Network→7210 MPLS LSP-EXP Mapping from the NFM-P main menu. The 7210 MPLS LSP-EXP Map Policies form opens.

3

Click Create or choose a policy and click Properties. The 7210 MPLS LSP-Exp Map, Global Policy (Create|Edit) form opens.

4

Configure the required parameters on the General tab.

5

Map profile states to MPLS LSP-EXP bits.

1. Click on the MPLS LSP-Exp Bits tab.
2. Select an LSP EXP bit and click Properties. The MPLS LSP-Exp Map Profile form opens.
3. Configure the Profile parameter.
4. Save your changes and close the form.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.84 To configure a 7210/7250 Dot1p classification policy

50.84.1 Purpose

A 7210/7250 Dot1p classification policy associates ingress dot1p values, forwarding classes, and profiles. The policy is assigned to 7210,7250 and 1830 SAP Access Ingress policies, 7210 and 1830 Network policies, and 7250 Ingress CoS policies.

The NFM-P provides a default 7210/7250 Dot1p classification policy. The default policy is applied if no user-configured policy is explicitly assigned.

The 7210 SAS-K supports 7210/7250 Dot1p classification policies assigned to access ingress and network policies on the device.

50.84.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7210/7250 Dot1p from the NFM-P main menu. The 7210/7250 Dot1p form opens.

2

Click Create or choose an existing policy and click Properties. The Dot1p Classification Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

The values configured in the Default Forwarding Class panel are applied to packets that do not match any of the associations on the Dot1p Classification tab; see Step 5.

4

Click on the Dot1p Classification tab.

5

Associate dot1p bits with FCs and profiles:

1. Click Create, or choose an existing entry and click Properties. The Ingress Dot1p, Dot1p Classification Policy (Create) form opens.
2. Configure the required parameters.
3. Save your changes.

6

Save your changes and close the forms.

END OF STEPS

50.85 To configure a 7210/7250 DSCP classification policy

50.85.1 Purpose

A 7210/7250 DSCP classification policy associates ingress DSCP values with forwarding classes and profiles. The policy is assigned to 7210, 7250 and 1830 SAP Access Ingress policies, 7210 and 1830 Network policies, and 7250 Ingress CoS policies. For table-based DSCP ingress classification, the policy is assigned to Ethernet ports and L3 interfaces configured for RVPLS; see 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

The NFM-P provides a default 7210/7250 DSCP classification policy. The default policy is applied if no user-configured policy is explicitly assigned.

The 7210 SAS-K supports 7210/7250 DSCP classification policies assigned to access ingress and network policies on the device.

The 7210 SAS-Mxp and 7210 SAS-R support 7210/7250 DSCP classification policies assigned to access ingress policies, access ports, and L3 access interfaces configured for RVPLS.

50.85.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7210/7250 DSCP from the NFM-P main menu. The 7210/7250 DSCP form opens.

2

Click Create or choose an existing policy and click Properties. The DSCP Classification Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

The values configured in the Default Forwarding Class panel are applied to packets that do not match any of the associations on the Ingress DSCP tab; see Step 5.

4

Click on the Ingress DSCP tab.

5

Associate DSCP bits with FCs and profiles.

1. Click Create, or choose an existing entry and click Properties. The Ingress DSCP, DSCP Classification (Create) form opens.
2. Configure the required parameters.
3. Save your changes.

6

Save your changes and close the forms.

END OF STEPS

50.86 To configure a 7210/7250 MPLS LSP-EXP classification policy

50.86.1 Purpose

A 7210/7250 MPLS LSP-EXP classification policy associates network ingress LSP-EXP values with forwarding classes and profiles. The policy is assigned to 7210 and 1830 network policies, and 7250 Ingress CoS policies.

The NFM-P provides a default 7210/7250 MPLS LSP-EXP classification policy. The default policy is applied if no user-configured policy is explicitly assigned.

The 7210 SAS-K12 and 7210 SAS-K30 ETR support MPLS LSP-EXP classification policies assigned to network policies on the device.

50.86.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7210/7250 MPLS LSP-EXP from the NFM-P main menu. The 7210/7250 MPLS LSP-EXP form opens.

2

Click Create or choose an existing policy and click Properties. The MPLS LSP-EXP Classification Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

The values configured in the Default Forwarding Class panel are applied to packets that do not match any of the associations on the Ingress LSP-EXP tab; see Step 5.

4

Click on the Ingress LSP-EXP tab.

5

Associate LSP-EXP bits with FCs and profiles.

1. Click Create, or choose an existing entry and click Properties. The Ingress LSP-EXP, MPLS LSP-EXP Classification Policy (Create) form opens.
2. Configure the required parameters.
3. Save your changes.

6

Save your changes and close the forms.

End of steps

50.87 To configure a 7210 FC Meter Map policy

50.87.1 Purpose

A 7210 FC Meter Map policy is assigned to a 7210, 7250 and 1830 SAP Access Ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544).

Service meters for SAP ingress provide an option to use meter resources from the ingress service-meter pool, which provides a larger number of meters/policers for use by access SAPs. This option is available only with table-based classification; it is not available when CAM-based classification is used. The Use SVC Meter Pool parameter must be set to true in the 7210, 7250 and 1830 SAP Access Ingress policy.

Note: Table-based classification uses meters from either the TCAM pool or the service meter pool, based on the SAP ingress policy type. If the SAP ingress policy is configured to use the use-svc-meter-pool parameter, the policy uses the service meter pool, otherwise the policy uses the TCAM meter pool.

7210 FC Meter Map policy attached to the 7210 SAP-ingress QoS policy supports only DSCP and Dot1p classifications. For the Policy ID parameter, the configurable values are 1 or 2. The default policy uses an ID values of 1; a user-defined policy uses an ID value of 2.

50.87.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7210 FC Meter Map from the NFM-P main menu. The 7210 FC Meter Map form opens.

2

Click Create or choose an existing policy and click Properties. The FC-meter-map (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Map FCs to meters.

1. Click on the Forwarding Classes tab.
2. Choose an FC in the list, or click Create. The Forwarding Class (Create|Edit) form opens.
3. Configure the Forwarding Class parameter and select the meter or meters to map to that forwarding class.

4. Click Apply to save the mapping and map another FC, or click OK to save your changes and close the form.
5. Save your changes and close the forms.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.88 To configure a 7210 Port Access Ingress Policy

50.88.1 Purpose

7210 Port Access Ingress Policy allows you to access ports as layer 2 uplinks. As a result, you can apply a single port-based access-ingress policy on ingress of an access port instead of using per SAP ingress policies. This allows a single policy definition to be used to classify and rate-limit all traffic received over access ports.

Meters are mapped to FCs according to the user-defined configuration in the policy. The maximum number of meters supported on a local policy varies depending on the chassis type.

50.88.2 Steps

1

Choose Policies→QoS→SROS QoS→Access Ingress→7210 Port Access Ingress Policy from the NFM-P main menu. The 7210 Port Access Ingress form opens.

2

Click Create or choose an existing policy and click Properties. The 7210 Port Access Ingress Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Configure ingress meters.

The Number of QoS Classifiers parameter, configured in Step 3 , affects the number of meters available.

Perform the following:

1. Click on the Meter tab.
2. Select a meter and click Properties, or click Create. The Meter form opens.
3. Configure the required parameters on the General tab.

You cannot choose ID values 1 or 9 for user-configured meters; these ID values are used by default meters.

4. Configure the CIR/PIR parameters:
   a. Click on the CIR/PIR tab.
   b. Configure the required parameters.
5. Configure the Burst Size parameters:
   a. Click on the Burst Size tab.
   b. Configure the required parameters.
      The parameters are configurable when the Default check box above each is deselected.
   c. Click OK.

5

Click on the Forwarding Classes tab.

6

Click Create or choose an existing forwarding class and click Properties. The Forwarding Class (Create | Edit) form opens.

7

Configure the required parameters and click OK.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.89 To configure a 7250 Ingress CoS policy

50.89.1 Before you begin

A 7250 Ingress CoS policy associates ingress CoS values with forwarding classes and profiles, using selected classification policies. The 7250 Ingress CoS policy provides a container for various combinations of ingress classification policies.

Before you configure a 7250 Ingress CoS policy, configure the following polices as required:

- 7210/7250 Dot1p classification policy; see 50.84 “To configure a 7210/7250 Dot1p classification policy” (p. 1637)
- 7210/7250 DSCP classification policy; see 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638)
- 7210/7250 MPLS LSP-EXP classification policy; see 50.86 “To configure a 7210/7250 MPLS LSP-EXP classification policy” (p. 1639)

The 7250 Ingress CoS policy is assigned to 7210,7250 and 1830 SAP Access Ingress policies; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544).

You cannot assign a 7250 Ingress CoS policy to a 7210,7250 and 1830 SAP Access Ingress policy if the 7250 Ingress CoS policy contains a 7210/7250 MPLS LSP-EXP classification policy.

The NFM-P provides two default 7250 Ingress CoS policies, with policy IDs 1 and 2. The default policies are not configurable. Default policy 1 is assigned to 7210,7250 and 1830 SAP Access Ingress policies when no user-configured policy is explicitly assigned. Default policy 2 is assigned to 7250 Network Ingress policies when no user-configured policy is explicitly assigned.

If you are configuring an ingress classification policy for a 7250 IXR with SR OS, see 50.45 “To configure a 7250 SROS Ingress Classification policy” (p. 1579). For more information about QoS policy support on the 7250 IXR, see 50.24 “7250 IXR QoS policies” (p. 1534).

50.89.2 Steps

1

Choose Policies→QoS→SROS QoS→Ingress Classification→7250 Ingress CoS from the NFM-P main menu. The 7250 Ingress CoS form opens.

2

Click Create, or choose an existing policy and click Properties. The Ingress CoS Classification Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

In the Classification policies panel, select the required classification policies:

- Dot1p Classification
- DSCP classification
- LSP-EXP Classification

See 50.89.1 “Before you begin” (p. 1642) for information about classification policies.

If you are configuring a 7250 Ingress CoS policy for a 7210,7250 and 1830 SAP Access Ingress policy, do not select an LSP-EXP Classification policy.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.90 To configure an OmniSwitch AOS QoS Network Group policy

50.90.1 Purpose

Perform this procedure to create or modify an AOS QoS Network Group policy. You can select network group policies to define a Source Network Group or Destination Network Group during configuration of IP match criteria in an AOS QoS Condition policy; see 50.92 “To configure an OmniSwitch QoS Condition policy ” (p. 1645).

50.90.2 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS Network Group, or choose QoS Network Groups (AOS QoS) from the object drop-down menu and select a Network Group, then click Properties. The AOS QoS Network Group - Global Policy (Create|Edit) form opens.

3

If you are creating an AOS QoS Network Group policy, configure the Displayed Name parameter.

4

Click on the Network Group Element tab and choose an item from the list, or click Create. The QoS Network Group Element (Create|Edit) form opens.

5

Configure the required parameters and click OK.

6

Create additional network group elements as required.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.91 To configure an OmniSwitch AOS QoS MAC Group policy

50.91.1 Purpose

Perform this procedure to create or modify an AOS QoS MAC Group policy. You can select MAC group policies to define a Source MAC Group or Destination MAC Group during configuration of MAC match criteria in an AOS QoS Condition policy; see 50.92 “To configure an OmniSwitch QoS Condition policy ” (p. 1645).

50.91.2 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS MAC Group, or choose QoS MAC Groups (AOS QoS) from the object drop-down menu and select a MAC Group, then click Properties. The AOS QoS MAC Group - Global Policy (Create|Edit) form opens.

3

If you are creating an AOS QoS MAC Group policy, configure the Displayed Name parameter.

4

Click on the MAC Group Element tab and choose an item from the list, or click Create. The QoS MAC Group Element (Create|Edit) form opens.

5

Configure the required parameters and click OK.

6

Create additional MAC group elements as required.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.92 To configure an OmniSwitch QoS Condition policy

Note: OmniSwitch QoS policies consist of a condition and an action. You must configure at least one condition and one action before you can create a policy.

50.92.1 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS Condition or choose an AOS QoS condition and click Properties. The AOS QoS Condition - Global Policy (Create | Edit) form opens.

3

Configure the Displayed Name parameter.

4

Click on the Filter Properties tab and configure the required parameters in the Port Match Criteria panel.

5

Click on the Layer 2 tab and configure the required parameters.

In the MAC Match Criteria panel, click Select to choose an AOS QoS MAC Group policy for the Source MAC Group and Destination MAC Group; see 50.91 “To configure an OmniSwitch AOS QoS MAC Group policy” (p. 1645).

6

Click on the Layer 3 tab and configure the required parameters.

In the IP Match Criteria panel, click Select to choose an AOS QoS Network Group policy for the Source Network Group and Destination Network Group; see50.90 “To configure an OmniSwitch AOS QoS Network Group policy” (p. 1644).

7

Click on the Layer 4 tab and configure the IP Protocol parameter, if required.

8

If you set the IP Protocol to a value other than HOPOPT (0) in the previous step, configure the required parameters in the IP Port Match Criteria panel.

9

Save your changes and close the forms.

END OF STEPS

50.93 To configure an OmniSwitch QoS policy action

Note: OmniSwitch QoS policies consist of a condition and an action. You must configure at least one condition and one action before you can create a policy.

50.93.1 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS Action or choose an AOS QoS action and click Properties. The AOS QoS Action - Global Policy (Create | Edit) form opens.

3

Configure the required parameters.

Note: You cannot use DSCP and ToS parameters in the same action.
You cannot use Maximum Bandwidth (Kbps) and Information Rate and Burst Size parameters in the same action.

4

Save your changes and close the form.

END OF STEPS

50.94 To create an OmniSwitch QoS policy

Note: OmniSwitch QoS policies consist of a condition and an action. You must configure at least one condition and one action before you can create a policy.

50.94.1 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS Policy or choose an AOS QoS policy and click Properties. The AOS QoS Policy - Global Policy (Create | Edit) form opens.

3

Configure the required parameters.

4

Select a default action and a default condition.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

50.95 To create an OmniSwitch QoS list

50.95.1 Purpose

A QoS list includes support of egress policy configurations to perform last minute operations on the data traffic before it egresses through the port.

See 50.92 “To configure an OmniSwitch QoS Condition policy ” (p. 1645) , 50.93 “To configure an OmniSwitch QoS policy action” (p. 1647) , and 50.94 “To create an OmniSwitch QoS policy” (p. 1647) to create AOS QoS policy actions, policy conditions, and policies.

50.95.2 Steps

1

Choose Policies→QoS→AOS QoS Policies from the NFM-P main menu. The AOS QoS Policies form opens.

2

Click Create→QoS List or choose an AOS QoS list and click Properties. The AOS QoS List - Global Policy (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click on the List Policies tab and click Create. The AOS QoS Policy List, AOS QoS List, Global Policy (Create) form opens.

5

Select an AOS QoS policy list.

6

Click Properties. The AOS QoS Policy - Global Policy (Edit) form opens.

7

Configure the Precedence parameter, if required.

8

Save your changes and close the forms.

END OF STEPS

50.96 To configure a Generic QoS Profile

50.96.1 Purpose

The Generic QoS Profile (GQP) provides an abstraction layer above the specific 7x50, 7210, and Multi Vendor QoS policies in the NFM-P. This functionality allows users to create QoS profiles that are generic for both proprietary NEs and third party NEs.

Note: Considerations for binding policies to GQPs
The following considerations apply when binding SAP Access Ingress/Egress policies or 7210 SAP Access Ingress/Egress policies to the GQP:

- When applying a SAP Access Ingress/Egress policy or 7210 SAP Access Ingress/Egress policy to the GQP, you can only apply those policies that have matching queues/meters/policers with the GQPs.
- Upon GQP creation and when the relevant GQP properties are left empty, relevant values from the assigned ingress policy and/or the egress policy are used to auto-populate the relevant GQP properties. This only occurs upon GQP creation and only if a policy is added. If any of the GQP properties are edited, this will not occur. Additional policies can be added after initial creation. The relevant properties are those set in the policy's queues, meters and policers.
- During the auto-populate, if a policy has the same queue ID entry as a policer/meter entry and the same CIR, PIR, CBS, or MBS characteristics, the ingress buffers are auto-populated, displaying the type as Not Applicable in the GQP. When the entries have different CIR, PIR, CBS, or MBS characteristics, the ingress buffers are auto-populated, displaying the type as Queue in the GQP.
- When comparing queue/meter/policers entries between the two policy types: only the first 8 queue entries are used for comparison; for meter and policer entries, only the first 32 entries are used.
- When you select an additional ingress or egress policy after the first ingress or egress policy has been selected, only policies with matching relevant properties are displayed in the select NFM-P window.
- When a GQP is cloned or modified using an API, the NFM-P creates a corresponding GQP and related ingress and egress QoS policies.

50.96.2 Steps

1

Choose Policies→QoS→Generic QoS Profiles from the NFM-P main menu. The Generic QoS Profiles form opens.

2

Click Create or choose a Generic QoS Profile and click Properties. The Generic QoS Profile (Create | Edit) form opens.

3

Configure the required parameters on the General tab.

If the MultiVendor Policies parameter is enabled, and the Generic QoS Profile is being created or modified, all appropriate multi-vendor global policies are auto-generated upon applying the change. The multi-vendor policies are automatically added to the GQP, but do not replace any non-auto-generated policies which are already there.

You can view the generated policies on the QoS Policies tab of the GQP and in the policy managers from the NFM-P main menu, Policies→Multi-Vendor, after the GQP changes are applied.

In a similar way to the auto-population of GQP values, the NFM-P extracts the values set in the GQP buffer (of type Policer only), Class of Service, and Match Criteria, and auto-generates Multi-Vendor QoS policies to match these properties. The new policies are created with a fixed naming convention, ending with a number.

Upon creation of these new policies, if the NFM-P finds an existing policy with the same name and number, it checks if any local definitions of that policy (or any component policies already bound to that policy) exist. If there are no local definitions, the NFM-P updates the existing policies and ensures they are assigned to the GQP. If a local definition does exist, the NFM-P creates a new set of policies with the same names as the existing policies, but with incremental numbers. The original policies in the GQP are then replaced with new policies. This is done to ensure that policy distribution to the NEs is as accurate as possible. The set of attributes which are supported by the auto-generation are described in the *NSP NFM-P Multi-Vendor Policy Guide*.

4

Configure an Egress QoS Buffer, Class of Service, Match Criteria, and Scheduling.

1. Click on the Egress tab, then the Buffer sub-tab.
2. Click Create or select an existing buffer and click Properties. The Egress QoS Buffer form opens.
3. Configure the required parameters.

   The Buffer ID is mandatory and can represent a queue, policer, or meter.

   Note: When a generic QoS profile contains a buffer with its Buffer Type set to N/A, then based on Buffer ID matching, both the queue and policer entries are updated in the policy. However, if the Buffer Type is configured as Policer, then only the corresponding policer that matches the Buffer ID is overridden in the policy. If the Buffer Type is configured as Queue, then only the corresponding queue that matches the Buffer ID is overridden in the policy.
4. Click OK to save the changes and close the form.
5. Click on the Class of Services sub-tab.
6. Select one of the eight default Class Names and click Properties. The Egress QoS Class of Service form opens.

7. Configure the Description parameter and select a Buffer ID to associate with this Class Name.
8. Click OK to save the changes and close the form.
9. Click on the Match Criteria sub-tab.
10. Click Create or select an existing entry and click Properties. The Egress QoS Match Criteria form opens.
11. Configure the parameters on the General tab.
12. Configure the parameters on the IPv4 Match Criteria, IPv6 Match Criteria, and MAC Match Criteria tabs, as required.
13. Click OK to save the changes and close the form.
14. Click on the Scheduling sub-tab.
15. Click Create or select an existing QoS Scheduler and click Properties. The Egress QoS Scheduler form opens.
16. Configure the Name parameter and select a scheduler Type, then configure the remaining parameters.
17. Click OK to save the changes and close the form.
18. Repeat substeps 2 to 17 to configure additional egress buffers.

5

Configure an Ingress QoS Buffer, Class of Service, Match Criteria, and Scheduling.

1. Click on the Ingress tab, then the Buffer sub-tab.
2. Click Create or select an existing buffer and click Properties. The Ingress QoS Buffer form opens.
3. Configure the required parameters.

   The Buffer ID is mandatory and can represent a queue, policer, or meter.

   Note: When a generic QoS profile contains a buffer with its Buffer Type set to N/A, then based on Buffer ID matching, both the queue and policer entries are updated in the policy. However, if the Buffer Type is configured as Policer, then only the corresponding policer that matches the Buffer ID would be overridden in the policy. If the Buffer Type is configured as Queue, then only the corresponding queue that matches the Buffer ID would be overridden in the policy.
4. Click OK to save the changes and close the form.
5. Click on the Class of Services sub-tab.
6. Select one of the eight default Class Names and click Properties. The Ingress QoS Class of Service form opens.
7. Configure the Description parameter and select a Buffer ID to associate with this Class Name.
8. Click OK to save the changes and close the form.
9. Click on the Match Criteria sub-tab.

10. Click Create or select an existing entry and click Properties. The Ingress QoS Match Criteria form opens.
11. Configure the parameters on the General tab.
12. Configure the parameters on the IPv4 Match Criteria, IPv6 Match Criteria, and MAC Match Criteria tabs, as required.
13. Click OK to save the changes and close the form.
14. Click on the Scheduling sub-tab.
15. Click Create or select an existing QoS Scheduler and click Properties. The Ingress QoS Scheduler form opens.
16. Configure the Name parameter and select a scheduler Type, then configure the remaining parameters.
17. Click OK to save the changes and close the form.
18. Repeat substeps 2 to 17 to configure additional ingress buffers.

6

Assign QoS policies to the profile.

Perform one of the following:

a. For Nokia equipment:
   1. Click on the QoS Policies tab.
   2. For an SROS NE, click the SROS Ingress/Egress sub-tab and then click the appropriate Select buttons to choose the required QoS policies. The policy types include:
      - SAP Access Ingress Policy
      - SAP Access Egress Policy
      - Ingress Scheduler Policy
      - Egress Scheduler Policy
      - Ingress Policer Control Policy
      - Egress Policer Control Policy
   3. For a 7210 NE, click the 7210 Ingress/Egress sub-tab and then click the appropriate Select buttons to choose the required QoS policies. The policy types include:
      - 7210 SAP Access Ingress Policy
      - 7210 SAP Access Egress Policy

b. For a 7250 IXR NE, click the 7250 SROS VLAN QoS sub-tab and then click the Select button to choose the required 7250 SROS VLAN QoS policy.

   Note: The 7250 SROS VLAN QoS policy is only applicable for 7250 IXR nodes.

c. For multi-vendor equipment:
   1. Click on the QoS Policies tab, then the Multi Vendor Ingress/Egress Policies sub-tab.
   2. Click the appropriate sub-tab to choose the required policies. The policy types include:
      - QoS Policies
      - Policers

- Filters
- Generic Policies

3. Click Add, then choose a policy from the Select form.
4. Click OK to close the form.

7

If you are viewing or configuring an existing Generic QoS Profile, you can click on the Dynamic Policies tab to view a list of dynamic policies associated with this profile. Such a policy is a local definition automatically created by the NFM-P when it receives a QoS policy override request from NSP.

Dynamic policies are essentially cloned SAP ingress or egress QoS policies that are used to override the CIR, PIR, CBS, and/or MBS values of the queues (and additionally, set policer overrides on 7x50 NEs). These policies are then deployed to specific NEs for service application. The Generic QoS Profile functionality is used to accomplish this. Dynamic policies are not listed in the standard QoS policy forms, since they are not created from the NFM-P or through CLI or the XML API.

You can only view dynamic policies and their implicit association to service SAPs, shown on the Access interfaces tab and its sub-tabs. A dynamic policy will automatically be removed from this list under the following conditions:

- the related service is deleted
- the related service SAP is deleted
- the QoS policy is cleared on the related service SAP
- the daughter card containing the service SAP port is removed

Dynamic policies are Local Edit Only policies that do not have any global reference. In addition, all local access ingress and egress policies contain a read-only attribute named Policy Mode that indicates if a policy is Dynamic or Static. In this context, Dynamic means the values have been modified by the Generic QoS Profile. This also includes policies for 7210 SAS variants where override support is not present by default.

8

Click OK to save the profile and close the form.

END OF STEPS

50.97 To configure QoS policy overrides on an L2 or L3 access interface

50.97.1 Purpose

Perform this procedure to override settings associated with an access ingress, access egress, or scheduler policy on an L2 or L3 access interface configured for a service. It is also applicable to multi-service sites.

See the following procedures for information about overriding access ingress, access egress, or scheduler policies that are associated with residential subscribers:

- 64.5 “To configure an SLA profile” (p. 1845) for the overrides associated with SLA profiles
- 64.4 “To configure a subscriber profile” (p. 1840) for the overrides associated with subscriber profiles

50.97.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Expand Access Interface (Service Management) in the object type drop-down menu and choose one of the following:

- L2 Access Interface (Service Management)
- L3 Access Interface (Service Management)

3

Choose an L2 or L3 access interface in the list and click Properties. The L2 | L3 Access Interface (Edit) form opens.

4

Click on the Override Policy Items tab and perform one of the following:

a. Create an override for an access ingress tagged IP/IPv6 match entry. Go to Step 5 .

b. Create an override for an access ingress or access egress policy queue. Go to Step 6 .

c. Create an override for an access egress policy HSMDA queue. Go to Step 7 .

d. Create an override for an ingress or egress policy policer. Go to Step 8 .

e. Create an override for an ingress or egress policer control policy. Go to Step 9 .

f. Create an override for an ingress or egress scheduler policy. Go to Step 10 .

g. Create an override for an access egress policy HS WRR group. Go to Step 11 .

5

To create an override for an access ingress tagged IP/IPv6 match entry, click on the Ingress Criteria tab and configure the required parameters.

6

To create an override for an access ingress or access egress policy queue, click on the Access Ingress or Access Egress Queues tab.

1. Click Create.The Access Ingress | Egress Queue Override form opens.
2. Select an ingress QoS policy queue.
3. Click on the Override tab.
4. Select a slope policy on the HS WRED Queue panel, if required.
5. Configure the required parameters.

   To override parameters, you must select the associated check box. De-select the associated Default or MAX check box to enter a specific value, or to select a drop-down option.

   Overrides for the PIR and CIR are configured in kbps if the Rate Type for the queue was originally set to kbps. The PIR and CIR overrides are configured as a percentage (%) if the Rate Type for the queue was originally set to Percent Port Limit or Percent Local Limit.

   The Port Average Overhead parameter is only configurable for access egress queues.

   For egress queues, you can configure queue depth monitoring in the Stats panel. For more information, see 50.21 “Queue depth monitoring” (p. 1527).
6. Save your changes and close the forms.

7

To create an override for an access egress policy HSMDA queue, click on the Access Egress HSMDA Queues tab.

1. Click Create. The Access Egress HSMDA Queue Override form opens.
2. Select the queue you want to override.
3. Click on the Override tab.
4. Configure the required parameters.

   If you select the Override check box for a parameter, the original parameter value is displayed.

   To specific another override value for the PIR (kbps) parameter, you must first deselect the associated MAX check box. To specific another override value for the Maximum Burst Size (bytes) parameter you must first deselect the associated Default check box.

   You can only specify an override value for WRR Weight for queues 1, 2, or 3.
5. Select the HSMDA Slope policy you want to override.
6. Save your changes and close the forms.

8

To create an override for an ingress or egress policy policer, click on the Ingress|Egress Policer tab.

1. Click Create. The Policer Override form opens.
2. Click Select and choose a policy policer, then click OK.

   Policy policers must be applied to the L2 Access Interface or the L3 Access Interface in order to be selected.
3. Configure the required parameters on the General tab.
4. Click on the Override tab.
5. Select the Override check box and configure any of the required parameters.

   To override any of the parameters, you must first select the associated check box. Deselect the associated Default or MAX check box to enter a specific value or to select a drop-down parameter option.

   Overrides for the PIR and CIR are configured in kbps if the Rate Type for the policer was originally set to kbps. The PIR and CIR overrides are configured as a percentage if the Rate Type for the policer was originally set to Percent Local Limit.
6. Save your changes and close the forms.

9

To create an override for an ingress or egress policer control policy, click on the Schedulers tab.

1. On the Ingress|Egress Policer Control Override panel, click Create. The Policer Control Override form opens.

   Policer control policies must be applied to the L2 or L3 access interface in order for the policer control override panels to be displayed.
2. Configure the required override check boxes and parameters.
3. Click OK to apply the changes and close the Policer Control Override form.
4. Click Apply in the L2 | L3 Access Interface form.
5. Select the policer control override you just created and click Properties. The Policer Control Override form opens.
6. Click on the Level Override Policy Items tab.
7. Select a level override from the list and click Properties. The Policer Level Override form opens.
8. Click on the Override tab.
9. Select the Override check box and configure the Maximum Cumulative Buffer Space parameter.
10. Save your changes and close the forms.

10

To create an override for an ingress or egress scheduler policy, click on the Ingress|Egress Schedulers tab. (Ingress or egress scheduler policies must be applied to the L2 or L3 access interface in order for the Ingress or Egress Scheduler tabs to appear under the Override Policy Items tab.)

1. Click Create. The Policy Scheduler Override form opens.
2. Select a scheduler policy.
3. Click on the Override tab.
4. Select the Override check box and configure any of the required parameters.

   To override any of the parameters, you must first select the associated Override check box to select a drop-down parameter option. For PIR (kbps) you must also deselect the MAX check box if you want to enter a specific value.
5. Save your changes and close the forms.

11

To create an override for an access egress policy HS WRR group, click on the HS WRR Groups tab.

1. Click Create. The SAP Egress QoS HS WRR Group override form opens.
2. Select an HS WRR group.
3. Click on the Overrides tab.
4. Configure the PIR Percent and Class Weight parameters, if required.

   To configure the parameters, you must first select the associated check boxes.
5. Save your changes and close the forms.

END OF STEPS

50.98 To configure QoS policy overrides on access ingress meters for the 7210 SAS

50.98.1 Purpose

Perform this procedure to override the meter settings of an access ingress policy for an L2 or L3 access interface on a 7210 SAS.

You cannot override the settings for QoS policies whose scope is set to exclusive.

Policy overrides on L2 access interfaces are not supported for the following:

- Uplink SAPs
- SAPs on mirror services
- SAPs with time of day suites attached

50.98.2 Steps

1

Open the properties form for the L2 or L3 access interface for which you need to configure an override:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose Site (Service Management) from the object drop-down menu and click Search.
3. Choose a 7210 SAS site from the list and click Properties. The properties form for the site opens.
4. Expand the site equipment tree as required, and click on an L2 | L3 access interface. The properties form for the access interface opens.

2

Click on the Override Policy Items tab, then on the Access Ingress Meters sub-tab.

3

Click Create. The Access Ingress Meter Override (Create) form opens.

4

Click Select in the Ingress QoS Policy Meter panel and choose a policy meter, then click OK.

5

Click on the Override tab.

6

Configure the required parameters.

To override any of the parameters, you must first select the associated check box. Deselect the associated Default or MAX check box to enter a specific value, or to select a drop-down menu option.

7

Click OK. The Access Ingress Meter Override (Create) form closes.

8

Click OK to save your changes and close the form.

END OF STEPS

50.99 To configure QoS policy overrides on access ingress queues for a 7210 SAS-X

50.99.1 Purpose

Perform this procedure to override the queue settings of an access ingress policy for an L2 or L3 access interface on a 7210 SAS-X.

You cannot override the settings for QoS policies whose scope is set to exclusive.

Policy overrides on L2 access interfaces are not supported for the following:

- Uplink SAPs
- SAPs on mirror services
- SAPs with time of day suites attached

50.99.2 Steps

1

Open the properties form for the L2 or L3 access interface for which you need to configure an override.

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose Site (Service Management) from the object drop-down menu and click Search.
3. Choose a 7210 SAS-X site from the list and click Properties. The properties form for the site opens.
4. Expand the site equipment tree as required, and click on an L2 | L3 access interface. The properties form for the access interface opens.

2

Click on the Override Policy Items tab, then on the Access Ingress Queues sub-tab.

3

Click Create. The Access Ingress Queue Override (Create) form opens.

4

Click Select in the Ingress QoS Policy Queue panel and choose a queue, then click OK.

Note: You can only configure an override on a queue that is mapped to a forwarding class.

5

Click on the Override tab and configure the required parameters.

To override any of the parameters, you must first select the associated check box. Deselect the associated Default or MAX check box to enter a specific value, or to select a drop-down menu option.

6

To override the queue management policy:

1. Select the check box for Displayed Name.
2. Click Clear, then click Select.
3. Choose a queue management policy from the list and click OK.

   Note:

   The selected queue management policy must be distributed to the NE that contains the SAP for which you are performing the policy override. Click on the Properties button to view the properties form for the queue management policy. You can use this properties form to distribute the policy to the NE, if required. See 49.2 “Policy distribution” (p. 1469) in Chapter 49, “Policies overview” .

7

Click on the OK button. The Access Ingress Queue Override (Create) form closes.

8

Click OK to save your changes and close the form.

END OF STEPS

50.100 To configure QoS policy overrides on port access egress queues for a 7210 SAS

50.100.1 Purpose

Perform this procedure to override the queue settings of a port access egress policy assigned to access ports on a 7210 SAS. When you enable overrides, you can change the settings for queues either to default settings or new values.

50.100.2 Steps

1

On the equipment tree, expand Network→NE→Shelf→Card Slot n→Daughter Card Slot n→Port n/n/n.

2

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

3

Click on the Ethernet tab, then click on the Port Access Egress tab.

4

Select a queue in the list, or click Create. The Port Access Egress Queue Override (Create|Edit) form opens.

5

If you are creating an override queue, configure the ID and Displayed Name parameters.

6

Click on the Override tab and configure the required parameters.

To override a parameter, you must first select the associated Override check box. If the Default check box for a parameter is enabled, the value for that parameter is set at the default value determined by the NFM-P. If the MAX check box for a parameter is enabled, the value for that parameter is set at the MAX value. If required, deselect the associated Default or MAX check box to configure a non-default value.

7

To override the queue management policy, click Select, choose a queue management policy from the list, and click OK.

Note: The required queue management policy must be distributed to the NE that contains the port for which you are performing the policy override. Click Properties to view the properties form for the queue management policy. You can use this properties form to distribute the policy to the NE, if required. See 49.2 “Policy distribution” (p. 1469) .

8

Save your changes and close the form.

9

Configure additional queue overrides, or save your changes and close the form.

END OF STEPS

50.101 To configure a shared policer policy

50.101.1 Purpose

A shared policer policy is a VLAN group policer for rate limit traffic from many VLANs. FCs from many VLANs are policed by the shared policer. The policy allows ingress policing to support a group of VLANs instead of a single VLAN.

50.101.2 Steps

1

Choose→Policies→QoS→SROS QoS→Shared Policer from the NFM-P main menu. The Shared Policer (Create|Edit) form opens.

2

Click Create or choose a policy and click Properties. The Shared Policer (Create|Edit) form opens.

3

Configure the required parameters.

4

Configure the CIR/PIR parameters:

1. Click on the CIR/PIR tab.
2. Configure the required parameters.

5

Configure the Burst Size parameters:

1. Click on the Burst Size tab.
2. Configure the required parameters.

   The parameters are configurable when the Default check box above each is deselected.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

51 Filter policies

51.1 Filter policies

51.1.1 Overview

You can use the NFM-P to configure a filter policy to define the network traffic filtering criteria used by routers and switches to allow or deny traffic on a device. You can specify a forwarding action for packets based on the matching criteria configured in the policy. For example, if you configure an ACL IPv6 filter policy, the network device analyzes the data that passes through the interface based on IPv6 matching criteria, compares it to the configuration set in the policy, and if the conditions are met, the traffic is forwarded as specified by the forwarding action. The process stops when the first complete match is found and the action defined in the entry is executed.

You can assign up to two filters to each applicable interface, circuit, or service the filter policy applies to; one in the receive direction and one in the transmit direction.

Note: When an interface, circuit, or service is not configured with a filter policy, all traffic is forwarded by the device. By default, no filters policies are applied by the NFM-P.

51.1.2 ACL IP and ACL IPv6 shared filter entries

RADIUS attributes can be used to enhance a locally configured IPv4 or IPv6 filter list with dynamic entries that can be shared between multiple subscriber hosts.

The ACL IP filter policy and ACL IPv6 filter policy include a High Watermark and Low Watermark configuration which specifies alarm thresholds for dynamically-inserted filter table counts.

51.1.3 Filter policy design considerations

Consider the following when you create NFM-P filter policies:

- The maximum number of filter entries varies depending on the NE and release. Confirm the allowable range of values for the intended device before you assign entry IDs.
- You can create up to 2000 MAC filter policies per managed NE.
- You can create up to 2000 IP filter policies per managed NE.
- The 7705 SAR has the following configuration limits for ACL filter entries:
  - For non-extended range filter entries, the match criteria cannot be defined to match more than 256 packets.
  - For extended range filter entries, there is no limit on the number of packet matches for the defined criteria; however, the total number of unique match criteria in use across all extended range filter entries within the same ACL IP filter cannot exceed eight.
- To expedite the creation of filter policies, you can copy the filter entries of one filter policy to another filter policy; see 51.19 “To copy filter policy filter entries” (p. 1700).

51.1.4 ACL MAC and ACL IP SAP and service tunnel forwarding

ACL MAC and ACL IP filters contain options for delivering packets to specific destination SAPs and service tunnels based on the match criteria. Because the forwarding action is specified separately from device configuration, the packet destination name must be entered directly in text form (for example, 1/1/2:500) and validated by the NFM-P.

Consider the following before you create a SAP or service tunnel forwarding filter:

- Although you can configure them on other service types, you should only create these filters on a VPLS. On other service types, packets are dropped.
- You can apply these filters on an ingress, but not an egress, SAP or service tunnel.
- The destination SAP or service tunnel must be on the same service as the device on which the filter is applied.
- The encapsulation type of the destination SAP and the associated port must be the same. Supported encapsulations are Ethernet Null, dot1q, QinQ, BCP, bridged Ethernet, FR, and ATM.

51.1.5 Embedded filter policy support

The NFM-P supports the use of embedded policies in ACL MAC, ACL IP, and ACL IPv6 filter policies that allow you to define a common set of filter policy rules that can then be embedded, or nested in one or more ACL filter policies.

When an exclusive or template filter policy that embeds one or more embedded filter policies is created, absolute values for all filter policy entries for both embedded and embedding filters are preserved. In the case of a conflict between entries (for instance, the same policy entry index), a higher priority filter entry is activated. The embedding filter has the highest priority, allowing its filter entries to always overwrite any embedded filter entries. This allows for customization of the common embedded filter policy rules within the embedding filter.

Any edits made to an embedded policy are automatically applied to all embedding filter policies that use the embedded filter policy.

The system ensures that system and hardware resources exist when a new embedded filter policy is created, changed or embedded by another filter policy, and either fails the configuration attempt or does not embed the embedded filter (depending on what resources are exhausted). An embedded filter is never embedded partially into another filter. If a change occurs to an already embedded filter fails due to a lack of hardware resources, the embedded filter will be removed from the embedding filter.

See 51.9 “To configure an embedding filter with embedded filter policies” (p. 1686) for information about configuring embedded filter policies.

51.1.6 ACL MAC, ACL IP, and ACL IPv6 web portal redirects

ACL MAC, ACL IP, and ACL IPv6 filter policies contain options for redirecting hosts to a URL address. The 7750 SR-7, 7750 SR-12, and 7950 XRS open a new connection to the specified web portal. The host can use the web portal to create or modify a service profile. The web portal updates the ACL policy to remove the redirection policy.

Note: The 7750 SR-1 does not support web portal redirect.

51.1.7 ACL MAC, ACL IP, and ALC IPv6 log filter configurations

You can optionally create the following policy types for ACL MAC, ACL IP, and IPv6 filter entries that can be used to specify where and how device related log information is collected and stored for a device. You can analyze the logs and extracted information on an ongoing basis to ensure your network is secure, or as required, to initiate the appropriate remediation action. See the appropriate SR node documentation for information about accessing log information on the SR device.

- Syslog policy that is used by the ACL Filter Log policy; the Syslog policy defines the destination details for log messages such as the target address and target UDP port, when the ACL Filter Log policy specifies a Syslog destination for storing log information. See 51.22 “To configure a Syslog policy” (p. 1704).
- ACL Filter Log policy that defines where log information for all actions performed on 7210 SAS, 7705 SAR, and 7x50 NEs which match ACL MAC, ACL IP, and ACL IPv6 filter entry criteria are written (memory or Syslog), how many log entries can be stored, and what action is performed when the log files meet the specified threshold. See 51.20 “To configure a ACL Filter Log policy” (p. 1702).

51.1.8 Redirect filters

A redirect filter policy allows operators to specifying multiple redirect target destinations and define health check test methods used to validate the ability for a given destination to receive redirected traffic. This destination monitoring allows a router to react to target destination failures. Operators can define IPv4 and IPv6 filter policies by specifying destinations to be IPv4 or IPv6 addresses respectively, and then referencing the redirect filter policy in the appropriate line card filter.

Reachability tests for specific destination addresses can be configured, including SNMP, URL, ping, and unicast route reachability. When a unicast route reachability test is configured, a destination becomes eligible for redirect policy destination selection only when the destination is reachable within the routing context that the policy is applied.

See 51.15 “To configure a Redirect Filter policy” (p. 1693) for information regarding configuring a redirect filter policy.

51.1.9 System filters

A system filter allows operators to configure a single set of filter policy rules that can then be activated in a system and used by other exclusive or template ingress IPv4 or IPv6 filter policies. When a system filter policy is activated (or when its configuration is changed), its entries are automatically downloaded to all line cards in the system. For packets to match system filter policy entries, an operator “chains” deployed IPv4/IPv6 ingress filter policies to the system filter. In this way, system filter entries are not duplicated into the chained filters. By employing such a chain, the active system filter policy rules are evaluated first. If a match occurs, the chained system filters are ignored. If no match occurs, only then are the rules of any chained filter policies evaluated. This allows a system filter policy to be used for implementing system blacklist rules, or security rules when multiple filter policies are required for other operational reasons (such as PBR).

Nokia recommends using a system filter policy with drop/forward actions. Other actions, for example, PBR actions or redirection to ISAs, should not be used unless the system filter policy is activated only in filters used by services that support such actions. Failure to observe this restriction

can lead to undesired behavior, since system filter actions are not verified against the services that the chaining filters are deployed for.

The system filter policy supports all IPv4/IPv6 filter policy match rules and actions, however, system policy entries cannot be LI or mirror sources. The system filter policy also does not support Radius, flowspec, or Gx inserted entries. Note that a system filter policy also requires chassis mode D to be set on an NE to which it is deployed.

See 51.18 “To configure a System Filter” (p. 1698) for information regarding configuring a system filter policy.

51.2 Supported filter policy types

51.2.1 Filter policies configurable in the NFM-P

Table 51-1 NFM-P filter policies

Filter policy type	Purpose	See
ACL Aggregate filter policy	This policy provides a container for configured filter policies. The following can be assigned to an ACL Aggregate filter policy: • ACL MAC filter policy • ACL IP filter policy • ACL IPv6 filter policy	51.3 “To configure an ACL Aggregate filter policy” (p. 1668)
ACL MAC filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface and service tunnel based on MAC matching criteria and the forwarding action.	51.4 “To configure an ACL MAC filter policy” (p. 1668)
ACL IP filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of a network or access IP interface or service tunnel based on IPv4 matching criteria and the forwarding action.	51.5 “To configure an ACL IP filter policy” (p. 1671)
ACL IPv6 filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface or service tunnel based on IPv6 matching criteria and the forwarding action.	51.6 “To configure an ACL IPv6 filter policy” (p. 1677)
ACL IP Exception filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface in an NGE or cellular domain based on protocol-specific matching criteria and source and destination information.	51.7 “To configure an ACL IP exception filter policy” (p. 1683)
ACL IPv6 Exception filter policy	This policy is an ACL IP exception filter for IPv6: it defines the network traffic filtering criteria used to allow or deny network traffic into or out of an interface based on protocol-specific matching criteria and source and destination information.	51.8 “To configure an ACL IPv6 exception filter policy” (p. 1684)
IP Prefix list filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface based on a defined list of IP Prefix list members.	51.11 “To configure an IP Prefix List policy” (p. 1688)

Table 51-1 NFM-P filter policies (continued)

Filter policy type	Purpose	See
Port list filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface based on a defined list of port list members.	51.12 “To configure a Port List policy” (p. 1690)
DHCP filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of IES and VPRN group interfaces based on DHCP matching criteria.	51.13 “To configure a DHCP Filter policy” (p. 1691)
DHCPv6 filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of IES and VPRN group interfaces based on DHCPv6 matching criteria.	51.14 “To configure a DHCPv6 filter policy” (p. 1692)
Redirect filter policy	This policy allows specifying multiple redirect target destinations and defining health check test methods used to validate the ability for a given destination to receive redirected traffic.	51.15 “To configure a Redirect Filter policy” (p. 1693)
Redirect policy binding	Redirect policy bindings allow configuration of an association between destination addresses in Rredirect policies, so that results of Ping tests can be shared.	51.16 “To configure a Redirect Policy Binding” (p. 1696)
ACL VLAN filter policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic on network ports of a ring card based on a defined VLAN range.	51.17 “To configure an ACL VLAN Filter policy” (p. 1697)
System filter policy	This policy allows operators to configure a filter chain for packet matching. In this chain, an active system filter policy rules are evaluated first. If no match occurs, then rules specified by any chained filter policies are evaluated.	51.18 “To configure a System Filter” (p. 1698)
Embedded and embedding filter policies	An embedded filter policy allows users to define a common set of filter policy rules that can then be nested in one or more other filter policies. The filters that support the embedded filter policies are referred to as embedding filter policies.	51.9 “To configure an embedding filter with embedded filter policies” (p. 1686)
Syslog policy	This policy defines the destination details for log messages such as the target address and target UDP port, when the ACL Filter Log policy specifies a Syslog destination for storing log information.	51.22 “To configure a Syslog policy” (p. 1704)
ACL Filter Log policy	This policy defines where log information for all actions performed on 7210 SAS, 7705 SAR, and 7x50 NEs that match ACL MAC, ACL IP, and ACL IPv6 filter entry criteria are written (memory or Syslog), how many log entries can be stored, and what action is performed when the log files meet the specified threshold.	51.20 “To configure a ACL Filter Log policy” (p. 1702)
GRE tunnel template policy	GRE tunnel template policies specify a set of GRE encapsulation parameters to use when the primary action for ACL IP and ACL IPv6 filter policies is set to Forward (GRE Tunnel). GRE Tunnel Templates are assigned to ACL IP filter policies and ACL IPv6 filter policies.	51.21 “To configure a GRE tunnel template” (p. 1703)
Protocol List policy	This policy defines the network traffic filtering criteria used to allow or deny network traffic into or out of an access interface based on a defined list of protocols.	51.10 “To configure a protocol list policy” (p. 1687)

51.3 To configure an ACL Aggregate filter policy

51.3.1 Important information

The ACL Aggregate Filter policy provides a container for configured filter policies. The following can be assigned to an ACL Aggregate filter policy:

- ACL MAC filter policy; see 51.4 “To configure an ACL MAC filter policy” (p. 1668)
- ACL IP filter policy; see 51.5 “To configure an ACL IP filter policy” (p. 1671)
- ACL IPv6 filter policy; see 51.6 “To configure an ACL IPv6 filter policy” (p. 1677)

51.3.2 Steps

1

Choose Policies→Filter→ACL Aggregate Filter from the NFM-P main menu. The ACL Aggregate Filter Policies form opens.

2

Click Create or choose an existing policy and click Properties. The Aggregate Filter (Create|Edit) form opens.

3

Configure the parameters as required.

In the Filter References panel, select an IP, IPv6, or MAC filter policy.

For information about IP, IPv6, and MAC filter policies, see Table 51-1, “NFM-P filter policies” (p. 1666).

Each ACL Aggregate Filter policy must contain only one filter reference. If the ACL Aggregate filter policy contains more than one filter reference, distribution to the NE will fail.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

51.4 To configure an ACL MAC filter policy

51.4.1 Steps

1

Choose Policies→Filter→ACL MAC Filter from the NFM-P main menu. The ACL MAC Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL MAC Filter (Create|Edit) form opens.

3

Configure the parameters as required.

The Default Action parameter specifies the action to be applied to packets when no action is specified in the MAC filter entries or when the packets do not match the specified criteria.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

4

If you need to configure the parameters on the Embedded Filters tab, refer to 51.9 “To configure an embedding filter with embedded filter policies” (p. 1686) .

5

Configure a filter entry.

1. Click on the Filter Entries tab and click Create. The Entry, ACL MAC Filter (Create) form opens.
2. Configure the required parameters.
3. Click Select to assign a Log ID to the ACL MAC filter entry.
4. Click Select in the Time Range panel to assign a time range for the ACL MAC filter entry or click Create to create a new time range. The Select Time Range - MacFilterEntry list form opens.
5. Select a time range entry and click OK. The ACL MAC Filter (Create) form refreshes with the time range information.

   Note: ACL filters that include ACL filter entries to which you have assigned a time range cannot be assigned to a time of day suite policy.

   Time ranges with which you have associated a ACL filter within a time of day suite policy cannot be assigned to ACL filter entries of that ACL filter.
6. Click on the Filter Properties tab.
7. Configure the Primary Action parameter. The Entry, ACL MAC Filter (Create) form refreshes to display the parameters, panels, and tabs applicable to the option you choose. As examples, if you choose Forward (SAP), the Forwarding Destination tab appears, along with the PBR Down Action Override and the Hold Time (seconds) parameters; if you choose HTTP Redirect, the Redirect URL and Allow RADIUS Override parameters appear; if you choose Forward (ESI), then the PBR Down Action Override parameter and the Forwarding ESI Details tab appear.

Note: The Forward (ESI) option provides the ability to steer traffic using an ESI value in an EVPN data center. The required traffic flow is identified using ACL IP, ACL IPv6, or ACL MAC filters, and then the action associated with the filter steers the traffic towards the service functions hosted on the EVPN data center.

Forward (ESI) is supported only if a device is in chassis mode D (for those NEs that have chassis mode support).

8. Configure the parameters associated with each Primary Action parameter option as required.
9. Configure the remaining parameters on the Filter Properties tab as required.

   Note: Configuring a Secondary Action to specify multiple PBR/PBF targets provides redundancy and load-sharing capacities on steered traffic. Choosing a Secondary Action will also display additional sub-tabs and parameters that must be configured. You must configure a Primary Action before being able to configure a Secondary Action.

   The Source MAC, Src Mask, Destination MAC, Dst Mask, Dot1p, Dot1p Mask, Low ISID and High ISID parameter pairs are configurable when the check box for each pair is selected.

   The Low ISID and High ISID parameters are configurable when the MAC Filter Type parameter is set to ISID.

   The Inner Tag Value, Inner Tag VID Mask, Outer Tag Value, and Outer Tag VID Mask parameters are configurable when the MAC Filter Type parameter is set to VID.

   The DSAP, DSAP Mask, SSAP, and SSAP Mask parameters are configurable when the Frame Type parameter value is set to e802dot2LLC and the MAC Filter Type parameter is set to Normal.

   The SNAP OUI and SNAP PID parameters are configurable when the Frame Type parameter value is e802dot2SNAP and the MAC Filter Type parameter is set to Normal.

   The Ether Type parameter is configurable only when the Frame Type parameter value is set to Ethernet II and the MAC Filter Type parameter is set to Normal.
10. Save your changes and close the form.

6

To add an additional filter entry, repeat Step 5.

7

To define the order in which the policy tries to match filter entries with packets, perform the following steps for each filter entry.

1. Click Refresh to find an existing filter entry. The list of filter entries is displayed.
2. Select a filter entry and click Renumber ID. The Renumber Entry ID form opens.
3. Configure the New Entry ID parameter.
4. Save your changes. The Entry ID column displays the new identifier assigned to the entry.

8

Save your changes. The ACL MAC Filter Policies form reappears.

9

CAUTION

Service Disruption

Distributing a global ACL MAC filter policy with no filter entries (either because none have been created or all existing ones have been deleted) can cause a service outage. You should ensure that the policy has at least one filter entry, or you must be certain that distributing an empty policy is what you really intend to do. A global policy will be distributed to all of the policy local definitions.

If you attempt the manual distribution of an empty policy, two warning confirmations will be issued. The first warning is issued when you change the policy's Configuration Mode on the General tab from Draft to Released. You can either choose to proceed by clicking Yes, or abort the Configuration Mode change by clicking No.

The second warning is issued if you changed the Configuration Mode to Released and then try to proceed with the actual distribution in the Distribute form. You can either choose to proceed by clicking Yes, or abort the distribution by clicking No.

If you attempt to release an ACL MAC filter policy that has been initialized from an NE, you will also receive a warning confirmation, since the global policy may be partially updated from the local policy. The Discovery State indicator on the General tab displays this Initialized condition, and the Origin indicator identifies the NE. You should manually synchronize with a specific local policy before changing the Configuration Mode from Draft to Released.

Click Search, select the policy in the list and click Distribute to manually distribute the policy locally to devices. See 49.6 "To release and distribute a policy" (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

10

Close the ACL MAC Filter Policies form.

END OF STEPS

51.5 To configure an ACL IP filter policy

51.5.1 Steps

1

Choose Policies→Filter→ACL IP Filter from the NFM-P main menu. The ACL IP Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL IP Filter (Create|Edit) form opens.

3

Configure the parameters as required.

Note: If you are creating an ACL IP filter that you will embed in another filter policy, then you must set the Scope parameter to embedded. This is referred to as an embedded filter. If you are creating an ACL IP filter that will contain an embedded filter or will enable flowspec as an embedded filter, then you must set the Scope parameter to either of the following: template, exclusive, or system. This is referred to as an embedding filter.
If you copy an embedded filter policy (Scope: embedded), the copied policy will also have a Scope of embedded. Filter entries from the embedded policy are also copied, and will have an Entry Type of Normal. See 51.19 “To copy filter policy filter entries” (p. 1700) for information on copying filter entries.

Note: If you are configuring an active system filter, then you must set the Scope parameter to system. The Chain to System Filter parameter must not be enabled. See 51.18 “To configure a System Filter” (p. 1698) for system filter information.
If you are configuring a chained system filter, then you must set the Scope parameter to either template or exclusive. The Chain to System Filter parameter must also be enabled. See 51.18 “To configure a System Filter” (p. 1698) for system filter information.
To change an existing filter policy’s Scope parameter to or from the system option, the policy must have no Filter Entries configured.

Note: To perform traffic management, you must set the Scope parameter to cpm. After the ACL IP and ACL IPv6 filter policies deploy to the 7250 IXR NEs with the Scope parameter set to cpm, then the IP Administrative Status and IPv6 Administrative Status parameters can be set to Up in the CPM filter policy; see the *NSP System Administrator Guide* for more information about configuring a CPM filter policy.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

Filtering for packet length can be configured as an action condition, or, on supporting NEs, as a match criterion. If you are configuring a policy with filter entries that use Match Criteria for packet length filtering, you must set the IP Filter Type parameter to the Packet-Length option.

4

If you need to configure the parameters on the Embedded Filters tab, including the enabling of flowspec, refer to 51.9 “To configure an embedding filter with embedded filter policies” (p. 1686).

5

Click on the Insertion Blocks tab.

6

If required, enable the check box in the Host Shared Filter Configuration panel and configure the High Watermark or Low Watermark parameters.

7

Configure the parameters as required in the Group Entry Insertion Configuration panel.

8

Configure a filter entry.

1. Click on the Filter Entries tab and click Create.The Entry, ACL IP Filter (Create) form opens.
2. Configure the required parameters.
3. Click Select to assign a Log ID to the ACL IP filter entry.
4. Click Select in the Time Range panel to assign a time range for the ACL IP filter entry or click Create to create a new time range. The Select Time Range - IP FilterEntry list form opens. Otherwise, go to Step 8 , substep 6.
5. Select a time range entry and click OK. The Entry, ACL IP Filter (Create) form refreshes with the time range information.

   Note: ACL filters that include ACL filter entries to which you have assigned a time range cannot be assigned to a time of day suite policy.

   Time ranges with which you have associated a ACL filter within a time of day suite policy cannot be assigned to ACL filter entries of that ACL filter.
6. Click on the Filter Properties tab.
7. Configure the Primary Action panel.

   Configure the Action parameter. The form refreshes to display the parameters, panels, and sub-tabs applicable to the option you choose.

   Configure the parameters (and sub-tab parameters, if applicable) associated with the chosen Action parameter option, as required.

   The Forward (ESI) option provides the ability to steer traffic using an ESI value in an EVPN data center. The required traffic flow is identified using ACL IP, ACL IPv6, or ACL MAC filters, and then the action associated with the filter steers the traffic towards the service functions hosted on the EVPN data center. Forward ESI is supported only if a device is in chassis mode D (as applicable).

   The Rate Limit option provides the ability to protect a network against DDoS attacks by specifying a TTL value (or hop-limit for IPv6), or packet length. When the specified value is exceeded, the transit traffic is dropped.

   The Forward (GRE Tunnel) option allows you to assign a GRE tunnel template that defines the encapsulation parameters; see 51.21 “To configure a GRE tunnel template” (p. 1703).

   The Forward (SAP) option requires VPLS L2 Access Interfaces. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) for information on associating an ACL IP filter to a VPLS SAP.

The Forward Next Hop (Router) option allows you to associate the filter to a VPRN L3 Access Interface. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) for more information.

The Forward (Pattern) option provides the ability to forward packets that contain a predefined UDP signature that conforms to the configured pattern parameters, essentially “whitelisting” authorized packets.

When the Action parameter is set to Drop, Rate Limit, or Forward (Pattern), parameters are available for pattern matching on supporting devices. Pattern matching can help identify DDoS attacks.

If you set the IP Filter Type parameter to the Packet-Length option in Step 3, do not configure packet length options and parameters in the Primary Action panel. Instead, configure packet length options in the Match Criteria panel in Step 8, substep 11.

8. For certain Primary Action options, you can set an Extended Action parameter (and associated parameters, as applicable) to enable a supplementary action to the Primary Action. Configure as required.
9. Configure a Secondary Action to specify two PBR targets as part of a single filter policy entry. This provides redundancy and load-sharing capacities on steered traffic. When primary and secondary actions are both configured, PBR uses the primary action if its target is operationally up, or the secondary action if its target is operationally up.

   If both PBR targets are down, the default behavior you configure for the primary action is used, unless you configure the PBR Down Action Override parameter otherwise. In addition, you can set a sticky destination Hold Time for a given redundant filter entry.

   Choosing a Secondary Action will also display additional sub-tabs and parameters that you must configure as required. Note that a Primary Action must be configured prior to setting a Secondary Action.

   When you configure the Secondary Action to Forward (VPRN Target), the Secondary VPRN Target tab opens.
10. For certain Secondary Action options, you can set an Extended Action parameter (and associated parameters, as applicable) to enable a supplementary action to the Secondary Action. Configure as required.
11. Configure the remaining parameters on the Filter Properties tab as required. Note the following:
    - You can assign a configured protocol list policy; see 51.10 “To configure a protocol list policy” (p. 1687) . When you assign a protocol list policy, you must set the Protocol parameter to NONE.
    - The Source Port, Destination Port, and Port related parameters are configurable when the Protocol parameter value is TCP or UDP.
    - If you select the Source and Destination option for the Configuration Type parameter, you can configure the Source and Destination ports separately, specifying either a Mask, Range, or Port List for each. If you select the Port option for the Configuration Type, then the Mask, Range, or Port List you specify will apply to both the source and destination.
    - Configuring the Src Mask and Src Net Mask parameters is mutually exclusive.
    - Configuring the Dst Mask and Dst Net Mask parameters is mutually exclusive.

- When the Protocol parameter is set to TCP, the TCP Properties panel is available for enabling TCP flags on supporting NEs. When you distribute the policy, ensure that the NE supports the required TCP flags.
- The ICMP Code and ICMP Type parameters are configurable when the Protocol parameter value is IPv6_ICMP.
- The Egress PBR parameter can only be configured when the Action parameter is set to one of the following: Forward (Redirect Filter), TCP MSS Adjust, Ignore Match, Forward (ESI), Forward Next Hop, or Forward Next Hop (Router).
- The Bonding Connection ID parameter must be configured when the Primary Action parameter is set to Forward (Bonding Connection).
- Match Criteria Packet Length options are supported only when the IP Filter Type is set to Packet-Length in Step 3. The Packet Length Option parameter is mutually exclusive with the DSCP, IP Option, IP Opt Mask, Option Present, Multiple Option, and Source Route Option match criteria.
- The Match Criteria Time-To-Live option is supported only when the IP Filter type is set to Packet-Length in Step 3. Before you change the IP Filter type, you must delete any filter entry that has the Time-To-Live option selected.
- Select a GRE tunnel template as required.

12. Click on the Cflowd tab and configure the parameters as required.

    For the Sample Profile ID parameter, enter the ID number of an existing Cflowd sample profile. For information about Cflowd sample profiles see 12.10 “To enable and configure global Cflowd sampling on an NE” (p. 347).
13. Click on the Forwarding VRPN Target tab.
14. Configure the required parameters.
15. Click on the Secondary VPRN Target tab.
16. Configure the required parameters and select a router and LSP.
17. Save your changes and close the form.

9

To create an additional filter entry, repeat Step 8.

10

To define the order in which the policy tries to match filter entries with packets, perform the following steps for each filter entry.

1. Click Refresh to find an existing filter entry. The list of filter entries is displayed.
2. Select a filter entry and click Renumber ID. The Renumber Entry ID form opens.
3. Configure the New Entry ID parameter.
4. Save your changes. The Entry ID column displays the new identifier assigned to the entry.

11

Save your changes. The ACL IP Filter Policies form reappears.

12

CAUTION

Service Disruption

Distributing a global ACL IP filter policy with no filter entries (either because none have been created or all existing ones have been deleted) can cause a service outage. You should ensure that the policy has at least one filter entry, or you must be certain that distributing an empty policy is what you really intend to do. A global policy will be distributed to all of the policy's local definitions.

If you attempt the manual distribution of an empty policy, two warning confirmations will be issued. The first warning is issued when you change the policy's Configuration Mode on the General tab from Draft to Released. You can either choose to proceed by clicking Yes, or abort the Configuration Mode change by clicking No.

The second warning is issued if you changed the policy's Configuration Mode to Released and then try to proceed with the actual distribution in the Distribute form. You can either choose to proceed by clicking Yes, or abort the distribution by clicking No.

If you attempt to release an ACL IP filter policy that has been initialized from an NE, you will also receive a warning confirmation, since the global policy may be partially updated from the local policy. The Discovery State indicator on the General tab displays this Initialized condition, and the Origin indicator identifies the NE. You should manually synchronize with a specific local policy before changing the Configuration Mode from Draft to Released.

Click Search, select the policy in the list and click Distribute to manually distribute the policy locally to devices. See 49.6 "To release and distribute a policy" (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

13

Select the distributed policy and click Properties. The ACL IP Filter Global Policy (Edit) form opens.

14

Click on the Local Definitions tab and select a local definition from the list and click Properties. The ACL IP Filter Local Policy (Edit) form opens.

15

Click on the Insertion Blocks tab.

16

If required, enable the check box in the Host Shared Filter Configuration panel and configure the High Watermark or Low Watermark parameters.

17

Configure the parameters as required.

Note: The Group Entries Inserted panel displays the number of entries inserted on this filter range.

18

Click Sort Group Insertions and click OK.

19

To view filter entry data.

1. Click on the Filter Entries tab, then the Credit Control Entries tab or the RADIUS Entries tab.
2. Click Search and select an entry from the list and click Properties.
3. Close the form.

20

Close all open forms.

END OF STEPS

51.6 To configure an ACL IPv6 filter policy

51.6.1 Steps

1

Choose Policies→Filter→ACL IPv6 Filter from the NFM-P main menu. The ACL IPv6 Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL IPv6 Filter (Create|Edit) form opens.

3

Configure the parameters as required.

Note: If you are creating an ACL IPv6 filter that you will embed in another filter policy, then you must set the Scope parameter to embedded. This is referred to as an embedded filter. If you are creating an ACL IPv6 filter that will contain an embedded filter or will enable flowspec as an embedded filter, then you must set the Scope parameter to one of the following: template, exclusive, or system. This is referred to as an embedding filter.
If you copy an embedded filter policy (Scope: embedded), the copied policy will also have a Scope of embedded. Filter entries from the embedded policy are also copied, and will have an Entry Type of Normal. See 51.19 “To copy filter policy filter entries” (p. 1700) for information on copying filter entries.

Note: To perform traffic management, you must set the Scope parameter to cpm. After the

ACL IP and ACL IPv6 filter policies deploy to the 7250 IXR NEs with the Scope parameter set to cpm, then the IP Administrative Status and IPv6 Administrative Status parameters can be set to Up in the CPM filter policy; see section the *NSP System Administrator Guide* for more information about configuring a CPM filter policy.

Note: You can only configure the IPv6 Address Format parameter when the system resource profile settings are appropriate. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .
To configure an active system filter, set the Scope parameter to system. The Chain to System Filter parameter must not be enabled. See 51.18 “To configure a System Filter” (p. 1698) for system filter information.
To configure a chained system filter, set the Scope parameter to either template or exclusive. The Chain to System Filter parameter must also be enabled. See 51.18 “To configure a System Filter” (p. 1698) for system filter information.
To change an existing filter policy’s Scope parameter to or from the system option, the policy must have no Filter Entries configured.

Note: NEs that support next-generation CLI use the policy name as the key identifier for internal system reference. For policies on these NEs, you must configure a policy name (typically the service name or a numerical string). Policy IDs are also supported. You must configure a numerical range on the NE for auto-assigned policy IDs; see 12.65 “To configure an Auto-ID range for policies” (p. 395).

Filtering for packet length can be configured as an action condition, or, on supporting NEs, as a match criterion. If you are configuring a policy with filter entries that use Match Criteria for packet length filtering, you must set the IPv6 Filter Type parameter to the Packet-Length option.

4

If you need to configure the parameters on the Embedded Filters tab, including the enabling of flowspec, refer to 51.9 “To configure an embedding filter with embedded filter policies” (p. 1686) for detailed information.

5

Click on the Insertion Blocks tab.

6

If required, enable the check box in the Host Shared Filter Configuration panel and configure the High Watermark or Low Watermark parameters.

7

Configure the parameters in the Group Entry Insertion Configuration panel as required.

8

Configure a filter entry.

1. Click on the Filter Entries tab and click Create.The Entry, ACL IPv6 Filter, Global Policy (Create) form opens.

2. Configure the required parameters.
3. Click Select to assign a Log ID to the ACL IPv6 filter entry.
4. Click Select in the Time Range panel to assign a time range to the ACL IPv6 filter entry. The Select Time Range - Ipv6FilterEntry list form opens. Otherwise, go to Step 8 , substep 6.
5. Select a time range entry and click OK. The Entry, ACL IPv6 Filter (Create) form refreshes with the time range information.

   Note: ACL filters that include ACL filter entries to which you have assigned a time range cannot be assigned to a time of day suite policy.

   Time ranges with which you have associated a ACL filter within a time of day suite policy cannot be assigned to ACL filter entries of that ACL filter.
6. Click on the Filter Properties tab.
7. Configure the Primary Action parameter.

   The Entry, ACL IPv6 Filter (Create) form refreshes to display the parameters, panels, and sub-tabs applicable to the option you choose. For example, if you choose Forward (ESI), then the PBR Down Action Override parameter and the Forwarding ESI Details sub-tab appear.
8. Configure the parameters (and sub-tab parameters, if applicable) associated with the chosen Primary Action parameter option, as required.

   The Forward (ESI) option provides the ability to steer traffic using an ESI value in an EVPN data center. The required traffic flow is identified using ACL IP, ACL IPv6, or ACL MAC filters, and then the action associated with the filter steers the traffic towards the service functions hosted on the EVPN data center. Forward ESI is supported only if a device is in chassis mode D (if applicable to the device).

   The Rate Limit option provides the ability to protect a network against DDoS attacks by specifying a TTL value (or hop-limit for IPv6), or packet length. When the specified value is exceeded, the transit traffic is dropped.

   The Forward (GRE Tunnel) option allows you to assign a GRE tunnel template that defines the encapsulation parameters; see 51.21 “To configure a GRE tunnel template” (p. 1703).

   The Forward (SAP) option is only applicable to VPLS L2 Access Interfaces. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) for information on associating an ACL IPv6 filter to a VPLS SAP.

   The Forward Next Hop (Router) option allows you to associate the filter to a VPRN L3 Access Interface. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) for more information.

   The Forward (Pattern) option provide the ability to forward packets that contain a predefined UDP signature that conforms to the configured pattern parameters, essentially “whitelisting” authorized packets.

   When the Action parameter is set to Drop, Rate Limit, or Forward (Pattern), parameters are available for pattern matching on supporting devices. Pattern matching can help identify DDoS attacks.

   If you set the IPv6 Filter Type parameter to the Packet-Length option in Step 3, do not configure packet length options and parameters in the Primary Action panel. Instead, configure packet length options in the Match Criteria panel in Step 8, substep 12.

9. For certain Primary Action options, you can set an Extended Action parameter (and associated parameters, as applicable) to enable a supplementary action to the Primary Action. Configure as required.
10. Configure a Secondary Action to specify two PBR targets as part of a single filter policy entry. This provides redundancy and load-sharing capacities on steered traffic. When primary and secondary actions are both configured, PBR uses the primary action if its target is operationally up, or the secondary action if its target is operationally up.

    If both PBR targets are down, the default behavior you configure for the primary action is used, unless you configure the PBR Down Action Override parameter otherwise. In addition, you can set a sticky destination Hold Time for a given redundant filter entry.

    Choosing a Secondary Action will also display additional sub-tabs and parameters that you must configure as required. Note that a Primary Action must be configured prior to setting a Secondary Action.

    When you configure the Secondary Action to Forward (VPRN Target), the Secondary VPRN Target tab opens.
11. For certain Secondary Action options, you can set an Extended Action parameter (and associated parameters, as applicable) to enable a supplementary action to the Secondary Action. Configure as required.
12. Configure the remaining parameters on the Filter Properties tab as required. Note the following:
    - You can assign a configured protocol list policy; see 51.10 “To configure a protocol list policy” (p. 1687) . When you assign a protocol list policy, you must set the Protocol parameter to NONE.
    - The Source Port, Destination Port, and Port related parameters are configurable when the Protocol parameter value is TCP or UDP.
    - If you select the Source and Destination option for the Configuration Type parameter, you can configure the Source and Destination ports separately, specifying either a Mask, Range, or Port List for each. If you select the Port option for the Configuration Type, then the Mask, Range, or Port List you specify will apply to both the source and destination.
    - Configuring the Src Mask and Src Net Mask parameters is mutually exclusive.
    - Configuring the Dst Mask and Dst Net Mask parameters is mutually exclusive.
    - When the Protocol parameter is set to TCP, the TCP Properties panel is available for enabling TCP flags on supporting NEs. When you distribute the policy, ensure that the NE supports the required TCP flags.
    - The ICMP Code and ICMP Type parameters are configurable when the Protocol parameter value is IPv6_ICMP.
    - The Egress PBR parameter can only be configured when the Action parameter is set to one of the following: Forward (Redirect Filter), TCP MSS Adjust, Ignore Match, Forward (ESI), Forward Next Hop, or Forward Next Hop (Router).
    - The Bonding Connection ID parameter must be configured when the Primary Action parameter is set to Forward (Bonding Connection).
    - Match Criteria Packet Length options are supported only when the IPv6 Filter Type is set to Packet-Length in Step 3. Packet length configuration is mutually exclusive with Flow Label configuration.

- The Match Criteria Hop Limit option is supported only when the IPv6 Filter type is set to Packet-Length in Step 3. Before you change the IPv6 Filter type, you must delete any filter entry that has the Hop Limit option selected.

13. Configure the parameters on any additional tabs that appear as a result of the Primary Action parameter setting. The type of additional configuration depends on the primary action selected.
14. Click on the Cflowd tab and configure the parameters as required.

    For the Sample Profile ID parameter, enter the ID number of an existing Cflowd sample profile. For information about Cflowd sample profiles see 12.10 “To enable and configure global Cflowd sampling on an NE” (p. 347).
15. Click on the Next Hop Routing tab and configure the required parameters.
16. Save your changes and close the form.

9

To create an additional filter entry, repeat Step 8.

10

Click on the Forwarding VRPN Target tab.

11

Configure the required parameters.

12

Click on the Secondary VPRN Target tab.

13

Configure the required parameters and select a router and LSP.

14

To define the order in which the policy tries to match filter entries with packets, perform the following steps for each filter entry.

1. Click Refresh to find an existing filter entry. The list of filter entries is displayed.
2. Select a filter entry and click Renumber ID. The Renumber Entry ID form opens.
3. Configure the New Entry ID parameter.
4. Save your changes. The Entry ID column displays the new identifier assigned to the entry.

15

Save your changes, The ACL IPv6 Filter Policies form reappears.

16

CAUTION

Service Disruption

Distributing a global ACL IPv6 filter policy with no filter entries (either because none have been created or all existing ones have been deleted) can cause a service outage. You should ensure that the policy has at least one filter entry, or you must be certain that distributing an empty policy is what you really intend to do. A global policy will be distributed to all of the policy's local definitions.

If you attempt the manual distribution of an empty policy, two warning confirmations will be issued. The first warning is issued when you change the policy's Configuration Mode on the General tab from Draft to Released. You can either choose to proceed by clicking Yes, or abort the Configuration Mode change by clicking No.

The second warning is issued if you changed the policy's Configuration Mode to Released and then try to proceed with the actual distribution in the Distribute form. You can either choose to proceed by clicking Yes, or abort the distribution by clicking No.

If you attempt to release an ACL IPv6 filter policy that has been initialized from an NE, you will also receive a warning confirmation, since the global policy may be partially updated from the local policy. The Discovery State indicator on the General tab displays this Initialized condition, and the Origin indicator identifies the NE. You should manually synchronize with a specific local policy before changing the Configuration Mode from Draft to Released.

Click Search, select the policy in the list and click Distribute to manually distribute the policy locally to devices. See 49.6 "To release and distribute a policy" (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

17

Select the distributed policy in the list and click Properties The ACL IPv6 Filter Global Policy (Edit) form opens.

18

Click on the Local Definitions tab and select a local definition and click Properties. The ACL IPv6 Filter Local Policy (Edit) form opens.

19

Click on the Insertion Blocks tab.

20

If required, enable the check box in the Host Shared Filter Configuration panel and configure the High Watermark or Low Watermark parameters

21

Configure the required parameters:

Note: The Group Entries Inserted panel displays the number of entries inserted on this filter range.

22

Click Sort Group Insertions and save your changes.

23

To view filter entry data:

1. Click on the Filter Entries tab, then the Credit Control Entries tab or the RADIUS Entries tab.
2. Click Search and choose an entry and click Properties.

   Note: Some parameters may not be supported on the local definition of the policy, depending on the device. See 49.2 “Policy distribution” (p. 1469) and 49.1.5 “Policy audits” (p. 1468) in Chapter 49, “Policies overview”.
3. Close the form.

24

Close all open forms.

END OF STEPS

51.7 To configure an ACL IP exception filter policy

51.7.1 Steps

1

Choose Policies→Filter→ACL IP Exception Filter from the NFM-P main menu. The ACL IP Exception Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL IP Exception Filter (Create|Edit) form opens.

3

Configure the parameters on the General tab as required.

4

Configure one or more filter entries:

1. Click on the Filter Entries tab and click Create. The Entry, ACL IP Exception Filter, Global Policy (Create) form opens.
2. Configure the required parameters on the General tab.

3. Click on the Filter Properties tab.
4. Configure the Protocol parameter.

   The form refreshes to display the parameters applicable to the option you choose. For example, if you choose ICMP, the ICMP Properties panel appears.
5. In the Match Criteria panel, configure the Source IP and Destination IP parameters.
6. Configure additional parameters as required.
7. Save your changes and close the form.

You can configure up to 256 filter entries.

5

To delete an entry, choose the entry from the list in the Filter Entries tab and click Delete.

6

To define the order in which the policy tries to match filter entries with packets, perform the following steps for each filter entry.

1. Click Search to find an existing filter entry. The list of filter entries is displayed.
2. Select a filter entry and click Renumber ID. The Renumber Entry ID form opens.
3. Configure the New Entry ID parameter.
4. Save your changes. The Entry ID column displays the new identifier assigned to the entry.

7

Save your changes. The ACL IP Exception Filter Policies form reappears.

8

Close all open forms.

END OF STEPS

51.8 To configure an ACL IPv6 exception filter policy

51.8.1 Steps

1

Choose Policies→Filter→ACL IPv6 Exception Filter from the NFM-P main menu. The ACL IPv6 Exception Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL IPv6 Exception Filter (Create|Edit) form opens.

3

Configure the parameters on the General tab as required.

4

Configure one or more filter entries:

1. Click on the Filter Entries tab and click Create. The Entry, ACL IPv6 Exception Filter, Global Policy (Create) form opens.
2. Configure the required parameters on the General tab.
3. Click on the Filter Properties tab.
4. Configure the Next Header parameter.

   The form refreshes as needed to display the parameters applicable to the option you choose.
5. In the Match Criteria panel, configure the Source IP and Destination IP parameters.
6. Configure additional parameters as required.
7. Save your changes and close the form.

You can configure up to 256 filter entries.

5

To delete an entry, choose the entry from the list in the Filter Entries tab and click Delete.

6

To define the order in which the policy tries to match filter entries with packets, perform the following steps for each filter entry.

1. Click Search to find an existing filter entry. The list of filter entries is displayed.
2. Select a filter entry and click Renumber ID. The Renumber Entry ID form opens.
3. Configure the New Entry ID parameter.
4. Save your changes. The Entry ID column displays the new identifier assigned to the entry.

7

Save your changes. The ACL IPv6 Exception Filter Policies form reappears.

8

Close all open forms.

END OF STEPS

51.9 To configure an embedding filter with embedded filter policies

51.9.1 Purpose

Perform this procedure to create an embedding filter that contains one or more embedded filters. An embedded filter is used to define a common set of rules. The filter that nests embedded filters is referred to as the embedding filter.

51.9.2 Steps

Create the embedding filter policy

1

Perform one of the following:

a. Create an ACL IP filter. See 51.5 “To configure an ACL IP filter policy” (p. 1671).

b. Create an ACL IPv6 filter. See 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

Note: You must set the Scope parameter of the embedding filter to either template or exclusive.

Associate a flowspec configuration with the embedding filter policy, if required

2

Click on the Embedded Filters tab and then the Flowspec Embedded sub-tab.

3

Enable the Enable Flowspec parameter.

4

Choose the Router Instance Type, either Base or VPRN.

If you selected VPRN, then click Select to choose a VPRN service and click OK.

5

Configure the Offset and Admin State parameters.

6

To limit the BGP flowspec embedded filters to specific interfaces, configure the Group ID parameter with the same value as the group-ID configured for the flowspec-set or flowspec-set-trans community members configured in the required routing community policy; see 54.8 “To configure a community policy” (p. 1752).

7

You should also define the maximum size for an embedded flowspec on the routing instances that this policy will be applied to. See configuring the IPv4 Max Size and IPv6 Max Size parameters under the Routing tab in 27.2 “To configure a routing instance or a VRF instance” (p. 826) for base routers, or 79.26 “To configure a routing instance on a VPRN site” (p. 2564)for VPRN services, as required.

Associate an embedded IP filter policy with the embedding filter policy, if required

8

Click on the Embedded Filters tab and then the IP Embedded sub-tab.

9

Click Create. The Embedded Entry form opens.

10

Select an embedded filter and click OK.

11

Configure the Offset and Admin State parameters.

12

Click OK. The Embedded Entry form closes.

13

Click OK to save your changes. The ACL MAC/IP/IPv6 Filter properties form closes.

END OF STEPS

51.10 To configure a protocol list policy

51.10.1 Purpose

A protocol list policy defines a list of network protocols to be used as filtering criteria for filter entries in ACL IP and IPv6 filter policies.

Protocol list policies are assigned during IP or IPv6 ACL filter entry configuration; see 51.5 “To configure an ACL IP filter policy” (p. 1671) and 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

51.10.2 Steps

1

Choose Policies→Filter→Protocol List from the NFM-P main menu. The Filter - Protocol Lists form opens.

2

Click Create or select an existing policy and click Properties. The Filter - Protocol List (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the Protocol List Members tab and click Add Protocols. The Select Protocols form opens.

5

Move the required protocols to the Selected Protocols panel, to a maximum of 32.

6

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

51.11 To configure an IP Prefix List policy

51.11.1 Steps

1

Choose Policies→Filter→IP Prefix List from the NFM-P main menu. The Filter - IP Prefix Lists form opens.

2

Click Create or select an existing policy and click Properties. The Filter - IP Prefix List (Create|Edit) form opens.

3

Configure the parameters as required.

4

Configure IP prefix list member entries as required:

1. Click on the IP Prefix List Members tab and click Create, or choose an entry in the list and click Properties. The IP Prefix List Member (Create) form opens.
2. Configure the parameters.
3. Click Apply to save your entry and create another, or click OK to save your entry and close the form. The newly-defined member appears in the list on the IP Prefix List Members tab.

5

Configure IP prefix list excluded member entries as required:

Note: You cannot configure excluded members and apply path members in the same policy.

1. Click on the IP Prefix List Exclude Members tab and click Create, or choose an entry in the list and click Properties. The IP Prefix List Exclude Member (Create) form opens.
2. Configure the parameters.
3. Click Apply to save your entry and create another, or click OK to save your entry and close the form. The newly-defined excluded member appears in the list on the IP Prefix List Excluded Members tab.

6

Configure apply path member entries as required:

Note: You cannot configure apply path members and excluded members in the same policy.

1. Click on the Apply Path Members tab and click Create, or choose an entry in the list and click Properties. The Apply Path Member (Create) form opens.
2. Configure the required parameters.
3. Click Apply to save your entry and create another, or click OK to save your entry and close the form. The newly-defined apply path member appears in the list on the Apply Path Members tab.

7

Save your changes and close the Filter - IP Prefix List (Create|Edit) form.

8

To release and distribute the policy, click Search on the Filter - IP Prefix Lists form, select the newly created policy and click Properties. The Filter - IP Prefix List (Edit) form opens.

9

Perform the required steps in 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy.

10

To associate a newly-created IPv4 IP Prefix List with an ACL IP Filter policy, see 51.5 “To configure an ACL IP filter policy” (p. 1671). For ACL IPv6 filters, see 51.6 “To configure an ACL

IPv6 filter policy" (p. 1677). In either of those procedures, you can specify an IP Prefix List by configuring the Source IP Prefix and/or the Destination IP Prefix parameters for a filter entry.

11

To associate a newly-created IPv4 or IPv6 IP Prefix List with a CPM Filter policy, see the *NSP System Administrator Guide*.

12

Save your changes and close the forms.

END OF STEPS

51.12 To configure a Port List policy

51.12.1 Steps

1

Choose Policies→Filter→Port List from the NFM-P main menu. The Filter - Port Lists form opens.

2

Click Create or select an existing policy and click Properties. The Filter - Port List (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the Port List Members tab and click Create. The Port List Member, Site (Create) form opens.

5

Configure the Port parameter. This can either be a specific port number or a range of port numbers. For a specific port number, use the EQUAL operator from the drop-down menu. For a range of port numbers, use the RANGE operator.

6

Save your changes and close the form. The newly-defined member appears on the list.

7

Repeat Step 4 to Step 6 to create additional members, as required.

8

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: To associate a newly-created Port List with a CPM Filter policy, see the *NSP System Administrator Guide*.

END OF STEPS

51.13 To configure a DHCP Filter policy

51.13.1 Steps

1

Choose Policies→Filter→DHCP Filter from the NFM-P main menu. The DHCP Filter Policies list form opens.

2

Click Create or select an existing policy and click Properties. The DHCP Filter (Create|Edit) form opens.

3

Configure the parameters as required.

4

Configure a filter entry.

1. Click on the Filter Entries tab and click Create. The Entry, DHCP Filter (Create) form opens.
2. Configure the required parameters.
3. Click on the Filter Properties tab.
4. Configure the required parameters:

   The Match Condition parameter is configurable only when a DHCP Option has been selected.

   The Match String Type, Match Type, and Match String parameter are configurable only when the Match Condition parameter is set to Match String
5. Save your changes and close the form.

5

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Repeat Step 1 and choose the distributed policy in the list and click Properties. The DHCP Filter Global Policy (Edit) form opens.

7

Click on the Group Interfaces tab to view IES or VPRN group interfaces that are bound to the DHCP filter policy. See 78.19 “To configure a group interface on an IES” (p. 2449) and 79.37 “To configure a group interface on a VPRN” (p. 2586) for more information on binding IES and VPRN group interfaces to DHCP filter policies.

8

To view DHCP filter entry data.

1. Click on the Local Definitions tab and choose a local definition from the list and click Properties. The DHCP Filter Local Policy (Edit) form opens.
2. Select an entry from the list and click Properties. An Entry, DHCP filter form opens.
3. Close the forms after viewing the filter entry data.

END OF STEPS

51.14 To configure a DHCPv6 filter policy

51.14.1 Steps

1

Choose Policies→Filter→DHCPv6 Filter from the NFM-P main menu. The DHCPv6 Filter Policies list form opens.

2

Click Create or select an existing policy and click Properties. The DHCPv6 Filter (Create|Edit) form opens.

3

Configure the parameters as required.

4

Configure a filter entry.

1. Click on the Filter Entries tab and click Create. The DHCPv6 Filter Entry (Create) form opens.
2. Configure the required parameters.
3. Click on the Filter Properties tab.
4. Configure the required parameters.

If the Action parameter is set to Bypass Host Creation, the Action Flags parameters must be configured.

After the DHCP Option parameter has been configured, the Match Condition parameter is configurable.

If the Match Condition parameter is set to Match String, the Match String Type, Match Type, and Match String parameters are configurable.

5. Save your changes and close the form.

5

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Repeat Step 1 and choose the distributed policy in the list and click Properties. The DHCPv6 Filter Global Policy (Edit) form opens.

7

Click on the Group Interfaces tab to view IES or VPRN group interfaces that are bound to the DHCPv6 filter policy. See 78.19 “To configure a group interface on an IES” (p. 2449) and 79.37 “To configure a group interface on a VPRN” (p. 2586) for more information on binding IES and VPRN group interfaces to DHCPv6 filter policies.

8

To view DHCP filter entry data.

1. Click on the Local Definitions tab and choose a local definition from the list and click Properties. The DHCPv6 Filter Local Policy (Edit) form opens.
2. Select an entry from the list and click Properties. An DHCPv6 filter form opens.
3. Close the forms after viewing the filter entry data.

END OF STEPS

51.15 To configure a Redirect Filter policy

51.15.1 Purpose

Redirect Filter policies allow specifying multiple redirect target destinations and defining health check test methods used to validate the ability for a given destination to receive redirected traffic. This destination monitoring allows a router to react to target destination failures.

51.15.2 Steps

1

Choose Policies→Filter→Redirect Filter from the NFM-P main menu. The Redirect Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The Redirect Filter Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

If you configured the Router Instance Type parameter to VPRN, click Select to assign a VPRN pointer to the policy.

5

Click on the Destination Addresses tab and click Create. The Destination (Create) form opens.

6

Configure the parameters as required.

If you enable the Enable Unicast Route Reachability Test parameter, no other health check test can be performed. Go to Step 8.

If you set the Address Type parameter to IPv4, go to Step 7. For IPv6, you can only perform Step 7 c.

7

Perform one or more of the following test types:

a. Configure an SNMP test:

1. Click on the SNMP Tests tab and click Create. The SNMP Test (Create) form opens.
2. Configure the required parameters.
3. Click on the SNMP Responses tab and click Create. The SNMP Response (Create) form opens.
4. Configure the required parameters.
5. Save your changes and close the form.
6. Repeat 3 to 5 to create addition SNMP responses for the SNMP test.
7. Repeat 1 to 6 to create addition SNMP tests for the redirect filter.

b. Configure a URL test:

1. Click on the URL Tests tab and click Create. The URL Test (Create) form opens.
2. Configure the required parameters.

3. Click on the URL Responses tab and click Create. The URL Response (Create) form opens.
4. Configure the required parameters.
5. Save your changes and close the form.
6. Repeat 3 to 5 to create addition URL responses for the URL test.
7. Repeat 1 to 6 to create addition URL tests for the redirect filter.

c. Configure a Ping test:

1. Enable the Enable Ping Test parameter.
2. Click on the Ping Test tab.
3. Configure the required parameters.
4. Save your changes and close the form.

For supporting NEs, the results of Ping tests can be shared among multiple destination addresses, when a redirect policy binding is configured for those addresses. See 51.16 “To configure a Redirect Policy Binding” (p. 1696).

8

Save your changes and close the form.

9

Repeat Step 2 to Step 8 to create additional IPv4 or IPv6 destination addresses as required.

10

See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

11

To associate a newly-created IPv4 redirect filter policy with an ACL IPv4 filter policy, see 51.5 “To configure an ACL IP filter policy” (p. 1671). For an IPv6 redirect filter policy, see 51.6 “To configure an ACL IPv6 filter policy” (p. 1677). You must set the Action parameter to the forward (Redirect Filter) option to enable an association. The ACL Filter Entries tab on the Redirect Filter Policy form lists all such associations for distributed policies.

12

When a Redirect Filter policy has been distributed to the required NEs, you can open local definitions from the Local Definitions tab of the global policy and examine the tests for each contained destination. A Refresh button is available on the individual test property forms to get the Last Response details.

13

Close the forms.

END OF STEPS

51.16 To configure a Redirect Policy Binding

51.16.1 Purpose

Redirect policy bindings allow configuration of an association between destination addresses in redirect filter policies, so that results of Ping tests can be shared.

The redirect filter policies and addresses configured in a redirect policy binding must already exist; see 51.15 “To configure a Redirect Filter policy” (p. 1693).

51.16.2 Steps

1

Choose Policies→Filter→Redirect Policy Binding from the NFM-P main menu. The Redirect Binding Policies form opens.

2

Click Create, or choose an existing redirect policy binding and click Properties. The Redirect Policy Binding (Create|Edit) form opens.

3

Configure the parameters on the General tab.

4

Click on the Redirect Policy Destination tab.

5

Configure redirect policy destination addresses for the binding. The required redirect policies and destination addresses must already exist.

1. Click Create, or choose an item in the list and click Properties. The Redirect Policy Binding Destination form opens.
2. Select a Redirect Policy and IP Address for the binding.
3. Save your changes and close the form.
4. Configure additional addresses as required.
5. To remove any address that is not required in the binding, choose the entry in the list and click Delete.

6

Save your changes and close the forms.

7

See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

51.17 To configure an ACL VLAN Filter policy

51.17.1 General Information

An ACL VLAN Filter policy can be configured on a 2-port ring MDA on the 7705 SAR-8 or 7705 SAR-18. Only the first two ports in network mode on the 2-port ring MDA can be configured with an ACL VLAN filter policy.

51.17.2 Steps

1

Choose Policies→Filter→ACL VLAN Filter from the NFM-P main menu. The ACL VLAN Filter Policies form opens.

2

Click Create or select an existing policy and click Properties. The ACL VLAN Filter, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Configure a filter entry.

1. Click on the Filter Entries tab and click Create.The Entry, ACL VLAN Filter, Global Policy (Create) form opens.
2. Configure the required parameters.
3. Click on the Filter Properties tab.
4. Configure the required parameters.
5. Save the changes and close the form.

5

Repeat Step 4 as required to create additional filters.

6

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

7

To associate the ACL VLAN Filter Policy with an Ethernet port on a 2-port ring MDA on a 7705 SAR-8 or 7705 SAR-18, see 16.24 “To configure Ethernet ports” (p. 599).

END OF STEPS

51.18 To configure a System Filter

51.18.1 Purpose

A system filter allows operators to configure a filter chain for packet matching. In this chain, the active system filter policy rules are evaluated first. If no match occurs, only then are the rules of any chained filter policies evaluated.

The system filter policies supports all IPv4/IPv6 filter policy match rules and actions. However, system policy entries cannot be LI or mirror sources. A system filter policy also does not support Radius, flowspec, or Gx inserted entries. In addition, a system filter policy also requires chassis mode D to be set on an NE to which it is deployed.

See 51.1 “Filter policies” (p. 1663) for more information on system filters.

51.18.2 Steps

1

Perform one of the following procedures to configure and distribute an IP filter policy for use as a system filter policy on the required NEs:

a. 51.5 “To configure an ACL IP filter policy” (p. 1671) for an ACL IP filter policy

b. 51.6 “To configure an ACL IPv6 filter policy” (p. 1677) for an ACL IPv6 filter policy

Whether the policy you create will be an active system filter policy or a chained system filter policy is determined by setting the Scope and Chain to System Filter parameters.

Note: If you are configuring an active system filter policy, then you must set the Scope parameter to system when you create the new filter policy.
If you are configuring a chained system filter policy, then you must set the Scope parameter to either template or exclusive when you create the new filter policy. The Chain to System Filter parameter must also be enabled.
To change an existing filter policy’s Scope parameter to or from the system option, the policy must have no Filter Entries configured.

2

Perform one of the following:

a. To create or configure an Active System Filter, go to Step 3.

b. To view or configure a Chained Filter, go to Step 10.

Create or configure an Active System Filter

3

On the equipment tree, right-click on the NE to which you want to apply the active system filter and choose Properties. The Network Element (Edit) form opens.

4

Click on the Globals tab and then the System Filter tab.

5

Click on the Active System Filters tab and then click Create. The System Filter (Create) form opens.

6

Select either IP or IPv6 as the required System Filter Type.

7

Select the IP filter policy you created or configured in Step 1 as the System Filter.

8

Click OK to apply the configuration and close the System Filter (Create) form. The new active system filter appears in the list.

Note: You can also delete or configure the properties of an existing active system filter from this form.

9

Close the form.

View or configure a chained system filter

10

On the equipment tree, right-click on the NE to which you distributed the chained system filter and choose Properties. The Network Element (Edit) form opens

11

Click on the Globals tab and then the System Filter tab.

12

Click on the Chained Filters tab and then either the IP or IPv6 tab, as required.

All ACL IP template or exclusive filter policies that have been distributed to this NE and have an enabled Chain to System Filter parameter are displayed. No further action is required. These are the chained system filters.

13

To change the configuration of an existing chained system filter, select it from the list and click Properties.

Refer to either 51.5 “To configure an ACL IP filter policy” (p. 1671) or 51.6 “To configure an ACL IPv6 filter policy” (p. 1677) for configuration details, as required.

Note: You cannot delete a chained system filter from this form. Chained system filters can only be deleted using the ACL IP or ACL IPv6 Filter Policy selection form.

14

Close the form.

END OF STEPS

51.19 To copy filter policy filter entries

51.19.1 Purpose

Perform this procedure to copy the filter entries of one filter policy to another filter policy. The supported filter policies include:

- ACL MAC Filter
- ACL IP Filter
- ACL IPv6 Filter
- DHCP Filter

51.19.2 Steps

1

Choose Policies→Filter→*Filter type* from the NFM-P main menu, where *Filter type* is the type of filter that you want to copy a filter entry from. The appropriate policies manager form opens.

2

Select a policy, then click Properties. The filter policy properties form opens.

3

Click on the Filter Entries tab, click Search and select one or more filter entries from the list.

4

Click Copy Entries. The Copy Filter type Filter Entries form opens and the copied filter entry appears in the Copy Policy Entry Instances panel.

5

Configure the following parameters as required.

- Enable Copied Entries to Override Existing Entries with the Same ID
- Distribute the Global policy after Applying Form

6

Select the local policies that you want to copy the filter entry to from the Available Policies list. Only local policies with a distribution mode of Local Edit Only will be listed.

7

Click on the right-pointing arrow to move the selected policies to the Copy Entries to these Targeted Policies panel.

Note: If you enabled the Distribute the Global policy after Applying Form parameter and one of the Targeted Policies is a global policy, then that global policy will subsequently be distributed to all NEs with local definitions of the policy.
However, if the Distribution Mode of a local policy is set to Local Edit Only, the associated NE will not be allowed to receive the distribution of the global policy. You must set the local policy's Distribution Mode to Sync With Global if you want the NE to receive the global policy distribution.
See 49.6 "To release and distribute a policy" (p. 1476) for detailed information on policy distribution.

8

Click OK. The filter entries are copied to the targeted policies and the Copy Filter type Filter Entries form closes.

Note: If you copy a filter entry with an Entry Type of Embedded, the copies sent to the targeted policies will have an Entry Type of Normal. See 51.9 "To configure an embedding filter with embedded filter policies" (p. 1686) for information on embedded filters.
If an NE that contains one of the targeted policies is down, it cannot receive the copied filter. A deployment failure message will be issued.

9

Close the filter policy properties form.

END OF STEPS

51.20 To configure a ACL Filter Log policy

51.20.1 Purpose

You can configure a ACL Filter Log policy to define:

- where log information for all actions performed on 7210 and 7x50 NEs that match ACL MAC, ACL IP, ACL IPv6, and CPM filter entry criteria are written (memory or Syslog)
- how many log entries can be stored
- what action is performed when the log files meet the specified threshold.

51.20.2 Steps

1

Choose Policies→Log Configuration. The Manage Log Configuration Policies form opens.

2

Click Create→ACL Filter Log or select an existing ACL Filter Log (Node Log Policy) and click Properties. The ACL Filter Log, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

The SysLog Id, Log Summary Enabled, and Log Summary Criterion parameters are configurable when you choose the SysLog option as the Log Destination.

The Maximum Number of Entries parameter is configurable when you choose the Memory option as the Log Destination.

4

Save your changes. The ACL Filter Log, Global Policy (Create|Edit) form refreshes displaying the Policy Configuration panel.

5

Click Switch Mode to release and distribute the policy to the required NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Close all open forms.

7

Associate the Filter log with a filter entry for any of the following filter types by configuring the Log ID parameter in the associated procedure:

- CPM Filter. See the “To configure a CPM filter” procedure in the *NSP System Administrator Guide*.

- ACL MAC Filter. See 51.4 “To configure an ACL MAC filter policy” (p. 1668).
- ACL IP Filter. See 51.5 “To configure an ACL IP filter policy” (p. 1671).
- ACL IPv6 Filter. See 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

  The Log ID you specify in any of these procedures will refer to the ACL Filter Log’s ID parameter you configured in Step 3.

END OF STEPS

51.21 To configure a GRE tunnel template

51.21.1 Before you begin

A GRE tunnel template specifies a set of GRE encapsulation parameters to use when the primary action for ACL IP and ACL IPv6 filter policies is set to Forward (GRE Tunnel).

GRE tunnel templates are assigned to ACL IP filter policies and ACL IPv6 filter policies; see 51.5 “To configure an ACL IP filter policy” (p. 1671) and 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

51.21.2 Steps

1

Choose Policies→Filter→GRE Tunnel Template from the NFM-P main menu. The Filter - GRE Tunnel Template form opens.

2

Click Create, or click Search, select an existing template, and click Properties. The GRE Tunnel Template (Create|Edit) form opens.

3

If you are creating a new template, configure the Template Name and Description parameters, and click Apply.

The Template Name and Description parameters are configurable only when creating a new template.

4

Click on the IPv4 tab and configure the required parameters.

5

Click on the IPv6 tab and configure the required parameters.

6

In the IP Destination Address panel, configure one or more IP destination addresses.

When the GRE tunnel template is assigned to an ACL IP or IPv6 filter policy, traffic matching the associated IPv4 or IPv6 filter is hashed across all available destination addresses. If no

destination address is available, then matching traffic follows the configured PBR Down Action Override action, if configured; see 51.5 “To configure an ACL IP filter policy” (p. 1671) and 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

You cannot modify an existing address in the template. You must delete the existing address and create another.

1. Click Create. The IP Destination Address, Site (Create) form opens.
2. Configure the IPv4 Destination Address parameter.

   You can modify address entries before the policy is saved, but not after.
3. Save your changes and close the form. The GRE Tunnel Template (Create|Edit) form is displayed.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

51.22 To configure a Syslog policy

51.22.1 Purpose

You can configure a Syslog policy that is used by the ACL Filter Log policy to define the destination details for log messages such as the target address and target UDP port.

51.22.2 Steps

1

Choose Policies→Log Configuration. The Manage Log Configuration Policies form opens.

2

Click Create→Sys Log or select an existing Sys Log Target (Node Log Policy) and click Properties. The Sys Log Target, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click Apply. The Sys Log Target, Global Policy (Create|Edit) form refreshes displaying the Policy Configuration panel.

5

Click Switch Mode to release and distribute the policy to the required NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Close all open forms.

7

See 51.20 “To configure a ACL Filter Log policy” (p. 1702) to associate the SysLog with a ACL Filter Log.

END OF STEPS

51.23 To configure a Log ID Profile Policy

51.23.1 Purpose

You can configure a Log ID Profile policy that is used to define various source and destination details to retrieve specific log files after the node reboot.

51.23.2 Steps

1

Choose Policies→Log Configuration. The Manage Log Configuration Policies form opens.

2

Click Create→Log ID Profile or select an existing Log ID Profile (Node Log Policy) and click Properties. The Log ID Profile, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

- For the Destination parameter, you can assign one of the following: Console, Syslog, SNMPTraps, File, Memory, Netconf, and SubscribedCli.
- You can assign an Event Filter Log to the Log ID Profile.
- The Destination parameter can only be set once and cannot be modified. Once set, it becomes read-only and is not a part of the Audit operation. However. the associated fields and other attributes are highlighted, if different, and they can be modified during the Audit process.
- The Source and Event Filter Log can be modified.
- If Destination is set as Syslog, you can assign the Syslog and Python policy.
- If Destination is set as SNMPTraps, the Log ID should be similar as the SNMP group id. For example, if the SNMPTrap group 53 is created on the node, you can create the Log ID as 53 and set its destination as SNMPTrap.
- When the source is set as *Li*, the distribution gets successful only on the nodes that has or supports the *Li* configuration

4

Click Apply. The Log ID Profile, Global Policy (Create|Edit) form refreshes displaying the Policy Configuration panel.

5

Click Switch Mode to release and distribute the policy to the required NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Close all open forms.

END OF STEPS

51.24 To configure an LI log ID profile

51.24.1 Purpose

You can configure an LI log ID profile to define various source and destination details to retrieve specific log files after an NE reboot. You must be an LI user in order to configure an LI log ID profile.

51.24.2 Steps

1

Choose Policies→LI Log Configuration. The Manage LI Log Configuration Policies form opens.

2

Create Create. The Li Log ID Profile, Global Policy (Create) form opens.

3

Configure the required parameters and select the event filter log.

4

Click Apply. The Li Log ID Profile, Global Policy (Create|Edit) form refreshes displaying the Policy Configuration panel.

5

Click Switch Mode to release and distribute the policy to the required NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

Close all open forms.

END OF STEPS

51.25 To configure an Event Filter Log Policy

51.25.1 Purpose

You can configure an Event Filter Log policy to define an application by adding its details as an entry to retrieve specific log files after the reboot.

51.25.2 Steps

1

Choose Policies→Log Configuration. The Manage Log Configuration Policies form opens.

2

Click Create→Event Filter Log or select an existing Event Filter Log (Node Log Policy) and click Properties. The Event Filter Log, Global Policy (Create|Edit) form opens.

3

In the General tab, configure the parameters as required.

4

Configure the Entry tab.

1. Click Create or select an existing Event Filter Log, Global Policy and click Properties. The Event Filter Log Entry (Create|Edit) form opens.
2. Configure the parameters as required.

   Note: Multiple applications ids are not allowed in one entry, but you can modify the existing application id.
3. You can also define operators and regular expressions for the following parameters: Application, Number, Severity, Subject, Router, and Message.
4. Click Apply.

5

Click Apply. The Event Filter Log, Global Policy (Create|Edit) form refreshes displaying the Policy Configuration panel.

6

Click Switch Mode to release and distribute the policy to the required NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

7

Close all open forms.

END OF STEPS

52 Multicast policies

52.1 Multicast policies

52.1.1 Overview

You can use the NFM-P to create multicast policies that reduce bandwidth requirements and regulate multicast traffic flow by simultaneously delivering a single stream of information to multiple receivers. Multicast traffic flow is based on the concept of a group; a multicast group is an arbitrary group of receivers that receive the same data traffic.

You can configure the following NFM-P multicast policy types:

- Multicast package
- Egress multicast group
- Multicast CAC
- Ingress multicast path management
- SSM translate
- 7210 SAS multipoint bandwidth management
- MCAC interface

52.2 Multicast package policies

52.2.1 Overview

Multicast package policies are used to assign a common set of multicast groups to NEs in an MVR VPLS. Multicast package policies ensure that the multicast groups on the devices of one or more MVR VPLS instances are consistent. See Chapter 27, “NE routing and forwarding” for information about configuring multicast groups on individual devices using a multicast routing policy.

When a multicast package policy is distributed to an NE, a multicast routing policy is configured with the appropriate multicast prefix list, which is based on the multicast groups in the multicast package policy. Multicast package policies are only used to distribute the policies to one of the two NEs. Policy discovery and synchronization are not supported.

A multicast package policy has zero or more multicast groups associated with it. The multicast groups are distributed to the NE as part of the policy. Only root catalog multicast packages can be distributed to the NE.

52.2.2 Multicast package policy design considerations

Consider the following when creating multicast package policies:

- one root multicast package policy should be created for a ring or multiple rings, which contains the multicast addresses of all BTV channels distributed in the ring
- specify this root multicast package policy for the BTV (MVR) VLAN

- create other multicast package policies from the root package by clicking on the Copy button on the configuration form, or by clicking on the More Actions button and choosing Copy, if the Copy button is not visible.
- use the child policies to differentiate the types of broadcast TV services offered, for example, basic service and premium service

52.2.3 Distributing multicast package policies

When you distribute a multicast package policy to an NE, the NFM-P maps one multicast package policy to one routing policy statement. Each multicast policy group entry is mapped to one routing policy prefix list member.

For example, the ID of a multicast package policy is 2. When the policy is distributed to an NE, the NFM-P automatically creates a routing policy statement for the device. The name created for the routing policy statement is SAM_MPP_2, where 2 is the multicast package policy ID.

When the policy is distributed to an NE, the ID of the routing policy statement entry is always 1. A name is generated by the NFM-P to identify the prefix list in the routing policy statement entry. The name is SAM_MPG_2_1, where 2 is the multicast package policy ID and 1 is the routing policy statement entry ID. The prefix list name is referenced in the routing policy statement entry as the multicast group prefix list name.

In the same example, the multicast address for the multicast group of a multicast package policy is 224.0.0.0. When the multicast package policy is distributed to an NE, the NFM-P automatically creates a routing policy prefix list member for the device. The prefix list member address is 224.0.0.0.

Since the multicast package policy is not resynchronized, you must avoid duplicate names for routing policies.

If a multicast group does not have a policy group entry, the multicast group is not distributed to the device as a prefix member.

If a multicast package policy is used by any MVR VPLS, it cannot be removed from the NFM-P.

52.3 Egress multicast group policies

52.3.1 Overview

With standard egress multicasting, the entire egress forwarding plane is used on a per destination basis to provide ACL, mirroring, QoS, and accounting for each path with the associated receivers. This process is very resource intensive. Also, when congestion occurs at the egress forwarding plane, unicast discards may be indiscriminate of the forwarding priority.

The NFM-P provides an alternative egress multicasting method that increases egress replication performance and scalability by combining multiple destinations into a single egress forwarding pass as an egress multicast group (EMG). You specify the EMG requirements using a global EMG policy

which defines the group destination SAPs in a chain. The chain receives a single multicast packet from the egress forwarding plane with only a single pass.

The EMG policy is applied to an L2 access interface that is created on an Ethernet port with null or dot1q encapsulation. SAP membership in an EMG policy has the following requirements:

- The SAP must belong to an SHG.
- The L2 access interface egress ACL filter must match that of the EMG.
- The access port and the EMG must have the same encapsulation type.

When a SAP becomes a member of an EMG, the common parameters that it shares with other EMG members cannot be changed. On the other hand, a modification to an EMG common requirement changes that parameter in all SAPs in the chain. The same common required traits apply to egress filters: individual SAP filters cannot be changed, and any change to the EMG egress filter parameters affects the egress filters in all member SAPs. In addition, if an egress filter is not defined in the EMG common parameter requirements, a new member SAP cannot have an egress filter.

The following restrictions also apply to an EMG:

- egress multicast group policies cannot be applied to SR NE variants at Release 15.0 R1 and higher.
- only IPv4 egress filters can be applied to an EMG
- up to 30 EMGs can be created on a router
- an EMG can be deleted only if it has no SAP members
- the number of members in an EMG (the chain limit) is restricted to 30, with an optimal length of 10 to 16

52.4 Multicast CAC policies

52.4.1 Overview

Multicast CAC policies control the bandwidth consumed by BTV distribution services. Bandwidth control helps manage network congestion and maintain QoS standards. The multicast CAC function is supported on IGMP, PIM, or MLD interfaces, IGMP or MLD group interfaces, and in the case of BTV distribution, on VPLS SAPs and SDPs where IGMP snooping is enabled.

A multicast CAC policy manages the bandwidth consumed by BTV services on both the access node link and specific links in the aggregation network. Routers in the path are configured to maintain certain limits on broadcast bandwidth and can limit the number of channels simultaneously sent on both the second mile link and the network link.

The total bandwidth usage of an IGMP host is calculated using the following logic:

1. Retrieve all channels (i.e. groups) from the IGMP host.
2. For each channel (i.e group), find the corresponding MCAC policy channel object defined with a matching source address.

The logic functions on the best match for the channel with the source address and prefix in the MCAC policy channel object. If the match fails, bandwidth is calculated from a matched MCAC channel object with a default (null) source address.

3. Retrieve the bandwidth assigned to each of the corresponding MCAC channels.
4. The sum of all bandwidth values is the total bandwidth usage of the host.

52.5 Ingress multicast path management policies

52.5.1 Overview

Ingress multicast path management policies are used to manage ingress multicast path bandwidth settings, path decisions, and destination server details that collect and analyze IGMP events associated with residential subscribers. You can create the following ingress multicast path management policy types:

52.5.2 Ingress multicast bandwidth policy

Ingress multicast bandwidth policies are used to manage the ingress multicast path bandwidth of the multicast forwarding paths into the switching fabric. When Multicast Path Management is disabled (the default), two paths are available:

- a high-priority path, on which packets from queues classified as “expedited” are forwarded
- a low-priority path on which packets from queues classified as “non-expedited” are forwarded

When Multicast Path Management is enabled, an ingress multicast bandwidth policy can be used to manage the flow of multicast traffic through an MDA (IOM, IOM 2, or XMA) or the forwarding plane of an 2 x XP MDA IOM 3, IMM, or XCM.

The maximum number of bandwidth policies per NE is 32, including the default policy.

MDAs can be configured to use previously defined ingress multicast bandwidth policies. However, any path limits specified in the selected policy can be overridden for each MDA, if required.

52.5.3 Ingress multicast info policy

Ingress multicast info policies are used to define how each multicast channel is handled by NEs. The policy is assigned to a VPLS/VPRN service site or default routing instance, but the policy is actually used by the Ingress Multicast Bandwidth Manager, the ECMP Path Manager, and the Egress Multicast CAC Manager to determine the path through the switch fabric and to make decisions on joins to multicast streams.

The maximum number of ingress multicast info policies per NE is limited to 32, including the default policy.

MDAs can be configured to use previously-defined ingress multicast bandwidth policies. However, any path limits specified in the selected policy can be overridden for each MDA, if required.

Ingress multicast forwarding on L2-snooped and L3-routed IP multicast traffic can be explicitly configured by applying an ingress multicast info policy to a VPLS site (all types of VPLS are supported), a VPRN site, or the default routing instance.

An ingress multicast info policy consists of one or multiple named Bundles, which in turn, contain Channel ranges with possible overrides for individual channels.

- Bundles

  A Bundle groups a set of explicit multicast channels (or channel ranges) into a common bandwidth context for CAC functions (such as join decisions) using common preferences. The channels in the bundle are managed as a defined percentage of the available bandwidth.

  Bundles also simplify provisioning, since the default characteristics of the bundle channels are specified on the bundle level. These characteristics can be overridden at the Channel range level, or explicitly per channel.

  Each ingress multicast info policy has a default bundle named “default”. It cannot be edited or removed. Any multicast channel that fails to match a channel range within an explicit operator-defined bundle is associated with this default bundle.

  Note:

  The maximum number of Bundles per Info policies is limited to 32. This limit includes the default bundle, leaving 31 operator-defined bundles per Info policy.

- Channels and channel ranges

  Channel ranges are used to define a set of multicast channels contained in a bundle and to override the default channel settings in the containing bundle.

  A channel range is defined by a start destination multicast IP address and an end destination multicast IP address (both IP v4 and IPv6 are supported, but the start and end addresses must be of the same type). A channel in this context is a channel range where the start address and end address are identical.

  Note:

  After you create a channel range, it is not possible to modify the start address or end address.

  A channel range can contain multiple channel overrides. A channel override is used to specify an explicit setting for a channel within the range. The channel override is identified by a destination multicast IP address which falls between the start and end IP addresses of the channel range.

  Note:

  The 7450 ESS does not support channel ranges and channel overrides having IPv6 addresses. Therefore, the NFM-P does not allow the creation of such channel ranges and channel overrides on local ingress multicast info policies. In addition, the NFM-P removes any such channel ranges and channel overrides when synchronizing local ingress multicast info policies with their respective global policies.

- Dual stream selection

  Dual stream selection support allows a single multicast stream to be duplicated into two different transmission paths. The two paths may have different transmission characteristics, such as latency and jitter. Rather than select one stream for retransmission to the client, the duplicate stream protection feature evaluates each stream packet-by-packet, selecting the packet that first arrives (and is valid) for retransmission.

- Video quality monitoring

VQM to provides the functionality to analyze the quality of a video stream just prior to reaching the STB of client in a IPTV network. Statistics reports and alarms can be generated for VQM, providing operators with a view of multicast video quality at the last mile of distribution.

52.5.4 Ingress multicast reporting policy

Ingress multicast reporting destination policies are used to specify the destination server for the collection of multicast reports. The policies are applied to residential subscriber IGMP policies.

52.6 SSM translate policies

52.6.1 Overview

Source-Specific Multicast policies allow the receiver of multicast packets to specify desired source addresses to the router, so that only packets originating from the specified sources are delivered. IGMP is used by the client to inform the router of the specified addresses.

The SSM translate policy allows you to create global sets of SSM translate items that can be mass-delivered to many routers at once. Every policy has one or more SSM translate items. When the policy is distributed, all of these items are deployed to the target base routers.

52.7 7210 SAS multipoint bandwidth management policies

52.7.1 Overview

A 7210 SAS Multipoint Bandwidth Management policy defines rate limits and burst sizes for multipoint queues, to balance flows with unicast traffic. The policy defines an aggregate rate for all multipoint traffic in the NE. Rates for individual queues are configured as a percentage of the aggregate rate.

In the 7210 SAS-X, SAP ingress and network ingress multipoint traffic is separated from unicast traffic and forwarded into multipoint queues before being replicated to the appropriate egress port. There are eight multipoint queues per NE, which are shared by all services. Packets are forwarded to multipoint queues based on the forwarding classes assigned during ingress. The FC-to-queue mapping is not user-configurable.

The NFM-P provides a default multipoint bandwidth management policy. The default policy is applied if no user-defined policy is explicitly assigned. The default policy specifies maximum rate and burst size settings. See the NE documentation for more information.

52.8 MCAC interface policies

52.8.1 Overview

An operator can configure one or more MCAC (Multicast CAC) interface policies under a global MCAC context to prioritize and limit the mandatory and unconstraint bandwidth values across the supported interfaces in one or more VPRN/EIS services via a Global Route Table (GRT).

When a single MCAC interface policy is used by multiple interfaces in the network, channel admittance on one interface impacts the available bandwidth on all interfaces that share the same MCAC interface policy. Changing policy values behaves the same way as if the change to

mandatory or unconstrained bandwidth was applied on an interface level.

If a MCAC interface policies is deployed together with per interface unconstrained and mandatory bandwidth constraints, a given multicast group must satisfy the MCAC conditions defined by the MCAC interface policy and by the interface to be admitted. If ESM is deployed on an interface and per-subscriber MCAC is enabled, then the per-subscriber MCAC must also be satisfied.

You can configure MCAC interface policies on 7x50 devices; this policy type can be associated with the following object/interface/policy types:

- VPLS SAP
- VPLS Spoke SDP and MVPL Spoke SDP
- IGMP Interface on base router and VPRN service
- MLD Interface on base router and VPRN service
- PIM Interface on base router and VPRN service
- IGMP Group Interface on IES and VPRN services
- MLD Group Interface on IES and VPRN services
- MSAP policy

52.9 To configure a multicast package policy

52.9.1 Purpose

You can use multicast package policies to:

- define a set of broadcast channels that are multicast across a ring group in a BTV VLAN
- assign a common set of multicast groups to all NEs in an MVR VPLS to ensure accuracy and consistency

52.9.2 Steps

1

Choose Policies→Multicast→Multicast Package from the NFM-P main menu. The Multicast Package Policies form opens.

2

Click Create or select an existing policy and click Properties. The Multicast Package Policy, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the Multicast Groups tab and click Create. The Multicast Group, Multiple Package Policy (Create) form opens.

Note: Each multicast group that is specified must be preconfigured in the PIM configuration. See the PIM configuration procedures in Chapter 28, “Routing protocol configuration”.

5

Configure the parameters as required.

6

Save the form. The Multicast Package Policy, Global Policy (Create|Edit) reappears.

7

Click Apply to save the policy. Additional buttons appears on the form.

8

Click Search, select a policy in the list and click Distribute to manually release and distribute the policy locally to devices in the ring. See 49.6 “To release and distribute a policy” (p. 1476) for additional information. Policies are also automatically distributed to devices when they are used by resources on the device.

Note: When any new devices are added to the ring, the policy is automatically distributed to any new MVR VLAN service or to any new MVR VLAN service sites in the ring.

9

Close the form.

Note: You can create other multicast package policies from the root policy by clicking on the Copy button, as required. These child policies of the root policy can be used to help differentiate the types of BTV services being created, for example, basic services and premium services.

END OF STEPS

52.10 To configure an egress multicast group policy

Note: An egress multicast group policy cannot be applied to SR NE variants at Release 15.0 R1 and higher.

52.10.1 Steps

1

Choose Policies→Multicast→Egress Multicast Group from the NFM-P main menu. The Egress Multicast Group Policies form opens.

2

Click Create or select an existing policy and click Properties. The Egress Multicast Group, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the SAP Common Requirements tab and configure the parameters as required.

5

Click Select in the Egress Filter panel to choose an egress ACL IP filter policy and click OK. See 51.5 “To configure an ACL IP filter policy” (p. 1671) for information about configuring an ACL IP filter policies.

Note: The egress ACL IP filter policy that you choose for this egress multicast group must be the same as the egress ACL IP filter policy specified for a member SAP.

6

Click Select in the Egress Filter IPv6 panel to choose an egress IPv6 filter policy and click OK. See 51.6 “To configure an ACL IPv6 filter policy” (p. 1677) for information about configuring IPv6 filter policies.

7

Click Apply to save the policy. Additional buttons appears on the form.

8

Click Search, select a policy in the list and click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

Note: To associate a SAP on the device with an EMG, the access port encapsulation type on the device must be the same as the encapsulation type that you specified in Step 4.

9

Close the form.

END OF STEPS

52.11 To configure a multicast CAC policy

52.11.1 Steps

1

Choose Policies→Multicast→Multicast CAC from the NFM-P main menu. The Multicast CAC Policies form opens.

Note: Multicast CAC policies are not supported on VSR or 7705 SAR-Hm chassis types.

2

Click Create or select an existing policy and click Properties. The Multicast CAC, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the Bundles tab and click Create to associate a multicast bundle with the multicast CAC policy.The Multicast CAC Bundle (Create) form opens.

5

Configure the parameters as required.

6

Click on the Channels tab and click Create. The Multicast CAC Channel, Multicast CAC Bundle (Create) form opens.

7

Configure the parameters as required.

8

Save your changes. The Multicast CAC Bundle (Create) form reappear and lists the newly configured channel.

9

Click Search, select a policy in the list and click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

10

Close the form.

End of steps

52.12 To configure an ingress multicast bandwidth policy

52.12.1 Steps

1

Choose Policies→Multicast→Ingress Multicast Path Management from the NFM-P main menu. The Ingress Path Management Policies form opens.

2

Click Create→Multicast Bandwidth Policy or select an existing policy and click Properties. The Multicast Bandwidth Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the T1 Paths tab and choose one of the primary, secondary, and ancillary paths and click Properties. The Bandwidth Policy Path (Create) form opens.

5

Configure the parameters as required.

6

Save your changes.

7

Click on the T2 Paths tab and choose the primary or secondary path and click Properties.The Bandwidth Policy Path (Create) form opens.

8

Configure the parameters as required.

9

Save your changes.

10

Click Apply to save the policy. Additional buttons appears on the form.

11

Click Search, select a policy in the list and click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for more information. Policies are also automatically distributed to devices when they are used by resources on the device.

12

Close the form.

END OF STEPS

52.13 To configure an ingress multicast information policy

52.13.1 Steps

1

Choose Policies→Multicast→Ingress Multicast Path Management from the NFM-P main menu. The Ingress Path Management Policies form opens.

2

Click Create→Multicast Info Policy or select an existing policy and click Properties. The Multicast Info Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the Bundles tab and click Create. The Info Policy Bundle, Multicast Info Policy (Create) form opens.

5

Configure the parameters as required.

6

Click on the Video Defaults tab.

7

Configure the parameters as required.

8

Configure the stream selection parameters, if enabled for the video group.

Note: Stream selection parameters are only available if you are performing the configuration on a local policy definition. After the changes are applied to the local policy, the policy can be switched back to Sync With Global mode and the local changes made to the policy are retained. See Chapter 49, “Policies overview” for information on policy distribution modes.

Proceed as follows:

1. Enter a unicast IPv4 address for the Source Address parameter.
2. Click Select to choose a primary access interface and click OK.
3. Click Select to choose a secondary access interface and click OK.

   The primary Interface and secondary Interface must be different.

9

Configure VQM if the Analyzer parameter in the VQM panel is enabled for the video group.

1. Configure the Description parameter.
2. Click on the VQM Config tab.
3. Configure the parameters as required.

10

Click on the Tunnel Interfaces tab.

11

Click Select to choose a Multicast Path Management Bundle and click OK.

12

Click Select to choose an LDP Tunnel Interface and click OK.

13

Configure the Ingress LER parameter.

You can configure either the P2MP LSP Name parameter or the P2MP ID for LDP parameter.

The Ingress LER parameter is automatically configured when the P2MP ID for LDP parameter is configured.

14

Click on the Channels tab and click Create.The Info Policy Channel Range, Info Policy Bundle (Create) form opens.

15

Configure the parameters as required.

The 7450 ESS does not support channel ranges and channel overrides having IPv6 addresses. Therefore, the NFM-P does not allow the creation of such channel ranges and channel overrides on local Ingress Multicast Info policies. In addition, the NFM-P removes any such

channel ranges and channel overrides when synchronizing local Info policies with their respective global policies.

16

Click on the Video tab.

17

Configure the parameters as required.

Note: If the Stream Selection parameter is enabled for the video group, repeat Step 8 . If the Analyzer parameter in the VQM panel is enabled for the video group, repeat Step 9 .

18

Click on the Tunnel Interfaces tab and repeat Step 11 to Step 13 if required.

19

Click on the Channel Override tab and click Create to specify overrides for explicit channels within the Channel Range. The Info Policy Channel Override (Create) form opens.

20

Configure the parameters as required.

21

Save the changes.

22

Click on the Video tab.

23

Configure the parameters as required.

Note: If the Stream Selection parameter is enabled for the video group, repeat Step 8 . If the Analyzer parameter in the VQM panel is enabled for the video group, repeat Step 9 .

24

Click on the Tunnel Interfaces tab and repeat Step 11 to Step 13 if required.

25

Click on the Video Interfaces tab and click Create if you need to configure a video interface. The Multicast Video Interface (Create) form opens.

Note: If you are configuring a video interface you must perform the configuration in the local policy definition. After the changes are applied to the local policy, the policy can be switched back to Sync With Global mode and the local changes made to the video

interface are retained. See Chapter 49, “Policies overview” for information on policy distribution modes. See Chapter 35, “ISA-Video” for more information about IPTV video features and configuration.

26

Configure the Address parameter.

27

Click on the Channel Config tab and select a channel and click Properties. The Multicast Video Interface Channel Config (Create) form opens.

28

Configure the parameters as required.

29

Save the changes. The configured channels are displayed.

30

Repeat Step 27 to Step 29 to configure additional channels, if required.

31

Save your changes. The Multicast Info Policy (Create|Edit) form reappears.

32

Click Distribute to manually release and distribute the policy locally to devices in the ring. See 49.6 “To release and distribute a policy” (p. 1476) for additional information. Policies are also automatically distributed to devices when they are used by resources on the device.

33

You can view information on VPLS sites (regular sites and I-Sites), VPRN sites, and Default Routing Instances affected by the Multicast Info Policy by clicking the appropriate tab. Run a search and view the properties for the desired item.

34

Click Search to display the newly created policy or policies in the bottom panel of the form.

END OF STEPS

52.14 To configure an ingress multicast reporting destination policy

52.14.1 Steps

1

Choose Policies→Multicast→Ingress Multicast Path Management from the NFM-P main menu. The Ingress Path Management Policies form opens.

2

Click Create→Multicast Reporting Destination or select an existing policy and click Properties. The Multicast Reporting Destination, Global Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click Apply to save the changes. The form refreshes to display additional tabs and buttons.

5

Click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for additional information. Policies are also automatically distributed to devices when they are used by resources on the device.

6

Close the forms.

END OF STEPS

52.15 To configure an SSM translate policy

Note: IGMP must be enabled on the routers to which you intend to distribute the SSM translate policy.

52.15.1 Steps

1

Choose Policies→Multicast→SSM Translate from the NFM-P main menu. The SSM Translate Policies form opens.

2

Click Create or select an existing policy and click Properties. The SSM Translate Policy (Create|Edit) form opens.

3

Configure the parameters as required.

4

Click on the SSM Translate Items tab and click Create. The SSM Translate Policy Item (Create) form opens.

5

Configure the parameters as required and save your changes. The SSM Translate Policy (Create|Edit) form reappears.

Note: The SSM Translate Items added to a policy should not have overlapping IPv4 address ranges, otherwise a configuration error message will be issued.

6

Click Apply to save your changes. The form refreshes to display additional tabs and buttons.

7

Click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for additional information. Policies are also automatically distributed to devices when they are used by resources on the device.

Note: When you distribute an SSM Translate policy to a router, it replaces any SSM Translate policy that may exist on that router. All existing SSM Translate Items on the router will be overwritten.

8

Close the SSM Translate Policy (Create|Edit) form. The SSM Translate Policies form reappears.

9

Click Search to display the newly created policy or policies in the bottom panel of the form.

END OF STEPS

52.16 To configure a 7210 multipoint bandwidth management policy

52.16.1 Steps

1

Choose Policies→Multicast→7210 Multipoint Bandwidth Management Policies from the NFM-P main menu. The 7210 Multipoint Bandwidth Management Policies form opens.

2

Click Create or choose a policy and click Properties. The 7210 Multipoint Bandwidth Management, Global Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Configure the multipoint queues.

1. Click on the Queues tab. Eight default queues are displayed.
2. Choose a queue and click Properties. The Multipoint Bandwidth Queue (Create) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

5

Save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Assign the policy to a 7210 SAS-X.

1. On the network navigation tree, right-click on a 7210 SAS-X and choose Properties.
2. Click on the Globals tab, then the Multipoint Management tab.
3. Click Select to assign the multipoint bandwidth management policy to the 7210 SAS-X. The Select Bandwidth Policy, Network Element Form opens.
4. Select a policy and click OK.
5. Save your changes and close the form.

END OF STEPS

52.17 To configure a Multicast CAC interface policy for IPv4/IPv6

52.17.1 Steps

1

Choose Policies→Multicast→Multicast CAC Interface from the NFM-P main menu. The Multicast CAC Interface Policies form opens.

Note: Multicast CAC Interface policies are not supported on VSR or 7705 SAR-Hm chassis types.

2

Click Create or choose a policy and click Properties. The Multicast CAC Interface, Global Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Save the policy and close the form.

5

Click Distribute to manually release and distribute the policy locally to devices. See 49.6 “To release and distribute a policy” (p. 1476) for additional information. Policies are also automatically distributed to devices when they are used by resources on the device.

6

As required, create the supported object type for this policy type and associate the policy to the object by clicking the Select button. Save your changes and close the form.

Supported object type	Base configuration procedure to create the supported object	Policy location on each object type
VPLS SAP	77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332)	Click on the IGMP Snooping tab, then the General sub-tab.
VPLS Spoke SDP MVPL Spoke SDP	77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386)	Click on the IGMP Snooping tab, then the General sub-tab.
IGMP Interface on a base router and VPRN service	78.8 “To add an IGMP interface to an IES” (p. 2433)	Click on the Multicast CAC tab, then the General sub-tab.
MLD Interface on a base router and VPRN service	28.123 “To configure an MLD interface on a base routing instance or VPRN routing instance” (p. 1035)	Click on the Multicast CAC tab, then the General sub-tab.
PIM Interface on a base router and VPRN service	28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013)	Click on the Multicast CAC tab, then the General sub-tab.

IGMP Group Interface on IES and VPRN services	79.42 “To add an IGMP group interface to a VPRN” (p. 2601)	For VPRNs, click on the Multicast CAC tab, then the General sub-tab. For IES service sites, click on the Multicast tab, then the IGMP Group Interfaces sub-tab or MLD Group Interface sub-tab.
MLD Group Interface on IES and VPRN services	28.125 “To configure an MLD group interface on a base routing instance or VPRN routing instance” (p. 1039)	Click on the Multicast CAC tab, then the General sub-tab.
MSAP policy	64.10 “To configure an MSAP policy” (p. 1852)	Click on the VPLS Configuration tab, then on the ICMP Snooping sub-tab.

END OF STEPS

52.18 To view NE multicast reporting destination statistics

52.18.1 Purpose

Perform this procedure to view multicast reporting destination statistics for an NE including frames and records lost and sent.

52.18.2 Steps

1

Choose Policies→Multicast→Ingress Multicast Path Management from the NFM-P main menu. The Ingress Path Management Policies form opens.

2

Choose Local from the Policy scope drop-down menu.

3

Click Select to choose an NE. The Select a Network Element form opens.

4

Choose an NE and click OK. The NE IP address is displayed in the Local Node IP Address field.

5

Choose Multicast Reporting Destination (Multicast) from the object drop-down menu and click Search.

6

Choose a policy and click Properties. The Multicast Reporting Destination, Local Policy (Edit) form opens.

7

Click on the Statistics tab and click Collect to collect on-demand statistics.

8

Choose a statistics record and click Properties to view the record. The Statistics Record - Multicast Reporting Destination Stats form opens.

9

View the statistics information.

10

Close the forms.

END OF STEPS

52.19 To view multicast CAC channel statistics

52.19.1 Purpose

You can view multicast CAC channel statistics to determine:

- the status of BTV channels on a managed NE
  - which channels are blocked
  - the channel location
  - the channel bandwidth availability
- the number of times channel requests are dropped on an interface
- the cause of the action based on the multicast CAC policy

52.19.2 Steps

1

Choose Tools→BTV Channel Monitor from the NFM-P main menu. The Select Site form opens.

2

Select a site and click OK. The BTV Channel Monitor form opens.

3

Choose one of the following from the object drop-down menu and click Search.

a. To list network protocol statistics, choose Multicast CAC Channel Statistics (multicast).

b. To list service statistics, choose Multicast CAC Channel Statistics on Services (multicast).

A list of multicast CAC channel statistics is displayed at the bottom of the form.

Note: To get the most recent channel information on a managed NE, click Resync before clicking Search.

4

Close the form.

END OF STEPS

53 Time-of-day policies

53.1 Time-of-day policies

53.1.1 Overview

The NFM-P supports time-of-day policies that define the times when policies, filters and schedulers are in effect. You can specify a time range and apply that range to QoS policies, ACL filters, and schedulers that are applied to aggregation schedulers and L2 and L3 access interfaces. For example, a time-of-day policy allows you to manage peer-to-peer traffic by limiting access during peak hours, such as evenings and weekends.

A time range policy consists of one or more schedules that specify a start and end day and time. You can configure ongoing time ranges that recur at a specified frequency.

The default time range policy is the NO-TIME-RANGE policy, which specifies that no time and day restrictions apply. There can be only one NO-TIME-RANGE entry in each type of time-of-day suite policy entry.

A time-of-day suite policy is a collection of ingress and egress policies, filter policies, and schedulers, to which time range policies have been assigned. You can apply time-of-day suite policies to aggregation schedulers and L2 and L3 access interfaces.

You can create a time-of-day suite policy by assigning time range policies to the following objects to apply time and day restrictions to their deployment and access:

- ACL MAC filters and filter entries
- ACL IP filters and filter entries
- ACL IPv6 filters and filter entries
- access ingress and egress QoS policies
- scheduler policies

Time-of-day policies are first created globally and then distributed to NEs.

53.2 Time range assignment analysis tool

53.2.1 General Information

The time range assignment analysis tool allows you to view the time range policies that have been assigned over a specific time period to the following objects:

- L2 access interfaces
- L3 access interfaces
- aggregation schedulers
- time-of-day suite policies

You can use the time range assignment analysis tool to verify the configuration of multiple schedules and the adequacy of a time-of-day suite policy on a specific NE. For example, you can use this tool to verify that multiple time ranges that are assigned to IP filters in a time-of-day suite policy cover an entire month.

See 53.5 “To perform a time range entry assignment analysis” (p. 1735) for information about configuring time range assignment analysis.

53.3 To configure a time range policy

53.3.1 Steps

1

Choose Policies→Time of Day→Time Range from the NFM-P main menu. The Time Range Policies form opens.

2

Click Create, or choose a policy in the list. The Time Range (Create|Edit) form opens.

3

Configure the required general parameters.

4

Configure one or more time ranges.

1. Click on the Time Range Entries tab.
2. Click Create, or choose an entry in the list and click Properties. The Time Range Entries (Create) form opens.
3. Configure the Time Settings parameters.

   When the Ongoing parameter is enabled, the Frequency parameter is available.

   When the Frequency parameter is set to Weekly, the Start Run Day and End Run Day parameters are available.
4. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

53.4 To configure a time-of-day suite policy

53.4.1 General Information

Consider the following when you create a time-of-day suite policy:

- when there are overlapping time range policy entries within a time-of-day suite policy, the time range policy entry with the highest priority is applied first
- only one time-of-day suite policy entry without a time range (default NO-TIME-RANGE) is allowed for each type of policy
- you cannot modify a time-of-day suite policy entry after you create the entry
- you cannot assign the same priority of a global policy within the same policy type
- you cannot assign the same priority of a local policy within the same policy type
- you cannot assign the same time range of a global policy within the same policy type
- you cannot assign the same time range a local policy within the same policy type
- you cannot distribute a time-of-day suite policy entry to a local policy if the local policy already has been assigned the same time range as the global policy
- you cannot distribute a time-of-day suite policy entry to a local policy if the local policy already has been assigned the same priority as the global policy
- all time-of-day suite policy entries of a local policy are replaced by the time-of-day suite policy entries of a global policy if none of the entries from the local policy have the same priority or time ranges as the global policy

53.4.2 Steps

1

Choose Policies→Time of Day→Time Of Day Suite from the NFM-P main menu. The Time Of Day Suite Policies form opens.

2

Click Create, or choose a policy and click Properties. The Time of Day Suite (Create) form opens.

3

Configure the required general parameters.

4

Click on the TOD Suite Entries tab.

5

Configure time-of-day suite entries for ingress traffic. Click on the Ingress tab and perform any of the following:

a. Associate time ranges with ingress filter policies.
   1. Click on the IP Filter, IPV6 Filter or MAC Filter tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Ingress *type_of_filter* (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy ID panel and choose a filter policy.
   6. Save your changes and close the form.

b. Associate time ranges with ingress QoS policies.
   1. Click on the QoS tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Ingress QoS (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy panel, or 7210 Policy panel, and choose a QoS policy.
   6. Save your changes and close the form.

c. Associate time ranges with ingress scheduler policies.
   1. Click on the Schedulers tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Ingress Scheduler (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy Name panel and choose a scheduler policy.
   6. Save your changes and close the form.

6

Configure time-of-day suite entries for egress traffic. Click on the Egress tab and perform any of the following:

a. Associate time ranges with egress filter policies.
   1. Click on the IP Filter, IPV6 Filter or MAC Filter tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Egress *type_of_filter* (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy ID panel and choose a filter policy.
   6. Save your changes and close the form.

b. Associate time ranges with egress QoS policies.
   1. Click on the QoS tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Egress QoS (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy ID panel and choose a QoS policy.
   6. Save your changes and close the form.

c. Associate time ranges with egress scheduler policies.
   1. Click on the Schedulers tab. A list of available time range policies is displayed.
   2. Click Create, or choose an entry in the list and click Properties. The Egress Scheduler (Create) form opens.
   3. Configure the Priority parameter.
   4. Click Select in the Time range panel and choose a time range policy.
   5. Click Select in the Policy Name panel and choose a scheduler policy.
   6. Save your changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

53.5 To perform a time range entry assignment analysis

53.5.1 Steps

1

Choose Tools→Time Range Entry Assignments from the NFM-P main menu. The Time Range Entry Assignments form opens.

2

Click Select and choose a site for the analysis.

3

Configure the Time Range Entry Container Type parameter, then click Select to choose a specific object of the configured container type.

4

If required, configure the Search by Time Of Day Entry Type and Time Of Day Entry Policy Type parameters.

5

Configure the Start Date and End Date parameters.

6

Click Search. A list of time range entries that meet the defined criteria is displayed.

7

Double-click on an entry to view information about the time range policy and time-of-day suite policy. The Time Range Entry Assignment: *Start Date End Date* form opens.

8

Close the forms.

END OF STEPS

54 Routing policies

54.1 Routing policies

54.1.1 Overview

Routing policies, also known as route redistribution policies, are used to make routing decisions based on policies that override the default routing protocol decisions. Normally, when a router receives a packet, where it is forwarded to is based on the destination address in the packet, which is then used to look up an entry in a routing table.

You may have a requirement to forward a packet based on other criteria to reach a destination. For example, you can selectively apply routing policies based on prefix and community access lists, packet size, advertised routes, the size and contents of the routing table, and other user-defined criteria. Additional user actions can also be applied to define user-defined routes, and to set the precedence, the type of service bits, and the damping method for the routes.

There are no default routing policies on the NFM-P. Each policy must be created and applied to an object, a routing protocol, or the forwarding table. Each set of rules that is associated with controlling routes are called routing policy statement entries on the client GUI. For example, use routing policies to:

- control routing protocols, such as BGP, to allow routes from another protocol, such as OSPF, into the routing table, which allows the routing table to redistribute packets into other protocols
- control the import and export of learned active routes of a protocol
- modify the characteristics of a route, for example, to change the community values of BGP AS path attributes
- prevent the routes for specific customers from being added to routing tables
- control BGP route flapping

54.2 Supported NE routing policies

54.2.1 General Information

The following table describes the NE routing policies you can configure using the NFM-P. The table below describes the additional functions you can perform with routing policies.

Table 54-1 NFM-P supported NE routing policies

Routing policy type	Purpose	See
Statement	This policy defines the flow control actions taken if a route matches all the conditions specified by the policy. Flow control actions determine whether to accept or reject the route or whether to evaluate the next term or routing policy. In addition, you can configure a multicast redirection interface as part of the routing policy statement to ensure that multicast IGMP traffic coming from the access network is only passed onto downstream devices once.	54.5 “To configure a routing policy statement” (p. 1745) 54.6 “To configure a multicast redirect interface on a local routing policy statement” (p. 1750)
Prefix List	This policy defines a list of IP prefixes members used by the Routing policy statement that specifies the matching criteria used for incoming routes. A prefix list is a named list of IP addresses. You can specify an exact match with incoming routes and apply a common action to all matching prefixes in the list.	54.7 “To configure a prefix list policy” (p. 1752)
Community	This policy defines the community members used for routing policy entries and provide a means of grouping destinations to which routing decisions can be applied, for example, to change the community values of BGP AS path attributes.	54.8 “To configure a community policy” (p. 1752)
Damping	This policy defines the damping mechanism used to control BGP route flapping which reduces the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.	54.9 “To configure a damping policy” (p. 1755)
AS Path	This policy defines the AS path to a destination and is used for both route selection and to prevent routing loops. An AS path consists of the AS numbers of the networks that a packet traverses if it takes the associated route to a destination.	54.10 “To configure an AS Path policy” (p. 1756)
AS Path group	This policy defines the match conditions based on an AS path group which consists of multiple AS paths to which a single routing policy can be applied.	54.11 “To configure an AS Path Group policy” (p. 1756)
Accounting Template Policy	This policy defines if accounting statistics are collected on IP interfaces for received packets and bytes that are sourced from or destined for routes marked with the class indexes contained in the policy. An accounting template policy can be associated with the following interfaces on base routers, IES services, and VPRN services: • network interfaces • L3 access interfaces • tunnel interfaces • subscriber group interfaces	54.12 “To configure an accounting template policy” (p. 1757)
Administrative Group	This policy defines the administrative groups that can be assigned to MPLS interfaces, LSPs, LSP paths, L3 interfaces for VPRN and IES services, or flexible algorithm definitions. You can also assign Administrative groups to network interfaces at the base routing instance. See 54.2.2 “Administrative Group policies” (p. 1740) for more information.	54.13 “To configure an administrative group policy” (p. 1758)

Table 54-1 NFM-P supported NE routing policies (continued)

Routing policy type	Purpose	See
Shared Risk Link Group	This policy defines which links in a network share a common fiber which, in the event one link fails, other links in the group may fail therefore they have a shared risk. See 54.2.3 “Shared Risk Link Group policies” (p. 1741) for more information. In addition, you can create a static configuration for a SRLG Policy to manually enter the SRLG membership for links in the entire network into the SRLG database.	54.14 “To configure a Shared Risk Link Group policy” (p. 1759) 54.15 “To create a static configuration for a SRLG Policy” (p. 1761)
Route Next Hop Template	This policy defines the policy control for the Loop-Free Alternate Shortest Path First (LSA SPF) feature which is used to reduce the routing transition time in case of link/router failures. A Route Next Hop Template policy can be applied to OSPFv2, OSPFv3 and IS-IS interfaces on base routing instances and VPRN services.	54.16 “To configure a route next hop template policy” (p. 1762)
Re-Assembly Profile	This policy defines the timers that assure all expected fragments of a packet are received within an expected time frame, on a per forwarding class basis. When the specified timer expires, all of the fragments of an incomplete frame are dropped. The system visits the reassembly queues every 64 ms in a constant loop, which causes up to a 63 ms variation between the user configured value and the actual detection time. For example, with the default configuration of 2000 ms, the system visits the reassembly queue timer at 1999 ms in which case there will not be a timeout. The timeout will occur during the next visitation at 2063 ms.	54.17 “To configure a re-assembly profile policy” (p. 1764)
Maintenance Policy	This policy provides BFD template information to support segment routing seamless BFD sessions. The maintenance policy must be associated with a segment routing static policy for a seamless BFD session to be created; see 28.77 “To create a segment routing policy” (p. 976).	54.18 “To configure a maintenance policy” (p. 1764)
Route Distinguisher	This policy defines the route distinguisher (RD). A route distinguisher is an 8-byte value that is prepended to an IPv4 or IPv6 prefix to ensure that the resulting BGP NLRI is unique within the routing domain that propagates BGP routes for this prefix. An RD set is a list of one or more route-distinguisher entries. Each route-distinguisher entry matches one or more route-distinguisher values depending on the use of wildcards.	54.25 “To configure a route distinguisher policy” (p. 1771)

Table 54-2 Additional routing policy functions

Policy function	Purpose	See
View policy variable usage in a routing policy statement	Allows you to view policy variable usage in a routing policy statement. Policy variables are defined in a (main) routing policy statement and are essentially pointers to certain other routing policy types. They are subsequently called by a subordinate routing policy statement that is specified in the main policy.	54.20 “To view policy variable usage in a routing policy statement” (p. 1766)
View routing policy usage	Allows you to display which NEs are currently using a global or local version of certain routing policy type.	54.21 “To view routing policy usage” (p. 1767)

Table 54-2 Additional routing policy functions (continued)

Policy function	Purpose	See
Shows the routing policy CLI configuration in the client GUI	Allows you to display the output of the **show router policy** *policy_type* CLI command on the NFM-P GUI.	54.22 “To show a routing policy CLI configuration in the client GUI” (p. 1768)
Verify BGP routes against a routing policy statement	Allows you to verify BGP routes against a routing policy statement to identify the prefixes that are matched or not matched by the policy, before attaching the prefix to a routing neighbor or instance.	54.23 “To verify BGP routes against a routing policy statement” (p. 1769)
Verify a BGP prefix list against a routing policy statement	Allows you to verify a BGP prefix list against a routing policy statement to identify the accepted and rejected route entries.	54.24 “To verify a BGP prefix list against a routing policy statement” (p. 1770)
Configure a global policy variable	Global variables function similarly to policy variables, but allow for more efficient configuration, because they can be configured once, then selected in multiple routing policy statements. This eliminates having to manually configure a policy variable in every routing statement. Support for global variables depends on the NE and release; see the NE documentation for information.	54.19 “To configure global variables” (p. 1765)

54.2.2 Administrative Group policies

Administrative group policies define administrative groups that can be assigned to MPLS interfaces, LSPs, LSP paths, and L3 interfaces in VPRN and IES services. You can also assign Administrative groups to network interfaces at the base routing instance, and to flexible algorithm definitions in NE global properties. After you configure the administrative groups, they can be assigned to interfaces and paths on their respective properties forms. Multiple administrative groups can be assigned to each of these objects.

When establishing LSP and LSP paths, devices only consider MPLS interfaces which are associated with the same administrative group as the LSP or LSP path. MPLS interfaces advertise administrative group associations using CSPF. This is done using the 32 bit mask which you configure using the Value parameter on the administrative group policy form.

An administrative group can also be assigned to be explicitly excluded from LSPs and LSP paths. The device cannot use MPLS interfaces in the administrative group to establish LSPs or LSP paths. Administrative group exclusion takes priority over administrative group inclusion.

CSPF must be enabled on LSPs for administrative groups to be relevant. You can enable and configure CSPF on the LSP properties form. When CSPF is enabled on an LSP, it is automatically enabled on associated LSP paths. LSP paths can be configured on the LSP path properties form to inherit additional CSPF, administrative group, and other parameters from LSPs.

The following table describes where to find information about assigning administrative groups.

Table 54-3 Administrative group assignments

To assign groups to	See Procedure
MPLS interfaces	31.7 “To create an MPLS interface” (p. 1120) in Chapter 31, “MPLS”
LSPs	31.10 “To create a static LSP” (p. 1124) in Chapter 31, “MPLS”
LSP paths	31.22 “To configure an LSP path” (p. 1144) in Chapter 31, “MPLS”
L3 interfaces	27.17 “To create an L3 network interface on a routing instance” (p. 856) in Chapter 27, “NE routing and forwarding”
Flexible algorithm definitions	12.5 “To modify NE properties” (p. 343) in Chapter 12, “Device object configuration”

54.2.3 Shared Risk Link Group policies

SRLG policies defines which links in a network share a common fiber which, in the event one link fails, other links in the group may fail therefore they have a shared risk. SRLGs are constructs which allow you to perform two operations that enhance overall system reliability.

SRLGs policies are associated with MPLS interfaces; an MPLS interface can belong to multiple SRLGs. SRLGs can also assigned to network interface at base routing instance and to L3 access interface of VPRN & IES service.

You can use SRLGs to establish a FRR LSP path and you also use SRLGs to establish a secondary LSP path which is disjointed from the primary LSP path. SRLG avoids the need to define MPLS path hops for secondary LSPs and creation of static bypass tunnels and management of these LSPs to achieve the same result dynamically.

Links that are members of the same SRLG represent resources which share the same risk. For example, fiber links sharing the same conduit, or multiple wavelengths sharing the same fiber.

An SRLG is modeled as a policy object. It therefore follows the normal policy behavior for creation, listing, updating, deletion, distribution, and re synchronization.

The SRLGs are used by the CSPF when computing a FRR detour/bypass path, or a secondary LSP path. SRLGs indicate to the CSPF which interfaces to avoid in the path’s computation. CSPF can include the SRLG penalty weight in the computation of a FRR detour or bypass for the protection of the primary LSP path at a PLR NE. Configured SRLGs can be associated with MPLS interfaces, and network, tunnel, and L3 access interfaces on base routers, VPRN services, and IES services.

The following table describes where to find information about related MPLS and LSP procedures.

Table 54-4 Related MPLS and LSP procedures

To:	See Procedure:
Configure an MPLS instance	31.6 “To configure an MPLS instance” (p. 1116) in Chapter 31, “MPLS”
Create an MPLS interface	31.7 “To create an MPLS interface” (p. 1120) in Chapter 31, “MPLS”
Configure LSP paths	31.22 “To configure an LSP path” (p. 1144) in Chapter 31, “MPLS”
Configure L3 interfaces	27.17 “To create an L3 network interface on a routing instance” (p. 856) in Chapter 27, “NE routing and forwarding”

54.3 Routing policy design considerations

54.3.1 General information

Careful planning is essential to implementing routing policies that affect the flow of routing information or packets traversing managed devices. Before configuring and applying a routing policy, consider the following:

- Develop an overall plan and strategy to accomplish your intended routing actions.
- Analyze the effect of what happens to a packet that meets the specified criteria in a routing policy statement entry to ensure the proper action is executed and that no routing loops are created.
- Ensure that routing policy statement entries are properly numbered, so they will be compared to packets in the correct order. Entries are compared in the numerical sequence of the Entry IDs, from lowest to highest. Nokia recommends staggering the numerical Entry ID values (for example, using 10, 20, 30… instead of 1, 2, 3…), to allow insertion of additional entries. The NFM-P supports renumbering existing entries on supporting NEs.
- When possible, redistribute from a lower-level routing protocol to a higher-level routing protocol; for example, from RIP to OSPF.
- Redistribute exported routes at a single device, when possible.
- To ensure routing configuration flexibility, the following routing policies can be configured as separate policies or can be cross-referenced to form a framework of policies depending on the network requirements:
  - Policy Statement
  - Prefix List
  - Community
  - Damping
  - AS Path
- Policies that are cross-referenced are distributed together. For example, a Policy Statement can reference the Prefix List policy by matching the From Criteria and the To Criteria for each policy type. In addition, policy variables can be defined in an upper-level routing Policy Statement, and then called by a specified subordinate routing Policy Statement. The variables are pointers to certain types of routing policies, including Community Lists, AS Paths, AS Path Groups, and Prefix Lists. To improve configuration efficiency, you can create global policy variables (for supporting NEs) and assign them to multiple routing policy statements. This routing policy parameterization allows operators a powerful and flexible approach to the configuration of routing policies.
- When IGMP join requests originate from subscriber host sessions, the NE can use a multicast redirection policy to redirect the requests to a plain (non-ESM) L3 access interface. Additionally, individual local routing policy statements can be configured with a multicast redirect action to redirect specific IGMP join requests for processing at the subscriber level.
- Imported BGF, OSPF, and RIP learned routes can be altered using import and export route policies. For example, for devices learning about routes from BGP or OSPF, you can create an import route policy that can limit the number and types of routes accepted and added to the routing tables. For exporting routes, some default behaviors exist:
  - internally learned routes are redistributed using the same protocol

  - externally learned routes are not automatically advertised to all neighbors or peers
- In the case of externally learned routes, you can create an export policy to determine how the routes are advertised; for example, configuring something learned by BGP to be redistributed using OSPF. This is useful in cases:
  - where legacy equipment does not support a specific protocol
  - when companies merge and they use different routing protocols on the devices

54.4 Routing policy decision sequence

54.4.1 General information

Routing policy statements and routing policy statement entries are compared against incoming packets in the following sequence.

1. The routing packets arrive.
2. The first routing policy is analyzed.
3. The first entry in the routing policy statement is analyzed. Entries are compared in sequence based on the Entry ID, from lowest to highest.

   If a match is found based on the first routing policy statement entry, the specified action is performed. You can use the Action parameter to specify the continued evaluation of route policy entries, or additional policy statements.
4. If no match is found, the packet is compared against the second routing policy statement entry. If a match is found based on the second routing policy statement entry, the specified action is performed.
5. If no match is found, the packet is compared sequentially against all remaining routing policy statements and routing policy statement entries. If a match is found, the specified action is performed.
6. If the packet does not match any routing policy statements or entries, the default action is performed.

   The following figure shows how routes are analyzed using routing policies.

Figure 54-1 Routing policy analysis workflow

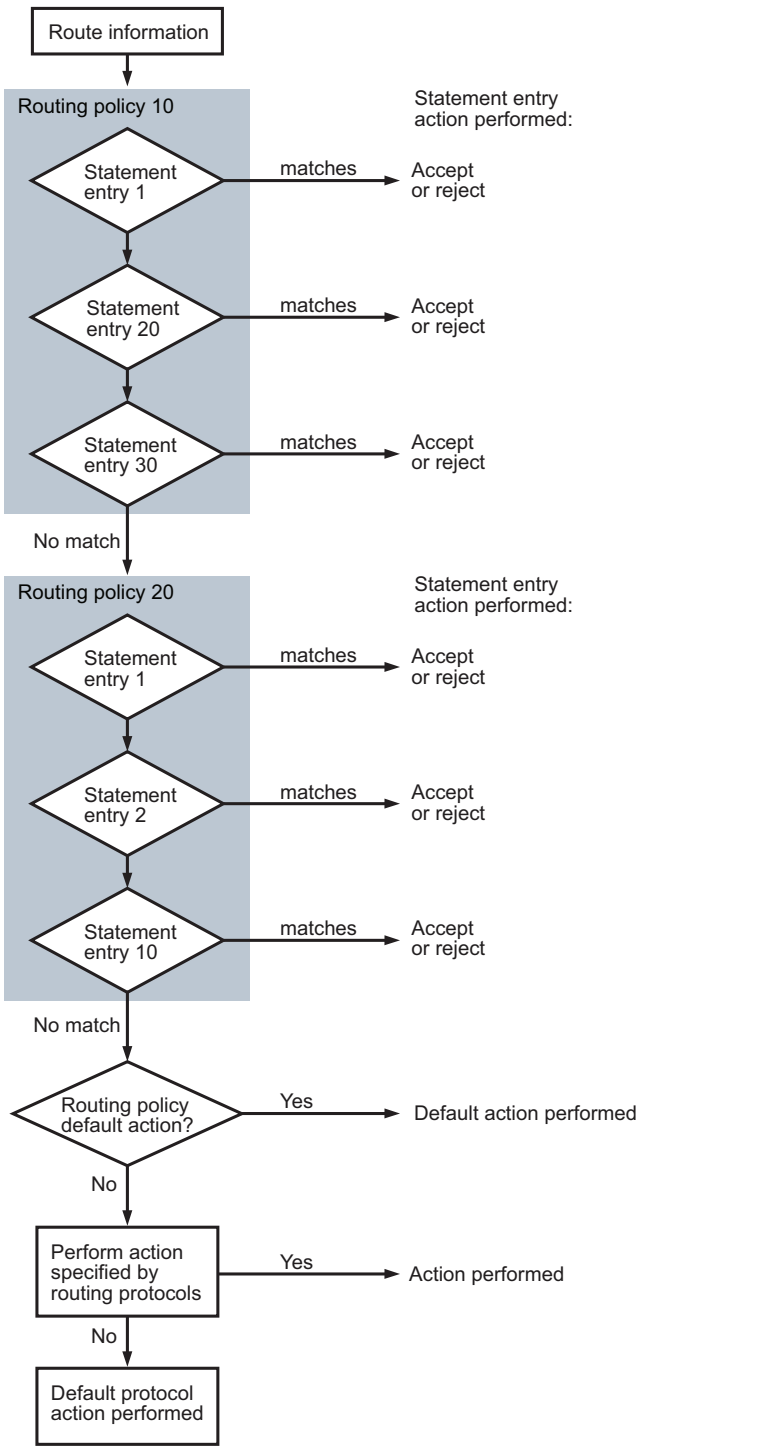

54.4.2 No action/default action routing policy decision sequence

When no action or default action in a routing policy is matched on a packet, then the routing protocols determine the default action for the packet. The action can be specified for the protocol. If the action is not specified, there are default routing protocol import and export actions performed, as described in the following table.

Table 54-5 Default routing protocol actions

Protocol	Default import route action	Default export route options
BGP	All routes from BGP peers are accepted and forwarded to the BGP route selection process.	All active internal BGP routes are advertised to BGP peers. All non-BGP learned routes are not advertised to BGP peers.
IS-IS	IS-IS route acceptance cannot be configured. All IS-IS routes are accepted from IS-IS neighbors.	All internal IS-IS routes are advertised to all IS-IS neighbors. All non-IS-IS learned routes are not advertised to IS-IS neighbors.
OSPF	OSPF route acceptance cannot be configured. All OSPF routes are accepted from OSPF neighbors.	All OSPF routes are advertised to all OSPF neighbors. All non-OSPF learned routes are not advertised to OSPF neighbors.
RIP	All RIP-learned routes are accepted from RIP peers.	All non-RIP learned routes are not advertised to RIP peers.

54.5 To configure a routing policy statement

54.5.1 Steps

1

Choose Policies→Routing→Statement from the NFM-P main menu. The Routing Policy - Statements form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy - Statement, Global Policy form opens.

3

Configure the required parameters on the General tab.

Note: Do not create a new routing policy statement with the following terms as the Policy Statement Name: N/A, n/a, or @. Creating a routing policy statement with @ in the Policy Statement Name results in Template Inconsistent/Configuration Mismatch alarms, which

do not clear after a synchronization with the global policy. Also, updating the policy statement name using node CLI as N/A is not valid.

If you set the Default Action parameter to Accept, Next Entry, Next Policy, or Reject/Drop, the Default Action tab is selectable. Go to Step 4.

If you set the Default Action parameter to None, go to Step 6.

4

Configure a default action for the routing policy statement. Click on the Default Action tab and perform the following on the General tab:

1. Configure the parameters, policy variables, and global variables on the General sub-tab, as required.

   Note:

   If you want to specify a policy variable for an action item on this tab, select the policy statement entry that defines the required policy variable.

   If you want to specify a policy variable for an action item on this tab, then you cannot configure a corresponding parameter Value for the same item. They are mutually exclusive.

   If you want to specify a global variable for an action item on this tab, select the global variable entry that defines the required global variable. Global variables are configured in 54.19 “To configure global variables” (p. 1765).

2. Click on the Path sub-tab and configure the BGP AS path parameters and variables, as required.

   Note:

   A BGP AS Path Name must be selected if the BGP AS Path Action is set to Add or Replace.

   You cannot specify both a BGP AS Prepend Number and a BGP AS Prepend Policy Variable. They are mutually exclusive.

   If you are specifying a BGP AS Prepend Policy Variable, this is done by selecting the policy statement entry that defines the required variable.

   If you want to specify a global variable for an action item on this tab, select the global variable entry that defines the required global variable. Global variables are configured in 54.19 “To configure global variables” (p. 1765).

5

Configure the default community actions for the routing policy statement.

1. Click on the Default Community Actions tab and click on the Add, Remove, or Replace sub-tabs as required to perform the specific action that you want to configure.
2. Click Select to add, remove, or replace a Community Name Action. Review the usage notes provided below the Community Name Action field and implement, as required.

Note: You can specify mid-string expansions such as "peer-@asname@" or "comm-@peeras@-add" for example, for the Community Name Action field. However, there is no validation on the parameterized mid-strings entered for communities.

6

Configure policy statement entries for the routing policy statement.

Note: Routing policy entries are compared against incoming packets. Entries are compared in the numerical sequence of the Entry IDs, from lowest to highest. When a match is found, the action specified using the Action parameter occurs. If no action is specified, the action specified using the Default Action parameter occurs. If no default action is specified, the default action for the protocol or the route configuration occurs. For this reason, Nokia recommends that you sequence the routing policy entries to ensure the first entry is the most explicit and the last entry is the least explicit.
Nokia recommends staggering the numerical Entry ID values (for example, using 10, 20, 30… instead of 1, 2, 3…), to allow insertion of additional entries without renumbering the entire set. The NFM-P supports renumbering existing entries on supporting NEs.

1. Click on the Policy Statement Entries tab.
2. To renumber an entry on an existing policy statement, choose the entry in the list and click Renumber ID. Repeat for each additional entry that requires renumbering.

   When the routing policy statement is redistributed to NEs that support entry renumbering, the new Entry ID values are used.
3. Click Create, or choose a policy in the list and click Properties. The Policy Statement Entry, Site form opens.
4. Configure the required parameters on the General tab.

   Select a policy variable for the SRv6 Sid Prefix Address parameter, configure the SRv6 Sid Prefix Length parameter, and select a global variable for the Policy/Global Variable parameter.

   You can choose to enable or disable either or both the To Criteria and From Criteria here, whether they are configured or not.

   If you set the Action parameter to Accept, Next Entry, or Next Policy, the Action tab is selectable. Go to Step 7 .

   If you set the Action parameter to None or Reject/Drop, go to Step 8 .

7

Click on the Action tab to configure actions for the policy statement entry.

1. Configure the parameters and select policy variables on the General sub-tab, as required.

   Notes:

   You must configure the SR Maintenance Policy parameter in order to select the policy variables in the SRv6 Return Path BFD and SR Return Path BFD Label panels.

   If you want to specify a policy variable for a subordinate policy action item on this tab, this is done by selecting the policy statement entry that defines the required variable. Such a policy statement entry is typically part of another existing (higher level) routing policy statement.
2. Click on the Path sub-tab and configure the BGP AS path parameters, as required.

   Note:

A BGP AS Path Name must be selected if the BGP AS Path Action is set to Add or Replace.

You cannot specify both a BGP AS Prepend Number and a BGP AS Prepend Policy Variable. They are mutually exclusive.

If you are specifying a BGP AS Prepend Policy Variable, this is done by selecting the policy statement entry that defines the required variable. Such a policy statement entry is typically part of another existing (higher level) routing policy statement.

3. Click on the Community Actions tab to configure a community action for the policy statement entry.
4. Click on the Add, Remove, or Replace sub-tabs as required to perform the specific action that you want to configure.
5. Click Select to add, remove, or replace a Community Name Action. Review the usage notes provided below the Community Name Action field and implement as required.

   Note:

   You can specify mid-string expansions such as "peer-@asname@" or "comm-@peeras@-add" for example, for the Community Name Action field. However, there is no validation on the parameterized mid-strings entered for communities.

8

Configure the From criteria for the policy statement entry.

1. Click on the From Criteria tab, select a route distinguisher policy, and configure the parameters on the General sub-tab, as required.

   Note:

   If you are performing this procedure to create a subordinate (lower level) routing policy that needs to specify an AS Path Name, AS Path Group Name, or Community List Name as a policy variable (defined in a higher level routing policy), then these policy variables are selected at this step.

   If you are performing this procedure to create a main (higher level) routing policy that contains policy variables that will be called by a subordinate routing policy, then that subordinate routing policy name must be specified at this step using the Policy Statement Name parameter. The policy variables for this higher level routing policy are defined in Step 10 .

   You cannot add policies that have nested sub-policies.

   Configuring a Community List Name in this step, either by selection or as a policy variable, is mutually exclusive with creating a Community List Expression in Step 9 .
2. Click on the Prefix Lists sub-tab to assign a prefix list to the policy statement entry.

   The Prefix Lists sub-tab displays if you did not assign a prefix list on the Prefix List Override sub-tab.
3. Configure the Prefix List Members as required.

   Note:

   You can specify mid-string expansions such as "prefixA-@prefixB@" or "prefixA-@prefixB@-add" for example, in the Prefix List field. However, there is no validation on the parameterized mid-strings entered for prefix lists.

4. Click on the Prefix List Override sub-tab.

   The Prefix List Override sub-tab displays if you did not assign a prefix list on the Prefix Lists sub-tab.

5. Select the prefix list and configure the required parameter.

   Depending on the option that you selected for the Type parameter, additional parameters may display.

9

Create a community expression list for the policy statement entry.

1. Click on the Community Expression tab and click Create. The Community Expression, Site form opens.
2. Select or enter a Community Name.
3. Click the Add button to add an entry into the Community List Expression field. The following conditions apply to this parameter:
   - The entire Community List Expression is limited to 900 characters. Use Boolean operators (AND, OR, and NOT) to form expressions, and use parentheses to form more complex expressions. For example:

     `[comlist1] AND ([comlist3] OR [comlist4])`
   - If you make an error in the construction of the string, use the Clear All button to start again, or place your cursor in the field and manually edit the text.
4. Save your changes and close the form.

10

Create a policy variable for the policy statement entry.

This step is required if you are creating a main (higher level) routing policy that will be called by a subordinate (lower level) routing policy. A policy variable in this context may be an integer, an IPv4/IPv6 address, or a string (that is, a pointer to another routing policy type). If you want to select the Type as a String, then you must create the other routing policy type (for example, Prefix List or Community policy) before you start this step. You can create up to ten policy variables for each policy statement entry.

1. Click on the Policy Variables tab and click Create. The Policy Variable, Site form opens.
2. Configure the Name parameter for the policy variable. The name must begin and end with an “@” sign; for example: @variable-name_1@.
3. Configure the Type and Value parameters. If you configure Type as String, then click Select to choose and assign a routing policy to the policy variable.
4. Click OK. The Policy Variable, Site form reappears.
5. Click OK to save your changes and close the form.

Note: You can configure global policy variables (for supporting NEs) that you can assign to multiple routing policy statements to improve configuration efficiency; see 54.19 “To configure global variables” (p. 1765).

11

Configure the To criteria for the policy statement entry.

1. Click on the To Criteria tab and configure the parameters on the General tab, as required.
2. Click on the Prefix Lists sub-tab and configure the Prefix List Members as required.

12

Configure a conditional expression for the policy statement entry.

1. Click on the Conditional Expression tab and click Create. The Conditional Expression, Site form opens.
2. Specify a conditional expression for the Route Exists String parameter, or click on the Select button to select an existing prefix list policy.
3. Save your changes and close the form.

13

To create an additional policy statement entry, go to Step 6 otherwise go to Step 14 .

14

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with global. To configure the distribution mode, see 49.6 “To release and distribute a policy” (p. 1476) .

15

See 54.20 “To view policy variable usage in a routing policy statement” (p. 1766) to view policy variable usage in a distributed routing policy statement.

END OF STEPS

54.6 To configure a multicast redirect interface on a local routing policy statement

54.6.1 General Information

To perform this procedure, you must use a global routing policy statement that has been released and distributed to one or more NEs. See 54.5 “To configure a routing policy statement” (p. 1745) for more information.

54.6.2 Steps

1

Choose Policies→Routing→Statement from the NFM-P main menu. The Routing Policy - Statements form opens.

2

Choose a policy and click Properties. The Routing Policy - Statement form opens.

3

Click on the Local Definitions tab. The sites the Routing policy statement is distributed to are listed.

4

Choose a site and click Properties. The Routing Policy - Statement, Site Local Policy form opens.

5

Click on the Policy Statement Entries tab.

6

Choose a policy statement and click Properties. The Policy Statement Entry, Routing Policy, Site form opens.

7

Click on the Action tab.

The Action tab is only available if the Action parameter on the Policy Statement was set to Accept, Next Entry, or Next Policy.

8

Click Select to choose an interface in the Multicast Redirection panel. The Select Interface - Policy Statement Action - Site form opens.

9

Select an interface and save your changes. The Policy Statement Entry, Routing Policy, Site form reappears.

10

Save your changes and close the form. The Policy Statement Entry, Routing Policy, Site form reappears. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

54.7 To configure a prefix list policy

54.7.1 Steps

1

Choose Policies→Routing→Prefix List from the NFM-P main menu. The Routing Policy - Prefix Lists form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy Prefix List, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Click on the Prefix List Members tab and click Create. The Prefix List Member, Site form opens.

5

Configure the required parameters.

Note: The Begin Length and Through Length parameters are configurable when the Type parameter value is set to Range or Through.
A prefix list policy can be created without creating prefix list members. These can be added to the policy at a later date.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

End of steps

54.8 To configure a community policy

54.8.1 Steps

1

Choose Policies→Routing→Community from the NFM-P main menu. The Routing Policy - Communities form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy - Community, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Perform one of the following:

Note: You must create at least one Community Expression or Community Member for this policy type.

a. To create a Community Expression, go to Step 5 .

b. To create a Community Member, go to Step 6 .

5

Click on the Community Expression tab and click Create. The Community Expression, Site (Create) form opens.

1. Configure the Exact Match parameter as required.
2. Enter a string in the Community Expression text field. Use Boolean operators to define the policy actions between different values. Valid expression operators include: AND, OR, and NOT. For example, if you enter the following string:

   `target:5:1000 OR target:9:10`

   then the policy matches a route that has either target. Or if you enter:

   `2:2 AND 200:200`

   then the policy only matches a route that has both targets.

   The Boolean operators can also be used with parentheses and combined into a more complex expression, such as:

   `(8:8 AND NOT 9:9) OR (7:7)`

   In this case, the policy matches a route which contains the regular expression for target 8:8 with the exception of target 9:9, or just target 7:7 alone.

Any community, extended community, large community, or well-known community can be part of the Community Expression. Valid values for the Community Expression's components include:

- Community Values, such as 100:100
- Extended Community Values, such as target:1:1, origin 1:1, target:1.1.1.1:1, or origin 1.1.1.1:1. On supporting NEs, flowspec-set:<ext-asnum>:<group-id>, and flowspec-set-trans:<ext-asnum>:<group-id> are configurable; see the NE documentation.
- Regular expressions for origins, such as origin:.*&.*
- Regular expressions for communities, such as 100|200:100|200

- Regular expressions for extended community values, for example, Extended Community Bandwidth:<as-number>:<value in Mbps> to specify link bandwidth
- Large Community Values (on supporting NEs; see the NE documentation)
- Well-known Community values

Note:

When constructing a Community Expression, Community IDs must be limited to 75 characters maximum, and regular expressions must be limited to 72 characters maximum. The entire Community Expression must be limited to 900 characters maximum.

Community Expressions are supported on the 7750 SR, 7750 SR-c12, 7750 SR-12E, and 7450 ESS. IOM3 is required, in chassis mode B, C, or D.

3. Save your changes. The Routing Policy - Communities form reappears. Go to Step 7 .

Note: 7x50 NEs support non-matching lists for regular expressions.

6

Click on the Community Members tab and click Create. The Community Member, Site (Create) form opens.

1. Configure the Community Member parameter.

For extended communities, use the generic format of ext:<value1>:<value2> or ext:<regex1>&<regex2>, where:

- ext is a keyword similar to target and origin
- <value1> if 4 hex digits, then <value2> can have 1-12 hex digits
- <value1> if 2 hex digits, then <value2> can have 1-14 hex digits
- <regex1> is a regular expression that can match a pattern of 4 hex digits representing the first 2 bytes of the extended community
- <regex2> is a regular expression that can match a pattern of 1-12 hex digits representing the last 6 bytes of the extended community

Use of the <regex> attributes can provide flexibility for removing multiple extended communities from a BGP route.

Note:

If you are configuring an origin validation state extended community for BGP SIDR, configure ext:xx[yy]:value, where ext indicates an extended community and value indicates the new origin validation state, from 0 to 299.

The ext format has following restrictions:

- <value1> must have the value 4300, specifying a non-transitive opaque extended community
- <value2> can have values 0, 1, or 2, corresponding to the three possible origin valid states

If you are configuring a link bandwidth extended community, then configure the Community Member as: bandwidth:<as-number>:<value in Mbps>

If you are creating a flowspec-set or flowspec-set-trans community member that will be used in an embedded filter, set the group-ID value to the same value configured for the

Group ID parameter in 51.9 “To configure an embedding filter with embedded filter policies” (p. 1686). The format for flowspec-set and flowspec-set-trans values is <AS-number>:<group-ID>.

2. Save your changes. The Routing Policy - Communities form reappears. Go to Step 7 .

7

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

End of steps

54.9 To configure a damping policy

54.9.1 Steps

1

Choose Policies→Routing→Damping from the NFM-P main menu. A Routing Policy - Dampings form opens.

2

Click Create or choose a policy and click Properties. A Routing Policy - Damping, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

End of steps

54.10 To configure an AS Path policy

54.10.1 Steps

1

Choose Policies→Routing→AS Path from the NFM-P main menu. A Routing Policy - AS Paths form opens.

2

Click Create or choose a policy and click Properties. A Routing Policy - AS Path, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.11 To configure an AS Path Group policy

54.11.1 Steps

1

Choose Policies→Routing→AS Path Group from the NFM-P main menu. The Routing Policy - AS Path Groups form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy - AS Path Group, Global Policy form opens.

3

Configure the required parameters on the General tab.

Note: You must create at least one AS Path Group Entry for this policy type.

4

Click on the AS Path Group Entries tab and click Create. The AS Path Group Entry, Site form opens.

1. Configure the required parameters. If the Regular Expression you enter is invalid, then the OK and Apply buttons are disabled.
2. Save you changes and close the form. The Routing Policy - AS Path Group, Global Policy form reappears.

5

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

End of steps

54.12 To configure an accounting template policy

54.12.1 Steps

1

Choose Policies→Routing→Accounting Template Policy from the NFM-P main menu. The Routing Policy - Accounting Template Policy form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy - Accounting Template, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Configure a destination class index for the policy.

1. Click on the Destination Class Indexes tab and click Create. The Destination Index form opens.
2. Configure the required parameter.
3. Save your changes close the form.

5

Configure a source class index for the policy.

1. Click on the Source Class Indexes tab and click Create. The Source Index form opens.
2. Configure the required parameter.
3. Save your changes close the form.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.13 To configure an administrative group policy

54.13.1 General Information

See 54.2.2 “Administrative Group policies” (p. 1740) in this chapter for more information.

54.13.2 Steps

1

Choose Policies→Routing→Administrative Group from the NFM-P main menu. The Routing Policy - Administration Group Policy form opens.

2

Click Create or choose a policy and click Properties. The Administrative Group Policy, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Click Apply to save the policy and to access the additional tabs on the form.

5

Configure the parameters or view the information associated with each tab as required. Table 54-6, “Description of Administrative group form tabs” (p. 1759) describes the tabs you can configure/view.

Table 54-6 Description of Administrative group form tabs

Tab	Description
Local Definitions	Lists Administrative group policies and allows you to manage administrative group policy distribution.
LSPs	Lists and allows you to manage LSPs to which the administrative group has been assigned.
LSP-Path Bindings	Lists and allows you to manage LSP paths to which the administrative group has been assigned.
Interfaces	Lists and allows you to manage interfaces to which the administrative group has been assigned. You can access the following interface types on the sub-tabs: • MPLS • L3 access • network • tunnel
P2MP LSPs	Lists and allows you to manage P2MP LSPs to which the administrative group has been assigned.
Route Next-Hop	Lists and allows you to manage the next-hop to which the administrative group has been assigned.
Flexible Algorithm Definition	Lists and allows you to manage flexible algorithm definitions to which the administrative group has been assigned.
Faults	List and manage alarms related to the administrative group.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.14 To configure a Shared Risk Link Group policy

54.14.1 General Information

See 54.2.3 “Shared Risk Link Group policies” (p. 1741) in this chapter for more information.

54.14.2 Steps

1

Choose Policies→Routing→Shared Risk Link Group from the NFM-P main menu. The Routing Policy - Shared Risk Link Group Policy form opens.

2

Click Create or choose a policy and click Properties. The Shared Risk Link Group, Global Policy (Create|Edit) form opens.

3

Configure the required parameters on the General tab.

4

Click Apply to save the policy. The form refreshes with additional tabs.

5

Configure the parameters or view the information associated with each tab as required. Table 54-7, “Description of Shared Risk Link Group Policy form tabs” (p. 1759) describes the tabs you can configure/view.

6

Configure the parameters or view the information associated with each tab as required. Table 54-7, “Description of Shared Risk Link Group Policy form tabs” (p. 1759) describes the tabs you can configure/view.

Table 54-7 Description of Shared Risk Link Group Policy form tabs

Tab	Description
Local Definitions	Lists Shared Risk Link Group policies and allows you to manage policy distribution.
Interfaces	Lists and allows you to manage MPLS interfaces to which the SRLG policy has been assigned. You can access the following interface types on the sub-tabs: • MPLS • L3 access • network • tunnel
Faults	List and manage alarms related to the SRLG.

7

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.15 To create a static configuration for a SRLG Policy

54.15.1 General Information

This procedure differs from the standard SRLG policy configuration described in 54.14 “To configure a Shared Risk Link Group policy” (p. 1759) by allowing you to manually enter the SRLG membership into the SRLG database for links in the entire network.

54.15.2 Steps

1

Choose Policies→MPLS→Shared Risk Link Group Static Configuration from the NFM-P main menu. The Shared Risk Link Group Static Configuration Policies form opens.

2

Click Create or choose a policy and click Properties. The Static Configuration for SRLGs Policy form opens.

3

Configure the required parameters on the General tab.

4

Click Apply to save the policy.

The Static Configuration for SRLGs Policy form is refreshed and the tabs become selectable. The following table describes the tabs that you can choose to configure parameters or view fault information.

Table 54-8 Description of Static Configuration for SRLG Policy form tabs

Tab	Description
Local Definitions	Lists Static Configuration for SRLG policies and allows you to manage policy distribution.
Routers	Lists and allows you to manage existing static configuration entries or add new ones.
Faults	List and manage alarms related to the static configuration for SRLGs.

5

Click on the Routers tab and click Create. The Router for Static Config for SRLG policy form opens.

6

Configure the required parameters.

7

Save your changes and close the form. The router(s) you added appears in the list on the Routers tab.

8

Repeat Step 5 to Step 7 if you need to add other routers, otherwise, go to Step 9 .

9

Select one of the routers and click Properties. The Router for Static Config for SRLG policy form opens.

10

Configure the Admin State parameter if required.

11

Click on the Interfaces tab and click Create. The Static Configuration for SRLGs Policy, Global Policies form opens.

12

Configure the required parameters.

Note: You can associate the same Interface Ip Address parameter with more than one SRLG.

13

Click Apply and repeat Step 9 to Step 12 to add other entries or save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.16 To configure a route next hop template policy

54.16.1 Steps

1

Choose Policies→Routing→Route Next Hop Template Policy from the NFM-P main menu. The Routing Policy - Route Next-Hop Template Policy form opens.

2

Click Create or choose a policy and click Properties. The Routing Policy - Route Next-Hop Template, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Configure included groups for the policy.

1. Click on the Include Group tab and click Create. The Route Next-Hop Admin Group Include Entry, Create Admin Group Entry form opens.
2. Click Select to choose an Admin Group Policy. The Select Admin Group Name - Route Next-Hop Admin Group Include Entry form opens.
3. Click Create or choose a group. The Administrative Group Policy, Global Policy form opens.
4. Configure the required parameters.
5. Click OK. The Select Admin Group Name - Route Next-Hop Admin Group Include Entry form reappears.
6. Select an entry and click OK. The Route Next-Hop Admin Group Include Entry form reappears.
7. Click OK. The Routing Policy - Route Next-Hop Template, Global Policy form reappears.

5

Configure excluded groups for the policy.

1. Click on the Exclude Group tab and click Create. The Route Next-Hop Admin Group Exclude Entry, Create Admin Group Entry form opens.
2. Click Select to choose an Admin Group Policy. The Select Admin Group Name - Route Next-Hop Admin Group Exclude Entry form opens.
3. Click Create or choose a group. The Administrative Group Policy, Global Policy form opens.
4. Configure the required parameters.
5. Click OK. The Select Admin Group Name - Route Next-Hop Admin Group Exclude Entry form reappears.
6. Select an entry and click OK. The Route Next-Hop Admin Group Exclude Entry form reappears.
7. Click OK. The Routing Policy - Route Next-Hop Template, Global Policy form reappears.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.17 To configure a re-assembly profile policy

54.17.1 Steps

1

Choose Policies→Routing→Re-Assembly Profile Policy from the NFM-P main menu. The Re-Assembly Profile Policies form opens.

2

Click Create or choose a policy and click Properties. The Re-Assembly Profile, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Configure a forwarding class.

1. Click on the Forwarding Class tab.
2. Click Create. The Forwarding Class Reassembly Profile form opens.
3. Configure the parameters and click OK.

5

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.18 To configure a maintenance policy

54.18.1 Steps

1

Choose Policies→Routing→ Maintenance Policy from the NFM-P main menu. The Manage Maintenance Policies form opens.

2

Click Create or choose a policy and click Properties. The Maintenance Policy, Global Policy form opens.

3

Configure the required parameters on the General tab.

4

Configure a BFD template

1. Click Select in the BFD Template panel. The Select BFD Template form opens.
2. Click Search to populate the list of available BFD templates.
3. Select a template or click Create. See 28.25 “To configure a BFD template policy” (p. 911) to create a BFD template.
4. Click OK to close the Select BFD Template form.

5

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with the global policy. See 49.9 “To change the distribution mode of a policy” (p. 1482) to change the distribution mode of a routing policy.

END OF STEPS

54.19 To configure global variables

54.19.1 Before you begin

For supporting NEs, you can create global variables and assign them to routing policy statements. Global variables are similar to policy variables, but allow for more efficient configuration, because they can be configured once, then selected in multiple routing policy statements; see 54.5 “To configure a routing policy statement” (p. 1745).

54.19.2 Steps

1

Choose Policies→Routing→Global Variables from the NFM-P main menu. The Routing Policy - Global Variables form opens.

2

Click Create or choose a variable and click Properties. The Routing Policy - Global Variables (Create|Edit) form opens.

3

Configure the Name parameter for the global variable. The name must begin and end with an “@” sign; for example: @variable-name@.

4

Configure the required parameters.

When you configure the Type parameter to Prefix, you must configure the Value parameter to an IP or IPv6 address and also configure the Prefix parameter.

A global variable may be an integer, an IPv4/IPv6 address, or a string. A string is a pointer to another routing policy type (for example, a Prefix List or Community policy). If you choose STRING for the Type parameter, you must create the required routing policy first, and click Select to choose and assign the policy to the global variable.

5

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Close the Routing Policy - Global Variables form.

END OF STEPS

54.20 To view policy variable usage in a routing policy statement

54.20.1 General Information

To perform this procedure, you must use a global routing policy statement that has been released and distributed to one or more NEs; see 54.5 “To configure a routing policy statement” (p. 1745) .

54.20.2 Steps

1

Choose Policies→Routing→Statement from the NFM-P main menu. The Routing Policy - Statements form opens.

2

Click Search and choose the required policy, then click Properties. The Routing Policy - Statement form opens.

3

Click on the Local Definitions tab. A list of NEs is displayed to which the policy was distributed.

4

Choose a site and click Properties. The Routing Policy - Statement form for the local definition opens.

5

Click on the Policy Statement Entries tab. A list of policies statement entries appears.

6

Choose an entry and click Properties. The Policy Statement Entry, Site form opens.

7

Click on the From Criteria tab.

8

Click Expand Policy adjacent to the Policy Statement Name field. A window opens and displays the show command from the NE. The policy variable names defined in the main routing policy statement are seen to be replaced by the actual routing policy names they represent, provided that those routing policies exist. If a particular routing policy does not exist, the name of the policy variable calling it is shown instead.

Note: Policy mediation must contain the username and/or password for the selected NE to allow the Expand Policy function to execute.
The Policy Statement Name field displays the name of the subordinate routing policy statement that calls the policy variables defined in this main routing policy statement.

9

Close all open forms as required.

End of steps

54.21 To view routing policy usage

54.21.1 Purpose

You can use the NFM-P to display which NEs are currently using a global or local version of certain routing policy type. For some policy types, for example, a route next hop template policy, you can view the OSPF/ISIS interfaces on which the policy was applied.

Note: To perform this procedure, you must use a routing policy that has been released and distributed to one or more NEs; see 49.6 “To release and distribute a policy” (p. 1476) .
Not all NFM-P routing policies support this functionality. If you can select the Policy Usage tab associated with a distributed routing policy, then the policy is supported.

54.21.2 Steps

1

Choose Policies→Routing→*Policy type* in the NFM-P main menu where *Policy Type* is the type of policy that you want to view the routing policy usage of. The appropriate policy manager form opens.

2

Click Search and choose the required policy, then click Properties. The appropriate Routing Policy form opens.

3

Click the Policy Usage tab and click Search.

The location of Policy Usage details varies for some routing policies. For example, for Routing Policy Statements, click on the required protocol sub-tab such as LDP, BGP, or OSPF for policy usage information.

4

Choose an entry and click Properties. The properties form opens and provides

5

Close all open forms as required.

END OF STEPS

54.22 To show a routing policy CLI configuration in the client GUI

54.22.1 General Information

To perform this procedure, you must use a routing policy that has been released and distributed to one or more NEs; see 49.6 “To release and distribute a policy” (p. 1476) . You must also ensure that the mediation security policy associated with the NE is configured to allow CLI access; see 9.17 “To configure device mediation” (p. 301) .

54.22.2 Steps

1

Choose Policies→Routing→*Policy type* in the NFM-P main menu where *Policy Type* is the type of policy that you want to view. The appropriate policy manager form opens.

2

Click Search and choose the required policy, then click Properties. The appropriate Routing Policy form opens.

3

Click on the Local Definitions tab. A list of NEs is displayed to which the policy was distributed.

4

Choose a site and click Properties. The Routing *Policy_type*, Local Policy form opens.

5

Click Show Policy. A Routing Policy Show Policy window opens and the NFM-P executes a “show router policy *type*” command on the NE. The Routing Policy Show Policy window displays the command output.

6

View the policy.

7

Close all open forms as required.

END OF STEPS

54.23 To verify BGP routes against a routing policy statement

54.23.1 Purpose

Perform the following procedure to verify BGP routes against a routing policy statement. The output of the test returns the prefixes that are accepted and rejected by the policy. This test allows you to evaluate an existing routing policy against the RIB to identify the prefixes that are matched or not matched by the policy, before attaching the prefix to a routing neighbor or instance.

54.23.2 Steps

1

Perform one of the following:

a. Perform a routing policy test on a base router instance BGP site.

1. In the navigation tree Routing view, expand Network→*NE*→Routing Instance→BGP.
2. Right-click on the BGP icon and choose Routing Policy Test. The rtr BGP Routing Policy Test form opens.

b. Perform a routing policy test on a VPRN instance BGP site.

1. Choose Manage→Service→Services from the NFM-P main menu.
2. Choose a VPRN service and click Properties.
3. Navigate to the BGP Site icon in the service navigation tree. The path is Sites→Routing Instance - NE System ID→Routing Instance→Protocols→BGP Site.
4. Right-click on the BGP Site icon and choose Routing Policy Test.

2

Configure the required parameters.

You must enter the name of an existing and distributed routing policy for the Policy Name parameter.

3

Click Execute.

4

Click Close.

End of steps

54.24 To verify a BGP prefix list against a routing policy statement

54.24.1 Purpose

Perform the following procedure to verify BGP route attributes against the routing policy statement. This test allows a routing policy statement to be checked against a local definition on a user-specified NE. The supported matching criteria include community lists and BGP prefix lists (BGP NLRI).

The test results lists the accepted and rejected route entries. Associated actions for accepted entries are performed. Actioned attributes include Communities, MED, Local Preference, and Next Hop.

54.24.2 Steps

1

Choose Tools→BGP Prefix Routing Policy Test from the NFM-P main menu.

2

Click Create.

3

Configure the general parameters and select a policy statement.

4

Click on the BGP Prefix Input tab and click Create.

5

Configure the input address and BGP neighbor peer parameters. IPv4 and IPv6 addresses are supported.

6

Configure the attributes to action in the Attributes panel.

7

Click OK.

8

Repeat Step 4 to Step 7 to create another BGP Prefix Input, if required.

9

Click Execute or Execute All. Execute is used to evaluate the policy on selected BGP Prefix Inputs.

10

Click on the BGP Prefix Output tab to list the results.

11

Close the forms.

END OF STEPS

54.25 To configure a route distinguisher policy

54.25.1 Steps

1

Choose Policies→Routing→Route Distinguisher from the NFM-P main menu. The Route Distinguisher, Global Policy (Edit) form opens.

2

Configure the required parameters on the General tab.

3

Click on the Route Distinguisher Entries tab.

4

Click Create. The Route Distinguisher Entry (Create) form opens.

5

Configure the Route Distinguisher Entry parameter and click OK.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

Note: A routing policy is discovered as local edit only. To synchronize a routing policy with all NEs, you must change the distribution mode of the routing policy to sync with global. To configure the distribution mode, see 49.6 “To release and distribute a policy” (p. 1476) .

END OF STEPS

55 VRRP policies

55.1 Overview

55.1.1 General Information

VRRP priority control-policies manage VR backup router priorities. The policies override the base priority value, depending on NE events or conditions. You can configure a VRRP policy only for the non-owner VRRP instance; the same policy can be applied to IPv4 and IPv6 VRRP instances.

The main function of a VRRP priority-control policy is to define the conditions or events that affect the VR ability to communicate with outside hosts or portions of the network. When at least one of these events is true, the base priority for the VR instance is affected in one of two ways:

- an explicit value is overridden
- a value is subtracted from the sum of delta priorities

The result is the actual in-use priority for the VR instance. Any priority event may be configured as an explicit event or a delta event. The following table describes the policy events that you can configure using the NFM-P.

Table 55-1 VRRP policy events

Policy Event	Description
Host unreachable	Configures a host unreachable priority-control event that monitors the ability of a host to receive an ICMP echo reply packet from a specific IP host address. A host unreachable priority-control event creates a continuous ICMP echo request (ping) probe to the specified IP address. During ping failure, the event is considered to be set. During ping success, the event is considered to be cleared.
IPv6 Host unreachable	Configures an IPv6 host unreachable priority-control event that monitors the ability of a host to receive an ICMP echo reply packet from a specific IPv6 host address. An IPv6 host unreachable priority-control event creates a continuous ICMP echo request (ping) probe to the specified IPv6 address. During ping failure, the event is considered to be set. During ping success, the event is considered to be cleared.
LAG port down	Configures a LAG priority-control event that monitors the operational state of the links and each port in the LAG. When one or more of the ports enters the operational down state, the event is considered to be set. When all ports enter an operational up state, the event is considered to be clear.
Route unknown	Configures a route unknown priority-control event that monitors the existence of a specific active IP route prefix in the routing table. Route unknown defines a link between the VRRP priority-control policy and the RTM. The RTM registers the specified route prefix as monitored by the policy. If any change (add, delete, new next hop) occurs relative to the prefix, the policy is notified and takes action according to the priority event definition.

Table 55-1 VRRP policy events (continued)

Policy Event	Description
Port Down Events	Configures an override or adjustment to the base priority value of a VRRP VR instance depending on the operational state of the event. Port Down events can only be configured on the local definitions of VRRP policies.
Multi Chassis IPsec Non Forwarding Events	Configures an MC IPsec priority-control event that monitors the MEP state of a device. When the MEP state is down, the event is considered to be set. When the MEP state is up, the event is considered to be clear. For more information about MC IPsec, see Chapter 41, “MC IPsec” .

As priority events clear in the policy, the in-use priority can eventually be restored to the base priority value.

55.2 To configure a VRRP priority-control policy

55.2.1 Steps

1

Choose Policies→VRRP from the NFM-P main menu. The Manage VRRP Policies form opens.

2

Perform one of the following:

a. Specify a filter to search for and edit an existing policy. Use the filters to search for and open an existing policy by choosing a policy in the filtered list and clicking on the Properties button.

b. Click Create button. The VRRP Policy (Create) form opens.

3

Configure the required parameters.

4

Click Apply.

5

Perform one of the following steps:

a. For IPv4 hosts, click on the Host Unreachable tab.

1. Click Create. The Host Unreachable Event (Create) form opens.
2. Configure the required parameters.
3. Click OK and save your changes. The Host Unreachable Event form reappears with the newly created filter entries displayed.

b. For IPv6 hosts, click on the IPv6 Host Unreachable tab.
   1. Click Create. The IPv6 Host Unreachable Event (Create) form opens.
   2. Configure the required parameters.

      The Interface Name parameter can be configured only when the host IP Address parameter is a Link Local Address.
   3. Click OK and save your changes. The IPv6 Host Unreachable Event form reappears with the newly-created filter entries displayed.

6

Perform the following to create a Lag Port Down event.

1. Click on the Lag Port Down tab.
2. Click Create. The Lag Port Down Event (Create) form opens.
3. Configure the required parameters.
4. Click on the Number Down tab.
5. Click Create. The Number Down (Create) form opens.
6. Configure the required parameters.
7. Close the Number Down (Create) form. A dialog box appears.
8. Click OK and save your changes. The Lag Port Down Event form reappears with the newly created filter entries displayed.

7

Perform the following to create a Route Unknown event.

1. Click on the Route Unknown tab.
2. Click Create. The Route Unknown (Create) form opens.
3. Configure the required parameters.
4. Click on the Next Hop tab.
5. Click Create.

   If the IP Address parameter you specify on the General tab is an IPv4 address, the NextHop (Create) form opens; if it is an IPv6 address, the NextHopV6 (Create) form opens.
6. Configure the Interface Name parameter if the IPv6 Next Hop Address is a Link Local Address.
7. Configure the Hop Address parameter.
8. Click OK and save your changes.

8

Perform the following to create a Multi Chassis IPsec Non Forwarding event.

1. Click on the Multi Chassis IPsec Non Forwarding tab.
2. Click Create. The MC IPsec Non Forwarding Event form appears.

3. Configure the required parameters.
4. Click OK and save your changes.

9

Click OK. The Manage VRRP Policies form reappears.

10

Click Search to display the newly-created policy or policies.

11

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

12

Close the Manage VRRP Policies form.

END OF STEPS

56 Auto tunnel policies

56.1 Auto tunnel policies

56.1.1 Overview

The NFM-P supports the automatic creation of service tunnels for groups of NEs, using mesh, hub-and-spoke, or ring topology rules. These rules, or auto tunnel policies, define the conditions for the service tunnel creation. An auto tunnel policy has the following components:

- a least one tunnel template
- a rules-based group to define which NEs are part of the tunnel group
- the topology rules, for example mesh or ring, that specify the service-tunnel characteristics, such as the type of tunnel and the underlying transport.

56.1.2 Tunnel templates

To create a tunnel template, you can create a new XML API template for service tunnel use or you can modify a provided LSP or SDP tunnel template. Use an example as a starting point to create service or tunnel templates with or without format and range policies. When a template with format and range policies is associated with an auto tunnel rule, the format and range policies defined in the template override the name defined in the auto tunnel rule.

Depending on the tunnel type selected for an auto tunnel, an SDP or LSP tunnel template can be selected. For an RSVP LSP tunnel type, one LSP template must be selected. For an RSVP SDP tunnel type, 16 LSP templates can be selected. When an LSP parent template with more than one LSP path child templates is specified, more than one LSP path is created.

A template should be configured to be deployed to a majority of the NE types. The user must validate the template and enter appropriate changes to the script before the template can be deployed in an auto tunnel policy. The parameters that can be modified depend on the template type.

56.1.3 Tunnel groups

A tunnel group is defined as collection of network resources, such as NEs, that perform the same role in the network; for example, the way that the designation of a port as an access or network port defines the role of the port in the network.

By grouping resources by role, a common policy-based management framework can be used to ensure the appropriate configuration of resources, for example, create service tunnels for different topologies.

The following conditions apply to tunnel groups.

- An NE can belong to multiple groups.
- A group is ordered or unordered depending on the requirements. The members in a group must

be ordered when you create some topologies; for example, a ring topology. Topologies such as mesh or hub-and-spoke, in which NEs are not linearly arranged, do not need an ordered group.

- A non-NFM-P-managed NE can belong to the tunnel group to which it was added, however, it is not active as a destination NE during rule execution.
- When two or more tunnel groups share an NE, the shared NE performs the roles defined for each group. In Figure 56-1, “NE shared by two tunnel groups” (p. 1777), Group 1 contains routers NE1, NE2, NE3, and NE4, and Group 2 contains routers NE3, NE5, NE6, and NE7. NE3 is the intersection set of Groups 1 and 2, therefore NE3 assumes the roles defined for Groups 1 and 2.

Figure 56-1 NE shared by two tunnel groups

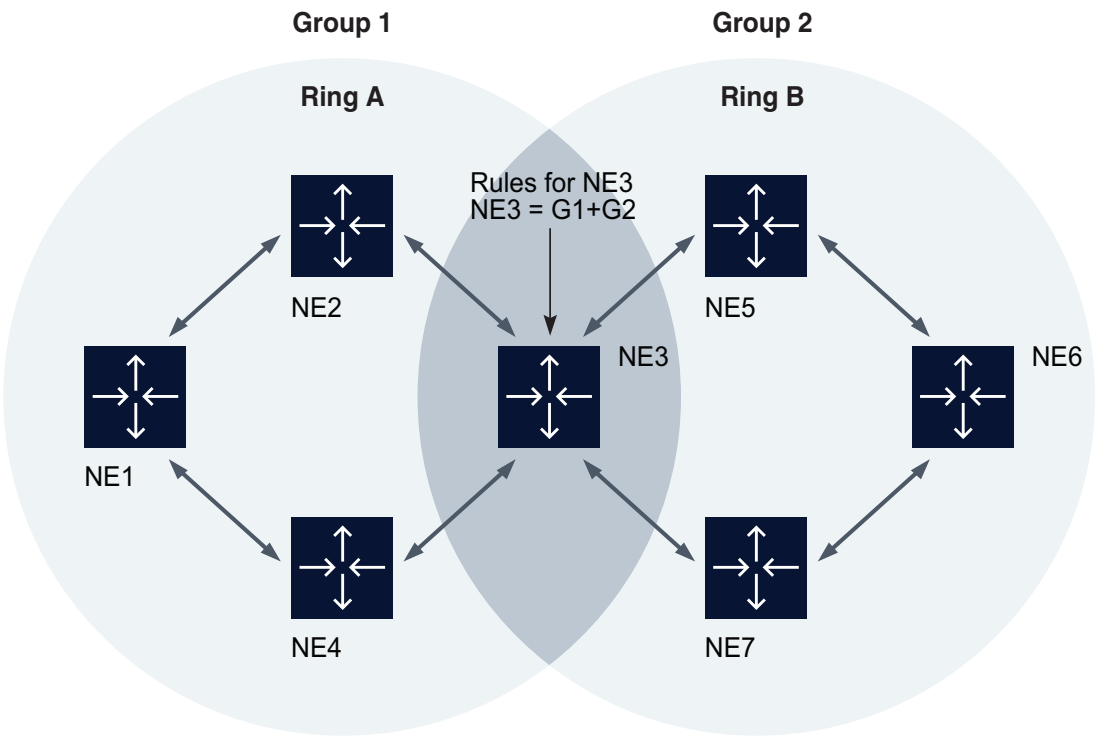

18983

56.2 Auto tunnel design considerations

56.2.1 General Information

The following applies to the configuration of auto tunnels:

- An NFM-P operator requires topology-management privileges to create or modify an auto tunnel.
- When an NE in a rule-based group is unmanaged or deleted, it is identified as unmanaged in the

rule-based group. Auto tunnels to the NE are not removed; however auto tunnels from the NE are removed from the NFM-P.

- When an auto tunnel rule is created for a group that contains non-NFM-P-managed NEs, auto tunnels cannot be created to and from the NE using that rule.
- When you add NEs to a tunnel group that is in use, new auto tunnels are created according to the rule definition.
- A group condition must contain at least one group of NEs.
- Auto tunnel elements are not updated dynamically.
- An unmanaged NE is removed from the tunnel groups to which the NE belongs. The associated service tunnels are also removed if they originate from a managed NE. A service tunnel that originates on an unmanaged NE is retained, but is not included in the tunnel group after reconfiguration.
- When a service tunnel created by a tunnel policy is removed, the policy does not attempt to recreate the tunnel.
- You can modify the tunnel definition in a tunnel policy when the tunnel policy has not yet created a service tunnel.

Not all NEs that are assigned to a topology rule share the same functionality. The topology rule runs capability checks on the source NEs to determine whether specific tunnels or tunnel configurations can be applied. The following table lists the auto tunnel rules that are supported by each NE.

Table 56-1 Auto tunnel rules and NE support

Auto tunnel rule	7450 ESS, 7750 SR, and 7950 XRS	7705 SAR
SDP LDP	X	X
SDP GRE	X	X
SDP LSP	X	X
Dynamic LSP	X	X
CoS-based SDP	X	—
LDP-over-RSVP	X	—

56.3 Workflow to configure auto tunnels

56.3.1 Stages

1

Configure a tunnel template from one of the following:

a. Choose Tools→Scripts from the NFM-P main menu to create a new XML API template.

b. Choose Manage→Templates from the NFM-P main menu to modify a provided LSP or SDP tunnel template.

2

Configure rule-based groups to define which NEs are part of the tunnel group; see 56.4 “To configure rule-based groups” (p. 1780).

3

Configure the required topology rules to allow for service-tunnel creation:

- for mesh or ring topology rules; see 56.5 “To create a mesh or ring topology rule” (p. 1782).
- for hub and spoke topology rules; see 56.6 “To create a hub-and-spoke topology rule” (p. 1783).

4

Configure a service and select the Automatic Mesh SDP Binding Creation parameter with the service to bind it to the previously created service tunnels. See the appropriate service chapter for more information.

Incidental tasks

5

Perform the following tasks, as required to manage auto tunnels:

- import tunnels not managed by topology rules; see 56.7 “To import tunnels not managed by topology rules” (p. 1785).
- display and delete tunnel elements; see 56.8 “To display and delete tunnel elements” (p. 1786).
- manually execute or reapply an auto tunnel topology rule; see 56.9 “To manually execute or reapply an auto tunnel topology rule” (p. 1786).
- display missing tunnel elements; see 56.10 “To view missing tunnel elements” (p. 1787).

56.4 To configure rule-based groups

56.4.1 Purpose

Perform this procedure to create or modify a rule-based group that defines which NEs are part of the tunnel group.

56.4.2 Steps

1

Choose Policies→Auto Tunnels→Rule-Based Groups from the NFM-P main menu. The Rule-Based Groups form opens.

2

Click Create and choose Rule-Based NE Group, or choose an entry and click Properties. The Rule-Based NE Group (Create|Edit) form opens.

3

Configure the required general parameters.

The Order parameter is configurable only during group creation.

4

If you are creating a new group, click Apply to save the configuration and create the group. The Group Members tab becomes available.

5

Click on the Group Members tab.

6

If you are configuring member NEs for an unordered group, perform the following. Otherwise, go to Step 7.

1. To delete existing member NEs, select the required NEs and click Delete.
2. To add member NEs, click Create. The Select Network Elements form opens.
3. Select the required NEs and click OK.
4. Go to Step 8.

7

To configure member NEs for an ordered group, perform any of the following:

a. To delete existing member NEs, select the required NEs and click Delete.

b. Add NEs in sequence.

   1. Click Create. The Select Network Elements form opens.
   2. Choose an NE and click OK.
   Add the NEs one at a time, in the required order.

   The order of NEs is shown in the Position in Group column on the Group Members tab. Position in Group is based on the sequence in which you add the NEs.

c. Insert an NE at a selected position in the group.

   1. Select a Position in Group in the list on the Group Members tab.
   2. Click Insert. The Select Network Elements form opens.
   3. Choose an NE and click OK. The NE is inserted in the list on the Group Members tab, at the position selected. From that position, existing members are displaced by one position.

8

Save your changes and close the form.

END OF STEPS

56.5 To create a mesh or ring topology rule

56.5.1 Purpose

Perform this procedure to create a mesh or ring topology rule for service tunnel creation. A template must be applied to an auto-tunnel policy. You must configure an LSP or SDP tunnel template before you can deploy a template to an LSP or SDP tunnel. .

56.5.2 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Click Create and choose one of the following:

a. Mesh Topology Rule. The Mesh Topology Rule (Create) form opens.

b. Ring Topology Rule. The Ring Topology Rule (Create) form opens.

3

Configure the required general parameters.

4

Add one or more rule-based groups.

1. Click on the LERs Definition tab.
2. Click Add and choose one or more tunnel groups.

5

Click on the Tunnel Definition tab and configure the Tunnel Type parameter.

6

Perform one of the following:

a. If you chose SDP as the tunnel type:

   1. Click on the SDP tab.
   2. Click Select and choose an SDP Template.
   3. Configure the required parameters.

4. If you chose RSVP-LSP or Mixed LSP Mode as the underlying transport type, configure the additional parameters. Otherwise, go to Step 7 .
5. Click on the LSPs tab, then click Add and choose an LSP template.

b. If you chose RSVP-LSP as the tunnel type:

1. Click on the LSP tab.
2. Click Select and choose an LSP Template.
3. Configure the required parameters.

7

Save your changes and close the form.

END OF STEPS

56.6 To create a hub-and-spoke topology rule

56.6.1 Purpose

Perform this procedure to create a hub-and-spoke topology rule for service tunnel creation. A template must be applied to an auto-tunnel policy. You must configure an LSP or SDP tunnel template before you can deploy a template to an LSP or SDP tunnel.

56.6.2 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Click Create and choose Hub and Spoke Topology Rule. The Hub and Spoke Topology Rule (Create) form opens.

3

Configure the required general parameters.

4

Add rule-based groups.

1. Click on the Hub LERs Definition tab.
2. Click Add and choose one or more tunnel groups.
3. Click on the Spoke LERs Definition tab.
4. Click Add and choose one or more tunnel groups.

5

Click on the Hub Tunnel Definition tab and configure the Tunnel Type parameter.

6

Perform one of the following:

a. If you chose SDP as the tunnel type:
   1. Click on the SDP tab.
   2. Click Select and choose an SDP Template.
   3. Configure the required parameters.
   4. If you chose RSVP-LSP or Mixed LSP Mode as the underlying transport type, configure the additional parameters. Otherwise, go to Step 7.
   5. Click on the LSPs tab, then click Add and choose an LSP template.
   6. If the Class Forwarding Capability parameter in 4 is enabled, the buttons in the Class Forwarding panel are available. If the buttons are not available, go to Step 7.
   7. Click Set as Default LSP to specify the LSP as the default, if required.
   8. Click Set as Multicast LSP to specify that the LSP is a multicast LSP, if required.
   9. Click Choose a Forwarding Class and choose a forwarding class from the contextual menu, if required.

b. If you chose RSVP-LSP as the tunnel type:
   1. Click on the LSP tab.
   2. Click Select and choose an LSP Template.
   3. Configure the required parameters.

7

Click on the Spoke Tunnel Definition tab and configure the Tunnel Type parameter.

8

Perform one of the following:

a. If you chose SDP as the tunnel type:
   1. Click on the SDP tab.
   2. Click Select and choose an SDP Template.
   3. Configure the required parameters.
   4. If you chose RSVP-LSP or Mixed LSP Mode as the underlying transport type, configure the additional parameters. Otherwise, go to Step 9.
   5. Click on the LSPs tab, then click Add and choose an LSP template.

b. If you chose RSVP-LSP as the tunnel type:
   1. Click on the LSP tab.
   2. Click Select and choose an LSP Template.

3. Configure the required parameters.

9

Save your changes and close the form.

END OF STEPS

56.7 To import tunnels not managed by topology rules

56.7.1 Purpose

Perform this procedure to import a tunnel that is not managed by a topology rule. After you import the tunnel, the tunnel is managed by the rule. However, the rule does not enforce the source and destination IDs to match the rule group. If the imported tunnel is not part of the group, the execute button or change notifications by the group members do not apply, even though the imported tunnel is part of the rule.

56.7.2 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Select a topology rule type from the object drop-down menu.

3

Choose a rule and click Properties. The properties form for the rule opens.

4

Click on the Created/Imported Tunnel Elements tab, then click Import Tunnels. The Import Tunnels form opens.

5

Click Add SDPs or Add LSPs and choose the required tunnels.

6

For topology rules where the tunnel definition is SDP and the underlying transport is RSVP-LSP, you can associate the imported LSP with an LSP template to avoid unnecessary LSP creation by the rule.

1. Choose an LSP from the list of imported LSPs.
2. Click Associate Template With Imported LSP and choose an LSP template from the list of LSP templates that are associated with the rule.

The Associate Template With Imported LSP button is available only for SDP-LSP rules that support the creation of multiple LSPs.

7

Click on the Import button. The Import Tunnels form closes.

8

Close the forms.

END OF STEPS

56.8 To display and delete tunnel elements

56.8.1 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Select a topology rule type from the object drop-down menu.

3

Choose a rule and click Properties. The properties form for the rule opens.

4

Click on the Created/Imported Tunnel Elements tab.

5

Click Search and perform one of the following to delete tunnel elements as required:

a. Click Delete Unused to delete tunnels that are no longer applied to a group.

b. Click Delete All to delete all tunnels that are not associated with a service.

6

Confirm your changes and close the forms.

END OF STEPS

56.9 To manually execute or reapply an auto tunnel topology rule

56.9.1 Purpose

Executing a topology rule causes an evaluation and reapplication of the rule.

Note: The NFM-P supports the reapply function, except in reference to hop-less path changes. If a hop-less path changes, it may not be corrected by a reapply operation.

You do not need to perform this procedure if the Auto-Rule Execution parameter was set to enabled when the rule was created. When the parameter is enabled, the rule is automatically applied to any objects that are added to the rule-based group.

56.9.2 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Select a topology rule type from the object drop-down menu.

3

Choose a rule and click Properties. The properties form for the rule opens.

4

Click Execute.

5

Close the form.

END OF STEPS

56.10 To view missing tunnel elements

56.10.1 Purpose

A missing tunnel element is one that the NFM-P is unable to create based on the specified tunnel rules, or a tunnel element that is deleted. Use this procedure to view such elements.

56.10.2 Steps

1

Choose Policies→Auto Tunnels→Auto Tunnel Rules from the NFM-P main menu. The Auto Tunnel Rules form opens.

2

Select a topology rule type from the object drop-down menu.

3

Choose a rule and click Properties. The properties form for the rule opens.

4

Click on the Missing Tunnel Elements tab.

5

Choose a missing tunnel element and click Properties. The properties form for the element opens.

6

View the element properties.

7

Close the forms.

END OF STEPS

57 AAA policies

57.1 AAA policy types

57.1.1 Overview

AAA policies provide security for network traffic on one or more NEs. You can use the NFM-P to create AAA policies that provide authentication, authorization, and accounting functionality. The NFM-P supports the following AAA policies:

57.1.2 Accounting on/off groups

An accounting on/off group defines a group of RADIUS server policies for the purpose of synchronizing accounting on/off messages across all policies in the group. It establishes which policy in the group acts as the controller policy, and which policies act as monitors. The accounting on/off state of the monitor policies is set by the controller policy.

The controller/monitor designation is configured on each RADIUS server policy in the group. Only one policy in a group can be designated as the controller.

57.1.3 ISA RADIUS policies

An ISA RADIUS policy defines the network security requirements used when an ISA server is configured to use a RADIUS server for AAA client requests.

The policy can also be configured to limit the sending rate of periodic logging messages to the RADIUS server, thereby reducing the chance of a backlog of pending messages accumulating at the RADIUS server.

57.1.4 L2TP RADIUS accounting policies

An L2TP RADIUS accounting policy defines the usage data collected for subscribers based on either an L2TP tunnel or an L2TP session and sends this data to a RADIUS server.

57.1.5 NAT RADIUS accounting policies

To increase logging performance (messages-per-second) and provide a reliable logging facility, a RADIUS accounting option is provided for Large Scale NAT. The configuration of a NAT RADIUS policy and its assignment to a NAT ISA group (NAT or WLAN GW) will cause the node to instruct the RADIUS user to start logging when a port-range-block is assigned to a user, or to stop logging when this block is released back into the pool.

57.1.6 RADIUS-based accounting policies

A RADIUS-based accounting policy is used to send accounting information to a RADIUS server.

When an NE uses a RADIUS-based accounting policy and the creation of a subscriber host invokes the policy, the NE generates an accounting-start packet that includes the policy parameters, and forwards the packet to a RADIUS server. Depending on the policy specifications, the NE also sends

interim update messages that contain usage statistics. When a policy is no longer used by any subscriber host, the NE sends an accounting-stop packet that contains the final usage statistics.

A RADIUS-based accounting policy specifies the type of accounting information that is forwarded to a RADIUS server. The information is one of the following:

- per-SLA profile—forwards start, stop, and interim update messages regarding SLA profile usage
- per-host and per-SLA profile—forwards the following:
  - start, stop, and interim update messages regarding SLA profile usage
  - start and stop messages for individual host sessions
- per-host only—forwards start, stop, and interim update messages for each subscriber host session

When a subscriber host disconnects, the NE sends an accounting-stop packet that contains the final usage statistics. An accounting-stop packet is also sent when a subscriber or subscriber host is deleted, or when an SLA profile instance (non-HSMDA) or subscriber instance (HSMDA) is changed.

You can reduce the volume of data that is generated by the accounting policy and sent to the RADIUS server by defining custom records. Custom records include specific counters in RADIUS accounting messages. Custom records eliminate queues or selected counters within these queues, that are not relevant for billing.

You can further decrease the number of accounting messaging by including only objects which have experienced a significant change in the specified counters. When this significant change is met, the record is generated.

57.1.7 RADIUS script policies

The RADIUS script policy references python scripts for RADIUS AAA packet manipulation in a subscriber management application. See 64.23 “To configure a RADIUS script policy” (p. 1865).

Operational states of primary and secondary RADIUS scripts

When the Up state is selected for primary or secondary scripts, the NE downloads the Python script from the specified path and copies it to the memory. If the path is invalid or the script is too long or contains syntax errors, the operational state of the script will change to the Down state. If the Python script installs correctly, the operational state of the script will remain in the Up state.

Script modifications are only applied by the NE after the script admin state has been changed to Down and then Up again.

57.1.8 RADIUS server policies

A RADIUS server policy is used to configure access to a group of RADIUS servers.

A maximum of five RADIUS server entries can be created under each RADIUS server policy. These RADIUS server entries must be created under the specified routing instance. For the global policy, the RADIUS server entries can be created within all base routing instances if Base routing instance is specified, or within the sites of the specified VPRN service if VPRN routing instance is specified. A RADIUS server entry cannot be created if Management routing instance is specified.

You can associate RADIUS server policies to Ethernet ports under 802.1x. See 57.8 “To associate a RADIUS Server policy to an Ethernet port for dot1x authentication” (p. 1799) for information.

57.1.9 Route download policies

A route download policy provides a mechanism for an NE to download (in advance) customer-assigned subnets from a RADIUS server, so that they can be re-advertised to the corresponding routing protocols. In this manner, subscriber connections can potentially be established faster because the routes are already in place. Routing protocol churn is reduced as subscribers connect and disconnect.

57.1.10 Subscriber authentication policies

A subscriber authentication policy uses RADIUS authentication to grant network access to a dynamic host. These DHCP-based policies define the parameters for dynamically-created subscriber host sessions and authenticate the sessions. They can be applied to a VPLS or IES SAP, or to a VPRN or IES group interface.

You can configure a subscriber authentication policy for LLID pre-authentication. The policy is specified as a pre-authentication policy on a local user database, as described in 74.9 “To configure a local user database for subscriber host authentication” (p. 2025).

57.1.11 Diameter peer policies

Diameter peer policies are used in a subscriber management context to provide a credit control mechanism. The diameter peer policy establishes a server/peer configuration. The NE functions as the credit control client, while the peer acts as the credit control server. The diameter peer policy is used to specify common diameter protocol parameters, while the diameter peer defines the relationship with an external diameter server. The diameter peer policy can be configured with a python policy, allowing the diameter policy to be modified by python scripts and diameter python messages.

A diameter peer policy can be configured to function as part of a redundancy configuration, based on a Gx proxy model. Two NEs run instances of Gx proxy, with one Gx proxy instance active and the other inactive. The active instance has a peering connection to DRA. The purpose of the redundancy is predictable operations recovery after an NE failure. A diameter peer policy with proxy configuration is associated with a diameter application policy which, in turn, is applied to a group interface. The diameter proxy monitors the group interface.

Peers

The diameter peer policy defines a set of peers with which to establish diameter sessions. Peers share the configuration of the policy with which they are associated, but can override individual timer parameters inherited from the policy. In addition, each peer defines transport and connection-specific parameters, values for destination-specific AVPs, and a preference value.

57.2 Workflow to configure AAA policies

57.2.1 Stages

1

As required, configure a RADIUS-based accounting policy. See 57.3 “To configure a RADIUS-based accounting policy” (p. 1793).

2

As required, configure a NAT RADIUS accounting policy. See 57.4 “To configure a NAT RADIUS accounting policy” (p. 1795).

3

As required, configure an L2TP RADIUS accounting policy. See 57.5 “To configure an L2TP RADIUS accounting policy” (p. 1796).

4

As required, configure a RADIUS server policy. See 57.6 “To configure a RADIUS server policy” (p. 1797).

5

As required, initiate an accounting on/off message from a RADIUS server policy. See 57.7 “To initiate an accounting on/off message from a RADIUS server policy” (p. 1798).

6

As required, configure an accounting on/off group. See 57.9 “To configure an accounting on/off group” (p. 1800).

7

As required, configure a route download policy. See 57.10 “To configure a route download policy” (p. 1801).

8

As required, force a route download. See 57.15 “To force a route download on an NE” (p. 1806).

9

As required, configure a subscriber authentication policy. See 57.11 “To configure a subscriber authentication policy” (p. 1802).

10

As required, configure an ISA RADIUS policy. See 57.12 “To configure an ISA RADIUS policy” (p. 1803).

11

As required, configure a diameter peer policy. See 57.13 “To configure a diameter peer policy” (p. 1804).

57.3 To configure a RADIUS-based accounting policy

57.3.1 Steps

1

Choose Policies→AAA Policies→RADIUS Based Accounting from the NFM-P main menu. The RADIUS Based Accounting Policies form opens.

2

Click Create or choose a policy and click Properties. The RADIUS Accounting Policy (Create|Edit) form opens.

3

Configure the general policy and RADIUS Attributes parameters.

If you enable the NAS Port ID checkbox, you must configure the Port Prefix Type, Port Prefix String, and Port Suffix Type parameters.

If you enable the NAS Port Type checkbox, you must configure the Port Type parameter.

If you enable the NAS Port Type checkbox and you set the Port Type parameter to Config, you must configure the Port Type Value parameter.

If you enable the Calling Station ID checkbox, you must configure the Calling Station ID Type parameter.

If you enable the NAS Port checkbox, you must configure the Port Binary Specification parameter.

4

Configure the required parameters in the Accounting and Interval panels.

Only two accounting modes can be applied in combination. The accounting modes are host accounting, queue instance accounting, and PPPoE session accounting. A subscriber host can have up to two simultaneously active accounting modes. You can configure Host Accounting Message, Queue Instance Accounting, and Session Accounting parameters to non-default values provided that SLA Only is chosen for one parameter, and not more than two accounting modes are combined.

If you configure the Delay Start Time parameter to reduce accounting messages in a dual stack PPP session configuration, you must also configure the Optimize for Session Accounting parameter on the subscriber profile to which the RADIUS accounting policy is applied; see 64.4 “To configure a subscriber profile” (p. 1840).

5

Select an accounting request script in the RADIUS Script Policy panel.

6

Select a RADIUS server policy in the RADIUS Server Policy panel.

7

Configure one or more RADIUS servers for accounting information forwarding:

1. Click on the RADIUS Servers tab and configure the required parameters.

   If you set the Router Instance parameter to VPRN, you must configure a VPRN service as a virtual router instance by selecting a VPRN service in the Service Type panel.
2. Click Create in the Servers panel.
3. Configure the required parameters.
4. Save your changes and close the form.

8

To create one or more custom queue counter configuration for a custom statistics record:

1. Click on the Custom Record tab, then on the Queue Counter Config tab.
2. Click Create or select a queue and click Properties. The Custom Queue Config (Create|Edit) form opens.
3. Configure the required parameters.

   You must select one or more ingress or egress counters in order to apply them to reference counters to the significant change delta]
4. Save your changes and close the form.

9

To create one or more custom override counter configurations you must select one or more ingress/egress counters in order to apply them to reference counters to the significant change delta.

1. On the Custom Record tab, click on the Override Counter Config tab.
2. Click Create or select an override and click Properties. The Custom Override Config (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

10

Configure record criteria that define the RADIUS statistics that the NFM-P collects per queue or for all queues:

1. Click on the Significant Change Criteria tab.
2. Configure the Significant Change Delta parameter.
3. Select a queue ID to monitor for the significant change or enable the All Queues parameter in the Reference Queue panel.
4. Configure the ingress and egress counters.

5. Select an override ID or enable the All Overrides parameter in the Reference Override panel.
6. Configure the ingress and egress counters.

11

Save the policy and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.4 To configure a NAT RADIUS accounting policy

57.4.1 Steps

1

Choose Policies→AAA Policies→NAT RADIUS Accounting from the NFM-P main menu. The NAT RADIUS Accounting Policies form opens.

2

Click Create or select an existing NAT RADIUS accounting policy and click Properties. The NAT RADIUS Accounting Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Configure the Router Instance parameter. If you set the Router Instance parameter to VPRN, perform Step 5. Otherwise, go to Step 6.

The Router Instance parameter is configurable on the 7450 ESS and 7750 SR.

5

To configure a VPRN service as a virtual router instance for the NAT RADIUS accounting policy, select a VPRN site in the VPRN ID panel.

6

Configure RADIUS servers for the policy:

1. Click on the RADIUS Servers tab.
2. Click Create or select an existing RADIUS server and click Properties. The RADIUS Server (Create|Edit) form opens.
3. Configure the required parameters.

7

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.5 To configure an L2TP RADIUS accounting policy

57.5.1 Purpose

Perform the following procedure to create an L2TP RADIUS accounting policy to collect usage data based on either an L2TP tunnel or an L2TP session and send this data to a RADIUS server. See 28.89 “To configure L2TP on a routing instance” (p. 989) for information on configuring an L2TP RADIUS accounting policy on a routing instance.

57.5.2 Steps

1

Choose Policies→AAA Policies→L2TP RADIUS Accounting from the NFM-P main menu. The L2TP RADIUS Accounting Policies form opens.

2

Click Create or select an existing L2TP RADIUS accounting policy and click Properties. The L2TP RADIUS Accounting Policy (Create|Edit) form opens.

3

Select a RADIUS Script policy.

4

Select a RADIUS server policy.

5

Configure the required parameters.

The Definition and Value parameters are only configurable when the NAS Port Type option is enabled.

The Bit Specification parameter is only configurable when the NAS Port option is enabled.

The Prefix Type, Prefix, and Suffix Type parameters are only configurable when the NAS Port ID option is enabled.

6

Click on the RADIUS Servers tab.

7

Configure the required parameters.

The Source Address parameter only appears when you are editing an existing local version of an L2TP RADIUS accounting policy.

8

Configure the Router Instance parameter. If you set the Router Instance parameter to VPRN, perform Step 9. Otherwise, go to Step 10.

The Router Instance parameter is configurable on the 7450 ESS and 7750 SR.

9

To configure a VPRN service as a virtual router instance for the L2TP RADIUS-based accounting policy, select a VPRN service in the VPRN ID panel.

10

Click Create or select an existing L2TP RADIUS server and click Properties. The L2TP RADIUS Server (Create|Edit) form opens.

11

Configure the required parameters.

12

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.6 To configure a RADIUS server policy

57.6.1 Steps

1

Choose Policies→AAA Policies→RADIUS Server from the NFM-P main menu. The Radius Server Policies form opens.

2

Click Create or select an existing RADIUS server policy and click Properties. The Radius Server Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Configure the Router Instance parameter. If you set the Router Instance parameter to VPRN, select a VPRN service in the VPRN ID panel.

5

Select a Python policy.

6

Configure RADIUS scripts for the policy:

1. Select a RADIUS Accept script.
2. Select a RADIUS authentication request script.
3. Select a RADIUS accounting request script.

7

Click on the Accounting On/Off tab to configure an accounting on/off group for the policy.

1. Configure the Admin State parameter. perform 2.
2. If you set the Admin State parameter to Enabled With State Change or Enabled With Monitoring, select an accounting on/off group.

   Only one RADIUS server policy in an accounting on/off group can be configured with an Admin State parameter of Enabled With State Change.

8

Click on the RADIUS Servers tab to configure RADIUS servers for the policy.

1. Click Create or select an existing RADIUS server and click Properties. The Server Entry (Create|Edit) form opens.
2. Configure the required parameters.

9

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.7 To initiate an accounting on/off message from a RADIUS server policy

57.7.1 Purpose

Perform the following procedure to manually send an Accounting On or Accounting Off message from a RADIUS server policy. The command is enabled only on a RADIUS sever policy that is configured as the controller in an accounting on/off group. In order to perform this procedure, the RADIUS server policy must be distributed to an NE (see Chapter 49, “Policies overview”).

57.7.2 Steps

1

Choose Policies→AAA Policies→RADIUS Server Policies from the NFM-P main menu. The Manage Radius Server Policies form opens.

2

Select a RADIUS server policy and click on the Properties button. The RADIUS Server Policy (Edit) form opens.

3

Click on the Send Accounting On/Off Message button and confirm the action.

4

In the Accounting On/Off Operation form, configure the Action parameter.

5

Click OK to send the Accounting On or Accounting Off message.

6

Close the forms.

END OF STEPS

57.8 To associate a RADIUS Server policy to an Ethernet port for dot1x authentication

57.8.1 Purpose

The RADIUS Server Authentication Policy and RADIUS Server Accounting Policy support IPv6 connectivity for dot1x. Perform this procedure to associate these Radius server policies to Ethernet ports under the 802.1x Port Authenticator tab on the Ethernet Physical Port configuration form.

57.8.2 Steps

1

Create one or more RADIUS servers.

2

On the equipment tree, expand Network→*NE*→Shelf→Card Slot *n*→Daughter Card Slot *n*→Port *n/n/n*.

3

Associate the RADIUS servers under AAA→Radius Server Policy and distribute to the NE

4

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

5

Click on the 802.1x Port Authenticator tab and then click Select beside the RADIUS Server Policy Authentication and RADIUS Server Policy Accounting parameters, as required. A list of the AAA RADIUS server policies displays.

6

Select and apply the policies, as required.

7

Save your changes and close the form.

END OF STEPS

57.9 To configure an accounting on/off group

57.9.1 Steps

1

Choose Policies→AAA Policies→Accounting On/Off Group from the NFM-P main menu. The Accounting On/Off Groups form opens.

2

Click Create or select an existing accounting on/off group and click Properties. The Accounting On/Off Group (Create|Edit) form opens.

3

Configure the Displayed Name and Description parameters.

4

To manage the RADIUS server policies assigned to the accounting on/off group, click on the RADIUS Server Policies tab.

The controller policy name is displayed in the Controller panel. The monitor policies are listed in the Monitors panel.

To change a controller policy to a monitor, open its Properties form and change its accounting on/off Admin State to Enabled With Monitoring (see 57.6 “To configure a RADIUS server policy” (p. 1797)).

To change a monitor policy to a controller, open its Properties form and change its accounting on/off Admin State to Enabled With State Change (see 57.6 “To configure a RADIUS server policy” (p. 1797)).

5

Save your changes and close the forms. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.10 To configure a route download policy

57.10.1 Purpose

In order to configure a route download policy, you must have a RADIUS server configured on each NE that is to be associated with the route download policy (see 27.8 “To configure a RADIUS server on a routing instance” (p. 846)), and you must have a RADIUS server policy configured (see 57.6 “To configure a RADIUS server policy” (p. 1797)).

57.10.2 Steps

1

Choose Policies→AAA Policies→Route Download from the NFM-P main menu. The Route Download Policies form opens.

2

Click Create or select an existing route download policy and click Properties. The Route Download Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click Apply. The Route Download Policy form refreshes with additional tabs.

5

Select a RADIUS server policy.

6

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.11 To configure a subscriber authentication policy

57.11.1 Purpose

Perform this procedure to create a DHCP-based subscriber authentication policy that defines the parameters for dynamically created subscriber host sessions and authenticates the sessions. The NFM-P supports up to 32 subscriber authentication policies.

57.11.2 Steps

1

From the NFM-P main menu, choose Policies→AAA Policies→Subscriber Authentication. The Subscriber Authentication form opens.

2

Click Create or select an existing subscriber authentication policy and click Properties. The Subscriber Authentication Policy (Create|Edit) form opens.

3

Configure the required parameters.

The NAS Port Bit Specification parameter is configurable when the NAS Port option is enabled for the RADIUS Attributes parameter.

The Calling Station ID Type parameter is configurable when the Calling Station ID option is enabled for the RADIUS Attributes parameter.

The Port Type parameter is configurable when the NAS Port Type option is enabled for the RADIUS Attributes parameter.

The Port Type Value parameter is configurable when the NAS Port Type option is enabled for the RADIUS Attributes parameter, and the Port Type parameter is set to Config.

The Port Prefix Type, Port Prefix String, and Port Suffix Type parameters are configurable when the NAS Port ID option is enabled for the RADIUS Attributes parameter.

4

If you need to configure PAP/CHAP user name re-writing, configure the required parameters in the User Name panel.

The Domain Name parameter is configurable when the User Name Operation parameter is set to Append, Replace, or Use As Default.

5

In the RADIUS Fallback panel, configure the required parameters.

The Local User DB Name parameter is configurable when the Fallback Action parameter is set to User DB.

6

In PPP User Name panel, configure the required parameters.

The Domain Name parameter is configurable when the User Name Operation parameter is set to Append, Replace, or Use As Default.

7

Configure RADIUS script policies.

1. Select an accept script.
2. Select a CoA script.
3. Select a request script.

8

Select a RADIUS server policy.

9

Click on the RADIUS Servers tab to configure the RADIUS servers for the policy.

1. Configure the required parameters:

   If you set the Router Instance parameter to VPRN, select a VPRN site on the VPRN ID panel.
2. Click Create. The RADIUS Entry form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

10

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.12 To configure an ISA RADIUS policy

57.12.1 Steps

1

Choose Policies→AAA Policies→ISA RADIUS from the NFM-P main menu. The ISA RADIUS Policies form opens.

2

Click Create or select an existing ISA RADIUS policy and click Properties. The ISA RADIUS Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a Python policy.

5

Save your changes and close the form. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

57.13 To configure a diameter peer policy

57.13.1 Steps

1

Choose Policies→AAA Policies→Diameter Peer from the NFM-P main menu. The Diameter Peer Policies form opens.

2

Click Create or select an existing diameter peer policy and click Properties. The Diameter Peer Policy (Create|Edit) form opens.

3

Configure the required parameters.

On the Connection Origin panel, set the Role parameter to Proxy if you intend to configure proxy entries on local definitions of the policy after it is distributed to NEs.

4

Select a Python policy.

5

Click Apply. The Diameter Peer Policy form refreshes with additional tabs.

6

Click on The Peers tab to configure diameter peers.

1. Click Create or select an existing diameter peer object and click Properties. The Diameter Peer (Create|Edit) form opens.
2. Configure the required parameters.
3. If you want to specify local override settings for any of the timer settings for the peer, deselect the Inherit Value check box and then specify a value for the corresponding timer parameter.

   When the Inherit Value check box is selected for a timer parameter on a peer object, the timer value is inherited from the parent diameter policy.

4. Save your changes and close the form.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

End of steps

57.14 To configure a proxy site on a local diameter peer policy

57.14.1 Purpose

The proxy configuration is only available on local definitions of the diameter peer policy if the parent (global) policy is configured with the connection origin role set to the Proxy option. See 57.13 “To configure a diameter peer policy” (p. 1804).

57.14.2 Steps

1

Choose Policies→AAA Policies→Diameter Peer from the NFM-P main menu. The Diameter Peer Policies form opens.

2

Select a diameter peer policy and click Properties. The Diameter Peer Policy (Edit) form opens.

3

Click on the Local Definitions tab.

4

Select a local definition item and click Properties. The Diameter Peer Policy, Local Policy (Edit) form opens.

5

Click on the Proxy tab.

6

Configure the Administrative State parameter and select a routing instance.

7

Click on the MC Peer tab to configure MC peer entries.

1. Click Create or select an existing peer entry and click Properties. The Multi Chassis Diameter Proxy Entry (Create|Edit) form opens.
2. Select an MC peer.

3. Configure the Peer Sync Tag parameter.
4. Save your changes and close the form.

8

Save your changes and close the forms.

END OF STEPS

57.15 To force a route download on an NE

57.15.1 Purpose

Use this procedure to force an immediate route download on an individual NE. In order to complete this procedure, you must have a route download policy configured (see 57.10 “To configure a route download policy” (p. 1801).

57.15.2 Steps

1

Choose Policies→AAA Policies→Route Download from the NFM-P main menu. The Route Download Policies form opens.

2

Select the route download policy that is configured on the NE for which you want to force a route download and click Properties. The Route Download Policy (Edit) form opens.

3

Click on the Local Definitions tab and click Search. A list of NEs associated with the route download policy is displayed.

4

Select an NE and click Properties. The Route Download Policy - Local Policy (Edit) form opens.

5

Click Force Download. The NE sends a route request to the RADIUS server.

6

Close the form.

END OF STEPS

58 Python policies

58.1 Python policies and Python script policies

58.1.1 Overview

A Python policy references one or more Python script policies in the NFM-P, and associates each with a message (RADIUS or other). Each message entry is configured to act on either ingress or egress traffic.

A Python script policy is configured on a Python policy. A Python script policy is used to specify the location of a Python script. The Python script policy specifies three possible URL locations for the script (for example, a local CF card or a remote FTP server). The system picks the first URL that establishes a valid connection. The Python script policy also specifies the method used to ensure the integrity and/or the confidentiality of the content of a Python script, and the action taken when the defined RADIUS message type fails.

A Python script policy can be configured with a script protection key that is shared with the script protection configuration on an NE. If Python script protection is configured, the NE is configured with the same protection key, along with the source URL of the Python script, and the destination URL where the protected script is stored.

A Python policy message entry can be used, for example, to modify RADIUS messages of a RADIUS proxy server, change the RADIUS attributes of the different RADIUS messages, or to process DHCP messages for an interface. A Python policy message entry of the type Syslog can be used to customize the formatting of syslog messages carrying NAT information. Python policies can be configured on the following object types:

- RADIUS server or RADIUS proxy server on a base routing instance (27.8 “To configure a RADIUS server on a routing instance” (p. 846) and 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847))
- RADIUS server policy (57.6 “To configure a RADIUS server policy” (p. 1797))
- IES group interface (78.19 “To configure a group interface on an IES” (p. 2449))
- VPRN group interface (79.37 “To configure a group interface on a VPRN” (p. 2586))

A Python policy can be configured with a cache in which script message strings are saved. Other scripts can be configured to retrieve the strings. This functionality is not limited to a particular service type. For example, a RADIUS Python script could save an SLA profile name provided in RADIUS Access-Accept message to the cache, and a DHCP Discovery script could retrieve it. The Python cache is configured with a maximum number of entries, and a maximum size for each entry. The lifespan of cache entries is also configurable.

Each local definition of a distributed (global) Python policy can be configured with a Python cache peer object, which specifies the IP address of an MC redundancy synchronization peer and its associated synchronization tag.

A Python policy can be associated with an ISA-WLAN GW group or and ISA-NAT group, in which case the policy acts as an ISA Python policy. This means that all Python scripts associated with the Python policy are loaded into ISA memory, depending on the state of the Python script (shutdown or

not shutdown), and its association with a Python policy.

58.2 To configure a Python script policy

58.2.1 Steps

1

Choose Policies→Python→Python Script from the NFM-P main menu. The Python Scripts form opens.

2

CAUTION

Service Disruption

Modifying an active Python script is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the Python script before you proceed.

Click Create or select an existing entry and click Properties. The Python Script form opens.

3

Configure the required parameters.

If the Protection parameter is set to HMAC-SHA-256, you must specify a protection key. The protection key value must match the protection key value specified on the NE to which the Python script policy is applied. The two values are not validated against each other until script protection is enabled on the NE.

4

Save your changes and close the form.

END OF STEPS

58.3 To configure a Python policy

58.3.1 Steps

1

Choose Policies→Python→Python Policy from the NFM-P main menu. The Python Policies form opens.

2

CAUTION

Service Disruption

Modifying an active Python script is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the Python policy before you proceed.

Click Create or select an existing entry and click Properties. The Python Policy form opens.

3

Configure the required parameters.

The ISA -WLAN GW group and ISA-NAT group parameters can only be configured at the time of policy creation. You can configure an ISA -WLAN GW group OR an ISA-NAT group on the policy, but not both. If the policy is configured with an ISA -WLAN GW group or an ISA-NAT group, you cannot configure a Python cache or message entries. Save your changes and close the form.

4

Click on the Cache tab to configure a Python cache.

1. Click Create or select the existing entry and click Properties. The Python Policy Cache form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click on the Message Entries tab to configure message entries.

1. Click Create or select an existing entry and click Properties. The Python Policy Message form opens.
2. Configure the required parameters.
3. Select a Python script policy.

6

Click OK to confirm and close the forms, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

58.4 To configure a Python policy cache peer

58.4.1 Purpose

A cache peer can only be configured on a local definition of a global Python policy that has been distributed to NEs.

58.4.2 Steps

1

Choose Policies→Python→Python Policy from the NFM-P main menu. The Python Policies form opens.

2

CAUTION

Service Disruption

Modifying an active Python script is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the Python policy before you proceed.

Select an entry and click Properties. The Python Policy form opens.

3

Click on the Local Definitions tab.

4

Select an entry and click Properties. The Python Policy, Local Policy (Edit) form opens.

5

Click on the Peer tab.

1. Click Create or select an existing entry and click Properties. The Python Policy Cache Peer form opens.
2. Select a peer address.
3. Configure the Sync Tag parameter.

6

Save your changes and close the forms.

END OF STEPS

59 802.1x policies

59.1 802.1x policies function

59.1.1 Overview

The NFM-P implementation of the 802.1X protocol provides 802.1X authentication on an individual port basis.

In an 802.1X environment, a user, called the supplicant, requests access to an access point, called the authenticator. The authenticator forces the supplicant into an unauthorized state, forcing them to send an EAP start message.

The authenticator returns an EAP message to request the user identity. The user returns the identity, which is forwarded by the authenticator to the authentication server. The server authenticates the user and returns an accept or reject message to the authenticator.

If an accept message is received, the authenticator changes the user state to authorized and user traffic is processed.

59.2 To configure an 802.1x policy

Note: Before you can assign an 802.1x policy to a port, 802.1x must be enabled on the device. See 12.17 “To enable or disable 802.1X” (p. 353) for information about enabling 802.1x.

59.2.1 Steps

1

Choose Policies→Ethernet→802_1x from the NFM-P main menu. The Manage 802_1x Policies form opens.

2

Click Create. The 802_1x Policy (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Click the RADIUS Servers tab, if required. Otherwise, go to Step 8.

5

Click Create. The RADIUS Server (Create) form opens.

6

Configure the required parameters.

7

Click OK to save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

9

Close the form.

END OF STEPS

60 PBB MRP policies

60.1 PBB MRP policies

60.1.1 Overview

PBB MRP policies limit the scope of MMRP advertisements to a specific network domain using ISID-based filters for both the MMRP control plane and the B-VPLS data plane. A PBB MRP policy contains a configurable option to instantiate an MMRP tree and related entry only when both an MMRP declaration and registration are received on a port.

A PBB MRP policy specifies either a forward or a drop action for the Group B-MAC attributes associated with the ISIDs that are specified in the match criteria. If the scope of the policy is set to template, you can apply the PBB MRP policy to multiple B-VPLS services.

See 77.2.4 “Provider Backbone Bridging in VPLS” (p. 2215) and “MRP and MMRP support” (p. 2218) in Chapter 77, “VPLS management ” for more information.

60.1.2 Workflow to manage PBB MRP policies

1

Create a PBB MRP policy. See 60.2 “To configure a PBB MRP policy” (p. 1813) .

2

Release and distribute the PBB MRP policy to NEs. See 60.2 “To configure a PBB MRP policy” (p. 1813) .

3

Configure MRP forwarding control on a B-site. See 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) .

4

Configure MRP forwarding control and apply the PBB MRP policy to a B-VPLS L2 access interface, or spoke or mesh SDP binding. See 77.105 “To configure learning protection parameters on a VPLS SDP binding” (p. 2403) and 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) .

60.2 To configure a PBB MRP policy

60.2.1 Steps

1

Choose Policies→Ethernet→PBB MRP from the NFM-P main menu. The PBB MRP Policies form opens.

2

Click Create. The PBB MRP Policy (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Click the Entries tab.

5

Click Create. The PBB MRP Policy Entry (Create) form opens with the General tab displayed.

6

Configure the required parameters.

7

Click the ISID Match Criteria tab.

8

Click Create to specify ISID matching criteria for an entry. Multiple ISID matching criteria can be specified for each entry. The ISID Match Criteria (Create) form opens.

9

Configure the required parameters and click OK.

10

Repeat Step 8 and Step 9 to add another set of criteria. Otherwise, go to Step 11 .

11

Click OK. The PBB MRP Policy Entry (Create) form closes.

12

Repeat Step 5 to Step 11 to add another entry. Otherwise, go to Step 13 .

13

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

14

Close the form.

END OF STEPS

61 AOS Ethernet Service policies

61.1 AOS Ethernet Service policy types

61.1.1 Overview

This chapter describes how to configure the UNI, SAP, UNP, and custom L2 profiles associated with OmniSwitch Ethernet services.

A UNI profile is associated with a UNI port and configures how a variety of protocol control packets are processed on the UNI port. Processing actions include Tunnel, Drop, Peer, and MAC Tunnel.

A SAP is associated with a stacked VLAN service. A SAP profile defines values for ingress bandwidth sharing, rate limiting, customer VLAN tag processing (translate or preserve), and priority mapping (inner to outer tag or fixed value).

A custom L2 profile is associated with a UNI profile for specific packet control and can be applied to specific actions, such as tunnel, MAC-tunnel, and discard.

61.2 To configure an OmniSwitch Ethernet service UNI profile

61.2.1 Steps

1

Choose Policies→Ethernet→AOS Ethernet Service from the NFM-P main menu. The Ethernet Service Policies form opens.

2

Click Create and choose Create UNI Profile. The Ethernet Service UNI Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

As required, bind the UNI profile to a custom L2 profile.

1. Click on the Custom L2 Profile Binding tab.
2. Select a custom L2 profile from the list and click OK.

   Note:

   If no custom L2 profiles exist, click Create. See 61.4 “To configure an OmniSwitch Ethernet service custom L2 profile” (p. 1816) for more information.

5

Click OK. The Ethernet Service UNI Profile (Create) form closes.

6

Close the Ethernet Service Policies form.

END OF STEPS

61.3 To configure an OmniSwitch Ethernet SAP profile

61.3.1 Steps

1

Choose Policies→Ethernet→AOS Ethernet Service from the NFM-P main menu. The Ethernet Service Policies form opens.

2

Click Create and choose Create SAP Profile. The Ethernet Service SAP Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Click OK. The Ethernet Service SAP Profile (Create) form closes.

5

Close the Ethernet Service Policies form.

END OF STEPS

61.4 To configure an OmniSwitch Ethernet service custom L2 profile

61.4.1 Before you begin

Support for OmniSwitch custom L2 profiles varies depending on the NE family, variant, release, and license; see the NE documentation for more information.

61.4.2 Steps

1

Choose Policies→Ethernet→AOS Ethernet Service from the NFM-P main menu. The Ethernet Service Policies form opens.

2

Click Create and choose Create Custom L2 Profile. The Custom L2 Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

The Protocol Mac parameter must be a multicast MAC address.

4

If you set the Protocol Entry Type parameter to a value other than Mac Type (the default value), then addition parameters are displayed. Configure the required parameters.

5

Click OK. The Custom L2 Profile (Create) form closes.

6

Close the Ethernet Service Policies form.

END OF STEPS

61.5 To clear custom L2 profile statistics

61.5.1 Steps

1

Right-click on the required OmniSwitch NE in the navigation tree and choose Properties from the contextual menu. The properties form for the NE opens with the General tab displayed.

2

Click on the Globals tab.

3

Click on the L2PT Clear Stats tab.

4

Configure the required parameters.

5

Click OK. The properties form for the NE closes.

END OF STEPS

61.6 To configure an OmniSwitch Ethernet UNP profile

61.6.1 Steps

1

Choose Policies→Ethernet→AOS UNP Profile from the NFM-P main menu. The AOS UNP Profile Policy form opens.

2

Click Create. The AOS UNP Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

62 VLAN Connection Profile policies

62.1 Overview

62.1.1 NE support

The VLAN Connection Profile policy is supported on 7450 ESS, 7750 SR and 7950 XRS NEs.

62.1.2 Description

The Ethernet policy VLAN Connection Profile and TPS Connection Profile enable you to configure SAPs to be part of multiple VLANs in a single command. You can specify all the VLANs that the SAPs needs to be part of by providing a range of VLAN IDs or separate, individual, VLAN IDs. When this policy is used under a SAP, the NFM-P distributes the connection profile to the NE first and then references the connection profile.

After creating a policy, you can use it in Epipe (terminating sites only) and VPLS (regular sites only) services. The VLAN Connection Profile policies are not supported in the following services: Etree, M-VPLS, B-VPLS, R-VPLS, I-VPLS and PBB-Epipe.

The NFM-P performs the validation of VLAN ID ranges to ensure that different connections profiles do not overlap.

62.1.3 SAP encapsulation

Only access and hybrid ports with the encapsulation type of Dot1 Q and Q-in-Q support VLAN Connection Profile and TPS Connection Profile policies. You need to be aware of the SAP encapsulation requirements and dependencies for the policy when using the policy in an Epipe or VPLS service.

The VLAN Connection Profile policy supports the following new SAP types:

- Dot1 Q encapsulation port
  - sap 1/1/1:cp-1

    A connection profile policy is required for the creation of this SAP type.
- Q-in-Q encapsulation port
  - sap 1/1/2:cp-2.*

    A connection profile policy is required for the creation of this SAP type, the VLAN Connection Profile tag needs to be Outer, and the inner encapsulation value needs to be set to 0, 4095, or * (asterisk).
  - sap 1/1/3:1234.cp-3

    A connection profile policy is required for the creation of this SAP type, the VLAN Connection Profile tag needs to be Inner, and the outer encapsulation value needs to be set to a value from 1 through 4094.

The assignment of SAPs with overlapping encapsulation types causes operation failure.

62.1.4 Constraints

This section describes the constraints that must be considered when configuring SAPs for VLAN connection profiles.

Services

VLAN connection profiles are not supported in the following types of services:

- Etree
- M-VPLS
- B-VPLS
- R-VPLS
- I-VPLS
- PBB-Epipe

Features

The following features are not supported in combination with VLAN connection profile services:

- Proxy ARP and Proxy ND
- Capture SAPs
- Ethernet tunnel SAPs
- Ethernet ring SAPs

  The VLAN connection profile SAPs can be used as Ethernet ring data SAPs, but do not support control G.8032 traffic.
- VLAN translation
- Xstp

  The VLAN connection profile SAPs can be managed by an M-VPLS, but services with CP SAPs do not support Xstp.
- L2PT
- BPDU translation
- Subscriber management features
- IGMP/MLD/PIM (v4 or v6)
- VLAN VC Tag under an SDP-binding sharing service with a VLAN connection profile SAP

62.2 To configure a VLAN Connection Profile policy

62.2.1 Steps

1

Choose Policies→Ethernet→VLAN Connection Profile or TPS Connection Profile from the NFM-P main menu. The Manage VLAN Connection Profile form opens.

2

Click Create. The Connection Profile (Create) form opens.

3

Configure the required parameters on the General tab..

4

Configure a VLAN Connection Profile range.

1. Click on the VLAN Connection Profile Range tab, and click Create. The VLAN Connection Profile Range, VLAN Connection Profile (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Save your changes and close the VLAN Connection Profile (Create) form.

6

Close the Manage VLAN Connection Profile form.

END OF STEPS

62.3 To view SAPs associated with a specific VLAN connection profile

62.3.1 Steps

1

Choose Policies→Ethernet→VLAN Connection Profile or TPS Connection Profile from the NFM-P main menu. The Manage Connection Profile form opens.

2

Perform a search to find the configured connection profile. You can search for all the configured VLAN connection profiles or search using the ID of a specific VLAN connection profile.

3

Select the VLAN connection profile and click Properties. The VLAN Connection Profile form opens.

4

Click on the L2AccessInterface tab and then on the VLL L2Access Interface or VPLS L2Access Interface tab.

5

Click Search. The SAPs associated with the VLAN connection profile are displayed.

END OF STEPS

63 Connection profile policies

63.1 Connection profile policies

63.1.1 Overview

You can configure an ATM-cell Apipe service with an ATM SAP type. The SAP type is an ATM connection profile name. The NFM-P ATM connection profile configures a list of VPI/VCI pairs.

You can apply a connection profile only to a SAP that is part of an Apipe VLL with a VC Type parameter that is set to ATM-cell. The ATM SAP can be on a regular port or APS port. You can apply a connection profile to any number of ATM SAPs.

Consider the following when you configure a connection profile:

- A SAP-to-SAP VLL service is not supported using ATM SAP to which a connection profile is assigned. You must configure each VPI/VCI pair into a separate SAP and create as many Apipe VLL services with the VC type ATM-vcc as required.
- You cannot configured an ATM SAP to which a connection profile is assigned on a port which is part of a MC-APS protection group.
- You cannot add an ATM SAP to which a connection profile to an Apipe VLL endpoint. The SAP cannot be configured on a MC-LAG or MC-APS port. However, this ATM SAP can be part of a VLL service which has an endpoint on a spoke SDP.

You can configure connection profile policies on the 4-port OC-3/STM-1:OC-12/STM-4 ATM MDA and on the 16-port OC-3/STM-1 ATM MDA, on the IOM3/IMM on the 7750 SR-1, 7750 SR-7, 7750 SR-12, 7750 SR-c4, and the 7750 SR-c4.

63.1.2 7210 SAS VLAN ranges

7210 SAS VLAN ranges allow you to group a range of VLAN IDs as a single entity. This allows you to provide a VPLS or VLL Epipe service configuration—such as forwarding, ACL, QoS, and accounting—to the group of VLAN IDs.

You can use a connection profile to specify either a range of VLAN IDs or individual VLANs to be grouped together in a single SAP. You can configure only one VLAN range SAP in a VPLS or Epipe service. In an Epipe service, the other endpoint must be a Q.* SAP or a spoke SDP binding. Access SAPs that use VLAN range values are allowed only for a Dot1Q-encapsulated port or LAG. You can configure multiple connection profiles for each port or LAG, as long as the VLAN values that are specified in each profile do not overlap.

For VPLS VLAN range SAPs, the following restrictions apply:

- BPDU translation, IGMP snooping, and MVR are not supported
- IPv4 and Ipv6 egress ACL filters are not supported
- You cannot configure the Maximum Entries parameter for the FIB

- The Learning Enabled parameter for the SAP cannot be set to false

VLAN ranges for VLL Epipe services are supported on the following 7210 SAS NEs:

- 7210 SAS-D
- 7210 SAS-K
- 7210 SAS-M
- 7210 SAS-Mxp
- 7210 SAS-R
- 7210 SAS-S
- 7210 SAS-Sx
- 7210 SAS-T
- 7210 SAS-X

VLAN ranges for VPLS are supported on the following 7210 SAS NEs:

- 7210 SAS-D
- 7210 SAS-K
- 7210 SAS-M
- 7210 SAS-T

63.2 Workflow to manage connection profiles

63.2.1 Purpose

The following workflow lists the high-level steps required to manage ATM connection profiles.

63.2.2 Stages

1

Create an ATM connection profile. See 63.3 “To configure a connection profile policy” (p. 1825).

2

Release and distribute the ATM connection profile to NEs. See 63.3 “To configure a connection profile policy” (p. 1825).

3

Apply the connection profile to an Apipe VLL L2 access interface. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174).

63.3 To configure a connection profile policy

63.3.1 Steps

1

Choose Policies→ATM→Connection Profile from the NFM-P main menu. The Manage Connection Profile form opens.

2

Click Create. The Connection Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Perform the following to add ATM members to the connection profile:

1. Click on the ATM Member tab.
2. Click Create. The ATM Member (Create) form opens.
3. Configure the required parameters and click OK. The Connection Profile (Edit) form refreshes with the VPI and VCI information displayed.
4. Repeat 2 and 3 to add additional members. You can add up to 16 members.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Close the forms.

End of steps

63.4 To configure a VLAN range for a 7210 SAS VPLS or VLL Epipe service

63.4.1 Purpose

Perform this procedure to create a connection profile to configure a range of VLAN values that you can assign to an Dot1q SAP in a VPLS or VLL Epipe service.

Note: You cannot modify an existing connection profile that is associated with a SAP. If you need to modify the connection profile, you must first remove the profile from all SAPs with which it is associated.

63.4.2 Steps

1

Choose Policies→ATM→Connection Profile from the NFM-P main menu. The Manage Connection Profile form opens.

2

Click Create. The Connection Profile (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Specify the list of VLAN ranges or individual VLAN IDs to be used for mapping the VLANs to the VPLS or VLL Epipe SAP.

1. Click on the VLAN Ranges tab.
2. Click Create. The VLAN Range form opens.
3. Configure the Range parameter and click OK. The VLAN Range form closes.
4. Repeat 2 and 3 to create additional ranges.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

6

Close the forms.

END OF STEPS

64 Residential subscriber policies

64.1 Primary residential subscriber policy components

64.1.1 Overview

This section describes the major policy components that are essential components in a residential subscriber configuration.

These required policies include:

- Subscriber identification policies
- Subscriber profiles
- SLA profiles
- Subscriber explicit maps

See Chapter 74, “Residential subscriber management” for information about configuring and managing residential subscribers.

64.1.2 Subscriber identification policies

Subscriber identification policies apply to SAPs and associate dynamic residential subscriber hosts with NE subscriber instances. Hosts that use the same subscriber identifier, called a subscriber identification string, belong to the same subscriber and receive common general HQoS and accounting treatment as defined by the customer SLA specifications in a subscriber profile. Hosts from multiple SAPs can belong to the same subscriber, but for HQoS scheduling purposes, all hosts of a subscriber must be active on the same IOM.

A subscriber identification policy includes:

- URLs for subscriber identification scripts
- subscriber profile map
- SLA profile map

Entries in the subscriber and SLA profile maps are optional and are used for the direct assignment of profiles to hosts when hosts include profile ID strings in the DHCP Option 82 field.

Default subscriber identification policy

Although the NFM-P is installed with no default subscriber identification policy, the user can create a subscriber identification policy and designate it as the default policy by giving it the case-sensitive name “default”.

Subscriber identification scripts

A URL in a subscriber identification policy points to the location of a subscriber identification script that an NE uses to parse DHCP option information. During initialization, an NE downloads scripts

from the URLs that are specified in the applied subscriber identification policies. An NE can store a maximum of four subscriber identification policies; each policy can contain up to three scripts.

A subscriber identification script derives the mandatory subscriber ID string, as well as a DSLAM identifier and optional profile ID strings, from the DHCP option information.

The user assigns a priority to each script in a subscriber identification policy. Only the operationally enabled script with the highest priority is active. If an NE encounters an error in a script or cannot find a script at the specified location, the NE marks the URL as operationally down, raises an alarm, and attempts to use the script that is next in priority.

The script-related alarms that the NFM-P raises against the local instance of a subscriber identification policy are:

- warning alarm when the primary URL is operationally up but a lower-priority URL is operationally down
- major alarm when the primary script is operationally down but one of the other scripts is operationally up
- critical alarm when all scripts are operationally down

A modification to a subscriber identification script or URL takes effect only after the URL is administratively disabled and then re-enabled, which causes an NE to reload the script. Replacing or modifying a subscriber identification script or a URL can be service-affecting if not done properly. To avoid a service disruption when you modify a subscriber identification script or URL, perform 74.19 “To modify the primary subscriber identification script or URL” (p. 2040) .

64.1.3 Subscriber profiles

A subscriber profile defines the aggregate HQoS and accounting characteristics for the hosts of a specific subscriber. During the creation of a subscriber profile, the user chooses ingress and egress scheduler policies that apply to all host queues of the subscriber. A subscriber profile permits the optional selection of SLA profiles to override the SLA profiles that are named in the subscriber identification policy.

A subscriber profile can be specified in the following components, which an NE searches in the order shown when it attempts to assign a subscriber profile to a host:

- subscriber explicit map
- subscriber identification policy

A subscriber requires an association with one and only one subscriber profile. If no subscriber profile is associated with a subscriber ID string or explicitly specified by a host, and there is no available default subscriber profile on an NE or on a SAP, the NE rejects the host.

When an accounting policy is associated with the subscriber profile, the user can configure the type of subscriber host accounting data that is reported in the RADIUS accounting message, and in the XML accounting records.

The user can define a default subscriber profile for an NE or a SAP by giving it the case-sensitive name “default” during profile creation. NE and SAP default subscriber profiles apply to hosts for which a subscriber profile is unspecified in the subscriber explicit map and subscriber identification policy.

64.1.4 SLA profiles

An SLA profile defines the resources that an NE assigns to a particular subset of subscriber hosts, such as VoIP telephones or BTV set-top boxes. These resources include network-access and ACL policies. An SLA profile also optionally defines the maximum number of hosts that use the profile and the action taken when the number of hosts reaches the maximum.

An SLA profile can be specified in the following components, which an NE searches in the order shown when it attempts to assign an SLA profile to a host:

- subscriber explicit map
- subscriber profile
- subscriber identification policy

An SLA profile includes:

- access ingress policy association
- access egress policy association
- local scheduler policy association for access egress policy
- ingress IP ACL filter policy association
- egress IP ACL filter policy association
- host limit that specifies the maximum number of hosts that use the SLA profile
- access ingress and access egress policy parameter override values
- access ingress and access egress policer override values
- egress scheduler override values

The queues in the access ingress and egress policies of an SLA profile must use a scheduler from the scheduler policy in the subscriber profile as their parent scheduler. Additionally, a local scheduler can be specified under the access egress policy configuration. The local scheduler policy QoS settings override the QoS settings of the parent scheduler specified in the subscriber profile.

All hosts that use the SLA profile must be active on the same IOM. When an SLA profile does not name an access ingress or access egress policy, an NE uses the SAP default policy.

The user can define a default SLA profile for a SAP. A default SAP SLA profile applies only to hosts for which an SLA profile is unspecified in the subscriber explicit map, subscriber profile, and subscriber identification policy.

64.1.5 Subscriber explicit maps

A subscriber explicit map is a table that directly associates dynamic subscriber hosts with subscriber profiles and SLA profiles. The entries in a subscriber explicit map override the default profile definitions. A subscriber ID string is the unique key for a subscriber explicit map entry.

A subscriber explicit map includes:

- subscriber ID string
- subscriber profile name
- SLA profile name

- AA group policy name
- subscriber ID alias

A subscriber explicit map does not allow the association of a subscriber ID string with the subscriber profile called “default”. However, if the explicit mapping omits a subscriber profile, the subscriber ID string is associated with the SAP or NE default subscriber profile. An attempt to delete a subscriber profile that is named in a subscriber explicit map fails. The user can remove explicit map entries at any time.

If the user creates a subscriber explicit map entry without using the NFM-P, for example, using CLI, the NFM-P creates a subscriber instance in the global subscriber explicit map. If the NFM-P subsequently discovers another NE with the same subscriber entry that has a mapping to different profile, the NFM-P treats the second mapping as a local mismatch to the global entry for the subscriber.

64.2 Secondary residential subscriber policy components

64.2.1 Overview

This section describes additional and optional policy components that can be applied to a residential subscriber configuration, depending on the specific application and service delivery model:

- ANCP policies
- PPPoE policies
- PPP policies
- MSAP policies
- Host tracking policies
- Category map policies
- Credit control policies
- IGMP policies
- IPoE session policies
- BGP peering policies
- Diameter policies
- Subscriber multicast CAC policies
- Mobile gateway/peer profiles
- Host lockout policies
- RADIUS script policies
- HTTP redirect policies
- MLD policies
- Diameter application policies
- Distributed subscriber management traffic policers
- Distributed subscriber management IP filter policies
- PIM policies
- SHCV policies
- RIP policies
- BRG profiles
- UPnP policies
- APN policies
- MAP-T domain policies

64.2.2 ANCP policies

An ANCP policy defines the behavior of the residential subscriber host with which the policy is associated.

An ANCP policy includes one of the following:

- static ANCP association
- static ANCP MSS association

- dynamic subscriber-profile association on VPLS, IES, VPRN and VLL services

An ANCP policy conveys status and control information based on port-up/port-down messages and changes to the current access line rate between the edge device and the access node. This allows supporting NEs to adjust the HQoS subscriber scheduler with the correct rate or raises an alarm when the rate goes below a set threshold. The policy can be changed if the rate drops below a minimal threshold value. The ANCP actual upstream synchronization rate is configured in the ingress panel while the ANCP actual downstream synchronization rate is configured in the egress panel.

64.2.3 PPPoE policies

PPPoE is used in subscriber networks to encapsulate PPP frames inside Ethernet frames. PPPoE combines the point-to-point protocol used with DSL sessions with the Ethernet protocol used to support multiple subscribers in a local area network. PPPoE takes advantage of the speed of a packet-based Ethernet network with the security and accounting functions of a PPP network. PPPoE allows service providers to use existing RADIUS authentication.

Since more than one subscriber is sharing the same connection to a service provider, PPPoE organizes subscribers during two stages:

- PPPoE discovery stage
- PPPoE session stage

During the discovery stage, the subscriber and service provider identify each other's MAC address and establish a PPPoE session ID.

The PPPoE discovery stage consists of the following steps:

1. PPPoE Active Discovery Initiation (PADI). The client initiates a session by broadcasting a PADI packet to the LAN to request a service.
2. PPPoE Active Discovery Offer (PADO). Any access device that can deliver the service requested by PADI packet, replies with a PADO packet that contains its name, unicast address and service requested.
3. PPPoE Active Discovery Request (PADR).The client selects one of the PADOs that it receives and sends a PADR packet to indicate the services required.
4. PPPoE Active Discovery Session Confirmation (PADS). When the selected device receives the PADR packet, it can accept or reject the PPPoE session. To accept the session the device sends the client a PADS packet with a unique ID and a service name. If the device rejects the session, it sends a PADS packet with a service name error and resets the session ID to zero.

During the session stage, PPPoE behaves as a peer-to-peer protocol. Each PPPoE session is identified by the MAC address of the peer and the session ID. Once a session is established, both end points build a point to point connection over the Ethernet and exchange packets. Once the connect is established, the RADIUS accounting policy can begin. LCP negotiates authentication parameters. After a session is authenticated IPCP sends an IP address to the PPPoE client. IP Addresses can be stored in the DHCP local user database, if configured.

After a session is established both end points monitor the session and can terminate a session after a configured number of keep alive intervals are exceeded. An alarm is raised by the NFM-P when a PPPoE session fails. Either peer can send a PPPoE Active Discovery Termination (PADT) packet.

See 74.2.4 “Local DHCP server” (p. 1996) in this chapter for more information on associating PPPoE with a local DHCP server.

PPPoA sessions

PPPoA is used in subscriber networks to encapsulate PPP frames inside ATM Adaption Layer 5 packets. Each PPP session is carried over a single ATM VC to the BNG. Two types of encapsulation are supported:

- AAL5_nlpid_PPP
- AAL5_Mux_PPP

PPPoEoA sessions

PPPoEoA is used in subscriber networks to encapsulate PPPoE messaging within ATM Adaption Layer 5 packets. Each PPP session can be directly terminated on any host within the customer Ethernet network and then transported over an ATM network to the BNG. Two types of encapsulation are supported:

- AAL5_muxed_bridged_EHT_no_FCS
- AAL5_snap_bridged

64.2.4 PPP policies

A PPP policy defines PPP parameters for an access interface. A PPP policy is applied to an access interface.

64.2.5 MSAP policies

An MSAP policy defines how parameters are applied in the creation of an MSAP. An MSAP policy is applied to capture SAP L2 access interfaces and MSAP L2 and L3 access interfaces. An MSAP policy can be configured to create sticky MSAPs, which are not automatically cleared by the system when they become idle. Sticky MSAPs must be cleared manually; see 64.11 “To clear idle sticky MSAPs” (p. 1854).

64.2.6 Host tracking policies

A host tracking policy is used to allow a subscriber's video traffic (multicast) to be included in the egress rate control for the subscriber. When a host tracking policy is specified in a subscriber profile, the egress traffic rate for the subscriber takes into account the unicast and multicast traffic in the aggregate egress rate or in the egress scheduler rate specified in the ANCP policy. There is no default host tracking policy.

When a host tracking policy is applied to a subscriber profile, all subscribers associated with the subscriber profile are tracked using that host tracking policy. You can view the tracked subscribers on the associated local host tracking policy property form.

You can also configure IGMP host tracking parameters on VPLS, IES, and VPRN service sites and SAPs.

The following on-demand host tracking information is available:

- on the Residential Subscriber Instance form, Host Tracking Info tab, which hosts are being

tracked for this particular subscriber

- on the Residential Subscriber Instance form, Host Tracking Status tab, the egress traffic rate reduction for this particular subscriber
- on the VPLS L2 access interface property form, Host Tracking Info tab, which hosts are being tracked
- on the IES and VPRN service access point property form, Host Tracking Info tab, which hosts are being tracked

On-demand historical and real-time host tracking statistics are supported for SAPs and residential subscribers. Only historical statistical plotting is supported. The statistics and read-only host tracking information can be cleared by the NFM-P.

64.2.7 Category map policies

A category map policy combines up to sixteen credit category mappings for association with SLA profiles and credit control policies. Each category specifies a group of queues and policers, as well as credit exhausting IP filter entries.

64.2.8 Credit control policies

A credit control policy is used to associate additional policy objects with a credit control server. If the credit control server is a RADIUS server, a category map policy is specified. If the credit control server is a diameter server, a diameter policy, diameter application policy, and category map policy are specified. A credit control policy also specifies out-of-credit and error handling actions. The credit control policy is configured on SLA profiles.

64.2.9 IGMP policies

An IGMP policy specifies IGMP group parameters and a list of static multicast group addresses and static source addresses. The IGMP policy is associated with subscriber profiles.

64.2.10 IPoE session policies

The IPoE session policy allows dual-stack IPoE subscriber sessions to function simlarly to dual-stack PPoE sessions. The policy specifies which session key information in the trigger packet is used to instantiate the IPoE session, and session timeout information. The IPoE session policy is bound to a VPLS capture SAP, and to a VPRN or IES group interface.

64.2.11 Diameter policies

The NFM-P supports diameter policies and associated diameter peers. Diameter policies are used in a subscriber management context to provide a credit control mechanism. The diameter policy establishes a server/peer configuration, as well as Diameter Credit-Control Application (DCCA) support. The NE functions as the credit control client, while the peer acts as the credit control server.

DCCA represents an alternative to RADIUS for providing a mechanism to support pre-paid service model in access networks. Under a DCCA configuration, the credit control server is used to grant service to a subscriber for pre-defined duration. Before a subscriber host is created, the BNG queries the credit control server about the credit for the given subscriber. If credit is granted, the

subscriber host is installed with the appropriate SLA level. If the subscriber has no credit, the credit control server denies access and no subscriber host is created.

The diameter policy is used to specify common diameter protocol parameters, while the diameter peer defines the relationship with an external diameter server. The diameter protocol parameters defined on the diameter policy describe connection and session characteristics, and indicate which AVPs to use in messages.

The diameter policy is configured on the credit control policy.

Peers

The diameter policy defines a set of peers with which to establish diameter sessions. Peers share the configuration of the policy with which they are associated, but can override individual timer parameters inherited from the policy. In addition, each peer defines transport and connection-specific parameters, values for destination-specific AVPs, and a preference value.

The 7750 SR and 7450 ESS provide read-only operational values and statistics for the peer definition of local diameter policies. The NFM-P can provide a snapshot view of operational values and statistics. Operational values include the peer's timer values, the current state of the peer's state machine, and the peer's order among the other peers within the policy.

DCCA

The 7750 SR supporst a DCCA, which can be customized through a set of parameters on the diameter policy in the NFM-P. The parameters allow customization of DCCA error handling, DCCA timer(s), and AVP value definitions.

64.2.12 Subscriber multicast CAC policies

The subscriber multicast CAC policy defines bandwidth constraints for multicast CAC. The subscriber multicast CAC policy is applied to a subscriber profile.

64.2.13 Mobile gateway/peer profiles

The mobile gateway/peer profile configures a signalling interface between the WLAN gateway and the mobile gateway. The mobile gateway/peer profile is applied to a PGW or a GGSN.

64.2.14 Host lockout policies

A host lockout policy implements the ability to hold off a misconfigured or malicious subscriber that is continuously retrying to connect. This conserves resources on the NE and on the RADIUS server.

A host lockout policy can be associated with the following objects:

- VPLS or MVPLS L2 access interface
- IES or VPRN L3 access interface
- IES or VPRN group interface

Host lockouts can be cleared at the policy level, at the NE level, or from the SAP to which they are applied.

64.2.15 RADIUS script policies

See 57.1.7 “RADIUS script policies” (p. 1790).

64.2.16 HTTP redirect policies

The HTTP redirect policy specifies a redirect URL and destination port, as well as a list of forwarding address/port pairs. An HTTP redirect policy is associated with a WLAN GW configuration on an IES or VPRN group interface.

64.2.17 MLD policies

The MLD policy specifies MLD group-related parameters, as well as a list of static multicast group/source address pairs. An MLD policy is configured on a subscriber profile.

64.2.18 Diameter application policies

The NFM-P supports diameter application policies and associated diameter peer policies. The diameter application policy is used to specify common diameter protocol parameters. The diameter protocol parameters defined on the diameter application policy describe connection and session characteristics, and indicate which AVPs to use in messages.

The protocols and AVP types vary, depending on the policy’s application type:

- Gx
- Gy
- NASREQ

The diameter application policy can be configured on local user database IPoE and PPP hosts, and on VPRN and IES group interfaces.

Diameter NASREQ

In network configurations where a Diameter infrastructure is available, Diameter NASREQ can be used an alternative to a local user database and RADIUS for ESM host authentication. If the diameter application policy is configured for NASREQ, it can be configured as a Diameter Authentication Policy on IES and VPRN group interfaces, on local user database IPoE and PPP hosts, and on VPLS capture SAPs.

64.2.19 Distributed subscriber management traffic policers

The distributed subscriber management traffic policer object specifies per-subscriber bandwidth policing on the anchor ISA. Distributed subscriber management traffic policers can be configured in two ways:

- Single-bucket policers monitor bandwidth using a single traffic flow rate (PIR) and burst parameter (MBS)
- Dual-bucket policers monitor bandwidth using dual traffic flow rates (PIR/CIR) and dual burst parameters (MBS/CBS).

Distributed subscriber management traffic policers are available to all the ISAs in a WLAN-GW group. A default policer template for the WLAN-GW group is always created on the ISA when a

WLAN-GW IOM is added to a WLAN-GW group. When a UE is authenticated, RADIUS returns a policer template in a VSA. On an anchor ISA, the ingress and egress policer is instantiated for the subscriber, based on the policer template and overrides for rates/bursts returned from RADIUS.

64.2.20 Distributed subscriber management IP filter policies

The distributed subscriber management IP filter policy installs shared per-SAP filters on SAPs between the anchor ISA and the IOM. The filters map the DSCP on ingress and egress traffic to per-SAP queues.

You can create a redirect filter with a drop action and configure RADIUS to perform a COA to associate this drop action filter with a given UE. The policy filter name is included in the RADIUS VSA.

64.2.21 PIM policies

The PIM policy is configured on a subscriber profile to enable PIM on residential subscriber hosts. For more information, see 78.1.11 “PIM on IES group interfaces” (p. 2425).

64.2.22 SHCV policies

See “SHCV policies” (p. 2000). To create an SHCV policy, see 64.31 “To configure an SHCV policy” (p. 1872).

64.2.23 RIP policies

The RIP policy is used to authenticate an associated host object on a RIP listener interface. A RIP policy can be configured on the following host objects:

- local user database PPP host
- local user database IPoE host
- IES SAP static host
- VPRN SAP static host

To create a RIP policy, see 64.32 “To configure a RIP policy” (p. 1872).

64.2.24 BRG profiles

The Bridged Residential Gateway (BRG) profile is used to apply default attributes to a virtual CPE, including SLA and subscriber profile strings, SHCV timers and hold times, RADIUS server policy, RADIUS proxy servers, and DHCPv4 configuration. A BRG profile can be configured on IES and VPRN group interfaces.

To create a BRG profile, see 64.35 “To configure a BRG profile” (p. 1875).

64.2.25 Trace profiles

The Trace profile allows debug information for IPoE sessions to be collected as a trace. It correlates debug information across multiple applications into a single trace, allowing you to easily find the root cause of an issue.

Use a trace profile to gather debug information for the following applications:

- Local user database
- RADIUS accounting
- RADIUS authentication
- Connectivity Management
- Python
- MSAP

To create a trace profile, see .

Workflow to configure call trace debugging

Complete the following procedures to deploy call trace debugging in a network.

1. Configure a trace profile; see 64.36 “To configure a trace profile” (p. 1876).

   Specify the applications for which you want to gather debug information, and specify the type of output for gathered information. Deploy the profile to NEs you want to trace.
2. Specify trace log file storage information on traced NEs; see 12.35 “To configure call-trace debug storage on an NE” (p. 369).

64.2.26 UPnP policies

The universal plug and play (UPnP) policy is configured on a subscriber profile to add UPnP port mapping functionality on residential subscriber hosts. UPnP provides user-friendly, automatic creation/deletion of port mappings for non-technical users.

UPnP performs NAT port mapping management by application (P2P, gaming, etc.). In a typical routed home, a CPE router (with NAT) acts as a UPnP IGD, accepting UPnP requests from home devices and applications to add or remove NAT port mappings as required.

64.2.27 MAP-T domain policies

The MAP-T domain policy defines a series of MAP rules to associate with a MAP domain. The MAP domain is configured as part of a NAT configuration on a VPRN service site.

Mapping Address and Port with Translation (MAP-T) functionality is used to connect a private IPv4 network and associated customer equipment (CE), over an IPv6 network, to a public IPv4 network connected to a centralized border router. MAP-T can also be used for direct communication between CEs.

Figure 64-1 MAP-T domain

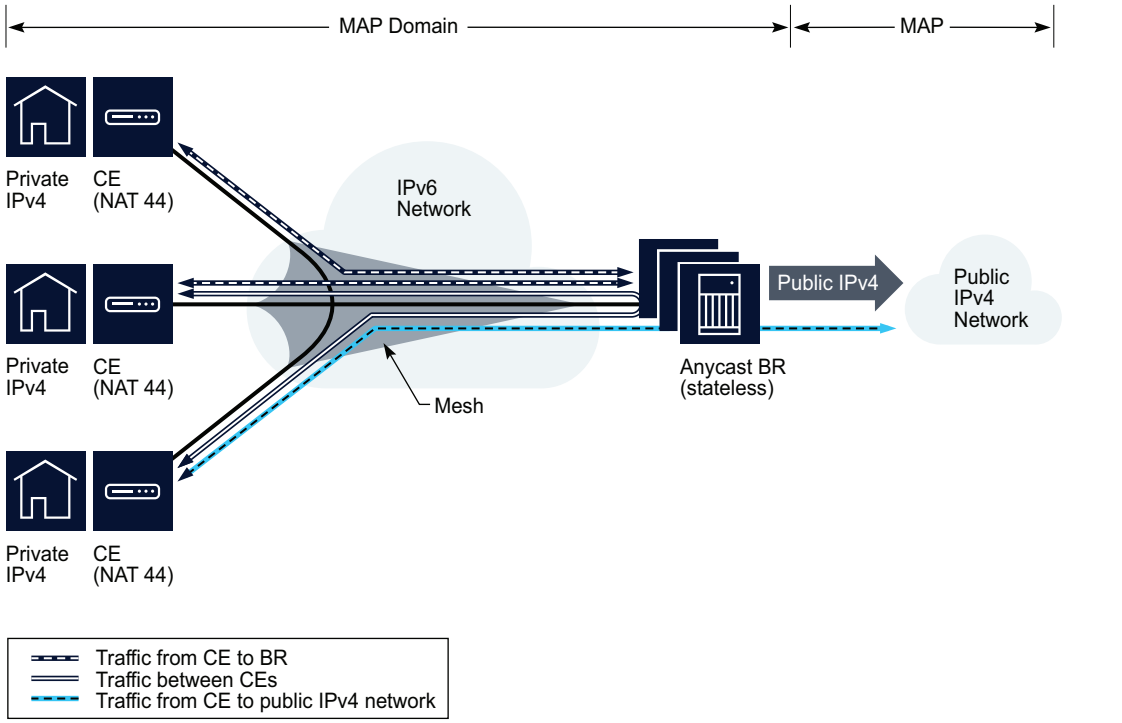

64.2.28 APN policies

See 74.2.15 “HSQ” (p. 2006).

64.2.29 Service chaining policies

For information on service chaining, see 74.2.18 “ISA service chaining” (p. 2007).

For information on configuring service chaining policies, see 64.33 “To configure a service chaining EVPN policy” (p. 1873) and 64.34 “To configure a service chaining VAS filter policy” (p. 1874).

64.3 To configure a subscriber identification policy

64.3.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

CAUTION

Service Disruption

Modifying an active subscriber identification policy is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the subscriber identification policy before you proceed.

Click Create→Subscriber Identification Policy or choose an existing subscriber identification policy and click Properties. The Subscriber Identification Policy (Create|Edit) form opens.

3

Configure the required parameters.

You can designate a subscriber identification policy as the default policy by specifying the case-sensitive name “default” for the Displayed Name parameter. The Displayed Name parameter is configurable only during subscriber identification policy creation.

4

Click on the Supported Service Models tab and specify the intended service model(s) for the subscriber identification policy.

5

Click on the SLA Profiles tab to configure SLA profiles for the subscriber identification policy.

1. On the General tab, configure the Use Direct Map as Default parameter.
2. Click on the Profiles tab.
3. Click Create or choose an SLA profile entry and click Properties. The SLA Profile Entry (Create|Edit) form opens.
4. Configure the SLA Profile String parameter.
5. Select an SLA profile to associate with the subscriber identification policy.
6. Save your changes and close the form.

6

Click on the Subscriber Profiles tab to configure subscriber profiles for the subscriber identification policy.

1. On the General tab, configure the Use Direct Map as Default parameter.
2. Click on the Profiles tab.
3. Click Create or choose a subscriber profile entry and click Properties. The Subscriber Profile Entry (Create|Edit) form opens.
4. Configure the Subscriber Profile String parameter.
5. Select a subscriber profile to associate with the subscriber identification policy.

6. Save your changes and close the form.

7

Click on the Application Profiles tab to configure application profiles for the subscriber identification policy.

1. On the General tab, configure the Use Direct Map as Default parameter.
2. Click on the Profiles tab.
3. Click Create or choose an application profile entry and click Properties. The App Profile Entry (Create|Edit) form opens.
4. Configure the required parameters.
5. Select an application profile to associate with the subscriber identification policy.

   Note:

   If the subscriber identification policy is a global policy, global AA group policies are listed. If the subscriber identification policy is a local policy, local AA group policies are listed.

   Global AA group policies must be manually distributed to the NE before a global subscriber identification policy using the application profiles can be distributed.

   When a new subscriber identification or subscriber explicit map is discovered from the NE, the AA group policy is not automatically resynchronized to the global subscriber identification policy.
6. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs

END OF STEPS

64.4 To configure a subscriber profile

64.4.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

CAUTION

Service Disruption

Modifying an active subscriber profile is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the subscriber profile before you proceed.

Click Create→Subscriber Profile or choose an existing Subscriber Profile and click Properties. The Subscriber Profile (Create|Edit) form opens.

3

Configure the required parameters.

You can designate a subscriber profile as the default profile by specifying the case-sensitive name “default” for the Displayed Name parameter. The Displayed Name parameter is configurable only during subscriber profile creation.

The Egress Aggregate Rate Limit and Frame Base Accounting parameters are configurable only when a scheduler is not specified in the Egress Scheduler panel.

The Virtual Port Hashing parameter applies only to standard LAG ports. The Secondary Shaper Hashing parameter applies only to HSMDA and HSQ ports. Both of these parameters can also be configured on a local definition of an existing subscriber profile, accessed on the Local Definitions tab.

4

Configure policies to associate with the subscriber profile.

1. Select an accounting policy in the Accounting Policy panel and configure the Volume Stats Type parameter.
2. Configure the required policy objects on the subscriber profile:
   - ANCP policy
   - host tracking policy
   - NAT policy
   - IGMP policy
   - MLD policy
   - subscriber multicast CAC policy
   - PIM policy
   - UPnP policy
   - NAT prefix list (NAT policy must also be configured)
   - Firewall policy

5

Click on the Scheduling tab and configure the required parameters.

An additional parameter, Encapsulation Offset Mode, can be configured using the XML API.

6

Configure LAG per-link hashing parameters, if required.

7

Select an ingress or egress scheduler in the Ingress or Egress panel.

You cannot specify an egress scheduler in a subscriber profile when the Egress Aggregate Rate Limit parameter is set to a value greater than zero.

8

Configure HS Egress Aggregate Rate Limit parameters, as required.

9

Click on the RADIUS Accounting tab and select first, second, and third RADIUS accounting policies. Depending on the NE version, the third RADIUS accounting policy option may not be displayed.

If you configure the Optimize for Session Accounting parameter to reduce accounting messages in a dual stack PPP session configuration, you must also configure the Delay Start Time parameter on the RADIUS accounting policies; see 57.3 “To configure a RADIUS-based accounting policy” (p. 1793).

10

Click on the SLA Profiles tab to configure SLA profiles for the subscriber profile.

1. On the General tab, configure the Use Direct Map as Default and HS SLA Profile Handling Mode parameters.
2. Click on the Profiles tab.
3. Click Create or choose an SLA profile entry and click Properties. The SLA Profile Entry (Create|Edit) form opens.
4. Configure the SLA Profile String parameter.
5. Select an SLA profile to associate with the subscriber profile.
6. Save your changes and close the form.

11

Click on the HSMDA QoS tab and perform one of the following:

a. Select an ingress policy in the Ingress panel.

b. Select an egress policy in the Egress panel and configure the Aggregate Rate Limit parameter.

12

Select a WRR policy in the Egress HSMDA Override panel and configure the Packet Byte Offset parameter, if required.

13

Click on the Policer Control tab to configure ingress and egress policer control policies.

1. Select a policer control policy in the Ingress Policer Control Policy panel.
2. Configure local override parameters for the selected ingress policer control policy, if required.
3. Select a policer control policy in the Egress Policer Control Policy panel.
4. Configure local override parameters for the selected egress policer control policy, if required.

14

Click on the Override Policy Items tab to specify local overrides for schedulers, queues, or policers assigned to the subscriber profile.

15

On the Ingress Schedulers tab, click Create or choose an ingress scheduler entry and click Properties. The Ingress Scheduler Entry Override (Create|Edit) form opens.

1. Select a scheduler policy.
2. Click on the Override tab and configure the required parameters.

   The CIR parameter can only be configured when Summed CIR is set to false.
3. Save your changes and close the form.

16

On the Egress Schedulers tab, click Create or choose an egress scheduler entry and click Properties. The Egress Scheduler Entry Override (Create|Edit) form opens.

1. Select a scheduler policy.
2. Click on the Override tab and configure the required parameters.

   The CIR parameter can only be configured when Summed CIR is set to false.
3. Save your changes and close the form.

17

On the Ingress Policer Level tab, click Create or choose an ingress policer level override entry and click Properties. The Ingress Policer Level Override (Create|Edit) form opens.

1. Select a policer level.

   The current maximum cumulative buffer space value for the policer level is displayed.

2. Click on the Override tab and then enable and configure the Maximum Cumulative Buffer Space (bytes) parameter.
3. Save your changes and close the form.

18

On the Egress Policer Level tab, click Create or choose an egress policer level override entry and click Properties. The Egress Policer Level Override (Create|Edit) form opens.

1. Select a policer level.

   The current maximum cumulative buffer space value for the policer level is displayed.
2. Click on the Override tab and then enable and configure the Maximum Cumulative Buffer Space (bytes) parameter.
3. Save your changes and close the form.

19

On the HSMDA Ingress Queues tab, click Create or choose an access ingress HSMDA queue override entry and click Properties. The Access Ingress HSMDA Queue Override (Create|Edit) form opens.

1. Select an access ingress policy queue.
2. To configure separate IPv4 and IPv6 statistics counters for the queue override, click on the Override tab and configure the Stat Mode parameter.
3. Save your changes and close the form.

20

On the HSMDA Ingress Policers tab, click Create or choose an access ingress HSMDA policer override entry and click Properties. The Access Ingress HSMDA Policer Override (Create|Edit) form opens.

1. Select an access ingress policy policer.
2. To configure separate IPv4 and IPv6 statistics counters for the policer override, click on the Override tab and configure the Stat Mode parameter.
3. Save your changes and close the form.

21

On the Access Egress HSMDA Queues tab, click Create or choose an access egress HSMDA queue override entry and click Properties. The Access Egress HSMDA Queue Override (Create|Edit) form opens.

1. Select an access egress policy queue.
2. To specify override values for the queue, click on the Override tab and configure the required parameters.

If you enable the Override check box for any of these parameters, then the original value of the parameter that was configured for the Access Egress HSMDA Queue will appear in that parameter's value field. You can then configure the override value for that parameter.

You can only specify an override value for WRR Weight for queues 1, 2, or 3.

3. Select an HSMDA slope policy.
4. Save your changes and close the form.

22

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 "To release and distribute a policy" (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.5 To configure an SLA profile

64.5.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

CAUTION

Service Disruption

Modifying an active SLA profile is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the SLA profile before you proceed.

Click Create→SLA Profile or choose an existing SLA Profile entry and click Properties. The SLA Profile (Create|Edit) form opens.

3

Configure the required parameters.

The Displayed Name parameter is configurable only during SLA profile creation.

4

Select a credit control policy to associate with the SLA profile.

5

Click on the QoS tab to associate ingress and egress QoS policies with the SLA profile.

1. Select an access ingress policy to associate with the SLA profile.

2. Configure the required ingress parameters.
3. Select an access egress policy to associate with the SLA profile.
4. Select a scheduler policy to associate with the SLA profile.
5. Configure the required egress parameters.

6

Click on the ACL tab to associate ingress or egress IP filters with the SLA profile. Select any of the following filters:

- ingress IP filter
- egress IP filter
- ingress IPv6 filter
- egress IPv6 filter
- one-time HTTP redirect filter

7

Click on the Supported Service Models tab to specify the intended service model(s) for the SLA profile.

8

Click on the Category Map tab.

1. Select a category map to associate with the SLA profile.

   If you selected a credit control policy in Step 4, you will not be able to select a category map.
2. Click Create or choose a category entry and click Properties. The Category (Create|Edit) form opens.
3. Select a category name and configure the Idle Timeout parameter.

   Note:

   The session will be disconnected if the value of the Activity Threshold parameter is not met during the period specified by the Idle Timeout parameter. See 64.13 “To configure a category map policy” (p. 1855) for information about configuring the Activity Threshold parameter.
4. Save your changes and close the form.

9

Click on the Bonding Selection tab to configure the SLA profile for use in a connection bonding configuration.

Configure the required selection bonding parameters.

10

Click on the Host Limits tab to specify a maximum host count for the SLA profile.

You can set an overall limit, or configure host limits at a more granular level for each of the IPv4, IPv6, and L2TP host type categories, and their respective subcategories.

Although the NFM-P specifies a Host Limit parameter range of 1-131071, the maximum host limit varies for different NE types and chassis types.

11

Click on the Override Policy Items tab to specify local overrides for queue parameters of the QoS policies assigned to the SLA profile.

12

On the Access Ingress Queues tab, click Create or choose an ingress queue override item and click Properties. The Access Ingress Queue Override (Create|Edit) form opens.

1. Select an access ingress policy queue.
2. Click on the Override tab and configure the required parameters.
3. Save your changes and close the form.

13

On the Access Egress Queues tab, click Create or choose an egress queue override item and click Properties. The Access Egress Queue Override (Create|Edit) form opens.

1. Select an access egress policy queue.
2. Click on the Override tab and configure the required parameters.
3. Configure overridden HS queue parameters and select an override slope policy in the HS WRED Queue, if required.
4. Save your changes and close the form.

14

On the Access Ingress Policer tab, click Create or choose an ingress policer override item and click Properties. The Ingress Policer Override (Create|Edit) form opens.

1. Select an access ingress policy policer.
2. Click on the Override tab and configure the required parameters.
3. Save your changes and close the form.

15

On the Access Egress Policer tab, click Create or choose an egress policer override item and click Properties. The Egress Policer Override (Create|Edit) form opens.

1. Select an access egress policy policer.
2. Click on the Override tab and configure the required parameters.
3. Save your changes and close the form.

16

On the Egress Scheduler tab, specify local overrides for traffic flow rate parameters of the egress schedulers assigned to the SLA profile.

1. Click Create or choose an egress scheduler override item and click Properties. The Egress Scheduler Override (Create|Edit) form opens.
2. Select the scheduler policy you want to override.
3. Click on the Override tab and configure the required parameters.

   The CIR parameter is available only when the Summed CIR parameter is set to False.
4. Save your changes and close the form.

17

On the HS WRR Groups tab, specify local overrides for traffic flow rate parameters of the egress schedulers assigned to the SLA profile.

1. Click Create or choose an HS WRR group override item and click Properties. The HS WRR Group Override (Create|Edit) form opens.
2. Select an HS WRR group.
3. Click on the Override tab and configure the required parameters.
4. Save your changes and close the form.

18

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.6 To configure a subscriber explicit map entry

64.6.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

CAUTION

Service Disruption

Modifying an active subscriber explicit map entry is potentially service-affecting.

Ensure that you consider the implications of reconfiguring the subscriber explicit map entry before you proceed.

Click Create→Subscriber Explicit Map Entry or choose an existing subscriber explicit map entry and click Properties. The Subscriber Explicit Map Entry (Create|Edit) form opens.

3

Configure the required parameters.

The Subscriber Identification parameter is configurable only during subscriber explicit map creation.

4

Select a subscriber profile to associate with the subscriber explicit map.

5

Select an SLA profile to associate with the subscriber explicit map.

6

Select an application profile to associate with the subscriber explicit map.

Note: The global AA group policy must be manually distributed to the NE before a global subscriber identification policy using the application profiles can be distributed.
When a new subscriber identification or subscriber explicit map is discovered from the NE, the AA group policy is not automatically resynchronized to the global subscriber identification policy.

7

Save your changes and close the form.

END OF STEPS

64.7 To configure an ANCP policy

64.7.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create ANCP or choose an existing ANCP policy and click Properties. The ANCP Policy (Create|Edit) form opens.

3

Configure the required parameters.

You can select an ingress or egress rate modification scheduler when the Rate Modification parameter is set to Scheduler.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.8 To configure an ANCP MSS static map

Note: You can configure an ANCP MSS static map when you have created an ANCP policy and an Aggregation Scheduler. See 64.7 “To configure an ANCP policy” (p. 1849) and 50.56 “To create an Aggregation Scheduler” (p. 1597) for more information.

64.8.1 Steps

1

Choose Manage→Service→Customers from the NFM-P main menu. The Manage Customers form opens.

2

Select a customer in the list and click Properties. The Customer (Edit) form opens.

3

Click on the Aggregation tab.

4

Select a site in the list and click Properties. The Aggregation Scheduler (Edit) form opens.

5

Click ANCP Static Map.

6

Select an ANCP policy in the list and click Create. The ANCP MSS Static Map (Create) form opens.

7

Configure the required parameters and click OK.

The ANCP Policy is listed in the ANCP Static Map tab of the Aggregation Scheduler (Edit) form.

8

Save your changes and close the forms.

END OF STEPS

64.9 To configure a PPP policy

64.9.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→PPP Policy or choose an existing PPP policy and click Properties. The PPP Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Apply button. The PPP Policy form refreshes.

5

Click on the Options tab to configure PPP options on the PPP policy.

1. Click Create or choose an existing PPP option entry and click Properties. The PPP Option (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

6

Click on the MLPPP tab to configure MLPPP on the PPP policy.

7

Configure the required parameters.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.10 To configure an MSAP policy

64.10.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form appears.

2

Click Create→MSAP Policy or choose an existing MSAP policy and click Properties. The MSAP Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Security tab to configure security policies on the MSAP Policy.

1. Select an NE DoS protection policy to associate with the MSAP policy.
2. Configure the MAC Monitoring and IP Src Monitoring parameters.
3. Select an NE DDoS protection policy to associate with the MSAP policy.

5

Click on the Subscriber Management tab to configure subscriber management parameters on the MSAP policy.

1. Configure the required parameters.
2. On the Policies panel, select the following policy objects:
   - default subscriber profile
   - default SLA profile
   - subscriber identification policy
3. On the Default Application Profile panel, select a default application profile.
4. On the Single Subscriber Configuration panel, select the following policy objects:
   - Non-Subscriber Traffic Subscriber Profile
   - Non-Subscriber Traffic SLA Profile
5. On the Non-Subscriber Traffic Application Profile panel, select a non-subscriber traffic application profile.

6

Click on the ATM tab.

1. Configure the ATM OAM Alarm Cell Handling and Periodic ATM OAM Loopback parameters.
2. Select an ingress ATM policy.
3. Select an egress ATM policy.

7

Perform the following, as required:

a. If the MSAP policy is for an MSAP on a VPLS, click on the VPLS Only Configuration tab.
   1. On the General tab, configure the required parameters.

      An egress multicast group policy cannot be applied to SR NE variants at Release 15.0 R1 and higher.
   2. Select an egress multicast group policy.
   3. Select an access Ingress policy.
   4. Select an access Egress policy.
   5. Click on the DHCP tab and configure the required parameters.
   6. Click on the IGMP Snooping tab and configure the required parameters.
   7. Select a multicast CAC policy.

b. If the MSAP policy is for an MSAP on an IES or VPRN, click on the IES VPRN Only Configuration tab.
   1. Set the Anti-Spoofing parameter to be Source IP and MAC Address.
   2. Configure the Ingress Queuing Type parameter.
   3. Select an access Ingress policy.
   4. Select an access Egress policy.

c. Click on the VPLS Multicast CAC Constraints tab.
   1. On the LAG Port Down tab, click Create or select an existing multicast CAC Lag entry and click Properties. The Multicast CAC Lag Entry (Create|Edit) form opens.

      Up to eight Multicast CAC LAG entries can be created for an MSAP policy
   2. Configure the Number of Ports Down and Level ID parameters:
   3. Save your changes and close the form.
   4. Click on the Levels tab.

      Up to eight Multicast CAC Level entries can be created for an MSAP policy.
   5. Click Create or select an existing muliticast CAC level entry and click Properties. The Multicast CAC Level (Create|Edit) form opens.
   6. Configure the Level ID and Bandwidth parameters:

      The Bandwidth range is 1 to 4 294 967 952 for an MSAP policy.
   7. Save your changes and close the form.

d. Click on the IGMP Host Tracking tab and configure the required parameters.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.11 To clear idle sticky MSAPs

64.11.1 Purpose

Perform this procedure to clear idle sticky MSAPs created by a local MSAP policy.

64.11.2 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Select the MSAP policy for which you want to clear idle sticky MSAPs and click Properties. The MSAP Policy (Edit) form opens.

3

Click on the Local Definitions tab.

4

Select a policy definition and click Properties. The Policy Definition Properties form opens.

5

Click Clear Idle MSAPs.

6

Click OK. The MSAPs are cleared.

7

Close the forms.

END OF STEPS

64.12 To configure a host tracking policy

64.12.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Host Tracking Policy or choose an existing host tracking policy and click Properties. The Host Tracking Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.13 To configure a category map policy

64.13.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form appears.

2

Click Create→Category Map Policy or choose an existing category map policy and click Properties. The Category Map Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Category tab to configure a category entry for the category map policy.

A maximum of 16 categories can be configured for each category map policy. If you must delete categories from the global policy in order to create space for new categories, you must also apply your changes to the policy after the deletions, and then distribute the policy as described in 49.6 “To release and distribute a policy” (p. 1476). This will clear out the deleted categories from all local instances of the policy, leaving room for the new categories when the global policy is re-distributed.

1. Click Create or select an existing category entry and click Properties. The Category (Create|Edit) form opens.
2. Configure the required parameters.

   The Rating Group parameter is only configurable when the Use Rating Group check box is enabled.

   The PIR (kbps) parameter is configured to its maximum value when the MAX check box is enabled.

5

Click on the Queues tab.

Select the required ingress and egress queues to be defined in the category.

6

Click on the Policers tab.

Select the required ingress and egress policers to be defined in the category.

7

Click on the Credit Exhausting Service Levels tab and perform one of the following:

a. To define credit exhausting service levels for IP filter entries, go to Step 8.

b. To define credit exhausting service levels for IPv6 filter entries, go to Step 9.

8

Click on the IP Filter Entries tab.

1. Click Create or select an existing filter entry and click Properties. The Exhausted IP Filter (Create|Edit) form opens.
2. Configure the required parameters on the General and Filter Properties tabs.

   If you are configuring the filter entry for an HTTP redirect override function, set the Action parameter to HTTP Redirect. The Web Redirect tab appears.
3. Click on the Web Redirect tab. Enable the Allow Override parameter and specify a redirect URL.
4. Save your changes and close the form.

9

Click on the IPv6 Filter Entries tab.

1. Click Create or select an existing filter entry and click Properties. The Exhausted IPv6 Filter (Create|Edit) form opens.
2. Configure the required parameters on the General and Filter Properties tabs.
3. Save your changes and close the form.

10

Save your changes and close the Category form.

11

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.14 To configure a credit control policy

64.14.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form appears.

2

Click Create→Credit Control Policy or choose an existing credit control policy and click Properties. The Credit Control Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

If you set the Credit Control Server parameter to Diameter, select a diameter policy and a diameter application policy.

5

Select a category map in the Category Map panel.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.15 To configure an IGMP policy

64.15.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→IGMP Policy or choose an existing IGMP policy and click Properties. The IGMP Policy (Create|Edit) form opens.

3

Configure the required parameters.

If any (or all) of the three query interval parameters are configured, the Query Interval parameter value must be greater than the Last Listener Query Interval and Response Query Interval parameter values.

4

Select a multicast reporting destination in the multicast reporting panel.

5

Click on the Static Group/Source tab. A list of multicast static group/static source address pairs appears.

6

Click Create or choose an existing address pair and click Properties. The IGMP Policy (Create|Edit) form opens.

7

Configure the Static Multicast Group and Static Source parameters.

8

Save your changes and close the form.

9

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.16 To configure an IPoE session policy

64.16.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→IPoE Session Policy or choose an existing IPoE Session policy and click Properties. The IPoE Session Policy (Create|Edit) form opens.

3

Configure the required parameters.

The Circuit ID or Remote ID option should only be selected for the Session Key parameter if all hosts in the IPoE session have one of these fields in their trigger packets. Do not enable both of these options.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.17 To configure a BGP Peering policy

Note: Most of the parameters in a BGP peering policy can be configured directly on the policy object, or the parameter value can be inherited from the parent BGP object by enabling the Inherit Value check box next to each parameter.

64.17.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→BGP Peering Policy or choose an existing BGP peering policy and click Properties. The BGP Peering Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Behavior tab and configure the required parameters.

5

Click on the AS Properties tab and configure the required parameters.

6

Click on the Import Policies tab and select import routing policy statements for the BGP peering policy.

7

Click on the Export Policies tab and select export routing policy statements for the BGP peering policy.

8

Click on the Authentication tab and configure the MD5 Authentication and Authentication Key parameters.

9

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.18 To configure a diameter policy

64.18.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Diameter Policy or choose an existing diameter policy and click Properties. The Diameter Policy (Create|Edit) form opens.

Modifying an active diameter policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the diameter policy before you proceed.

3

Configure the required parameters.

If you set the Virtual Router Type parameter to VPRN Service, you must also specify a VPRN service.

4

Click on the DCAA Parameters tab and configure the required parameters.

5

Click on the Peers tab to configure diameter peers.

1. Click Create or choose an existing diameter peer entry and click Properties. The Diameter Peer (Create|Edit) form opens.
2. Configure the required parameters.

   If you want to specify local override settings for any of the timer settings for the peer, disable the Inherit Value check box and then specify a value for any of the timer parameters:
   - Watchdog Timer (seconds)
   - Connection Timer (seconds)
   - Transaction Timer (seconds)

   When the Inherit Value check box is selected for a timer parameter on a peer object, the timer value is inherited from the parent diameter policy.
3. Save your changes and close the form.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.19 To configure a subscriber multicast CAC policy

64.19.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Subscriber Multicast CAC Policy or choose an existing subscriber multicast CAC policy and click Properties. The Subscriber Multicast CAC Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.20 To configure a mobile gateway/peer profile

64.20.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Mobile Gateway/Peer Profile or choose an existing mobile gateway/peer profile and click Properties. The Mobile Gateway/Peer Profile (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a python policy, if required.

5

If you want to specify non-default message-retransmit parameters, disable the Message-Retransmit - Default parameter and configure the Message-Retransmit - Timeout (seconds) and Message-Retransmit - Retries parameters.

6

If you want to specify non-default keep-alive parameters, disable the Keep-Alive - Default parameter and configure the Keep-Alive - Interval (seconds), Keep-Alive - Timeout (seconds), and Keep-Alive - Retries parameters.

7

On the Charging Characteristics panel, configure the rquired home charging characteristics and roaming charging characteristics.

8

Click on the PGW tab to configure packet data network gateway QoS parameters.

Configure the required parameters.

You must disable the Default check box for each parameter in order to configure the parameter.

9

Click on the GGSN tab to configure gateway GPRS support node QoS parameters.

Configure the required parameters.

You must disable the Default check box for each parameter in order to configure the parameter.

10

Click on the MME tab to configure gateway MME node parameters.

Configure QoS, Dowlink, Uplink, and AMBR parameters as required.

11

Click OK to save the profile and close the form, or click Apply to save the profile. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the profile to NEs.

END OF STEPS

64.21 To configure a host lockout policy

64.21.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Host Lockout Policy or choose an existing host lockout policy and click Properties. The Host Lockout Policy (Create|Edit) form opens.

3

Configure any applicable parameters:

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.22 To clear host lockouts

64.22.1 Purpose

Perform this procedure to clear host lockouts enforced by a specific host lockout policy. Lockouts can be cleared at the policy level, or they can be cleared from an NE with which the policy is associated. Host lockouts can also be cleared from IES or VPRN SAPs.

64.22.2 Steps

1

To clear host lockouts, do one of the following:

a. To clear host lockouts at the host lockout policy level or NE level, complete the sub-steps within Step 2.

b. To clear host lockouts from a SAP on an IES or VPRN, complete the sub-steps within Step 3.

2

To clear host lockouts at the host lockout policy level or NE level, do the following:

1. Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.
2. Select the host lockout policy for which you want to clear lockouts and click Properties. The Host Lockout Policy (Edit) form opens.
3. Click on the Local Definitions tab.
4. Select a policy definition and click Properties. The Policy Definition Properties form opens.
5. Click Clear Host Lockout Policy.
6. Choose one of the following options:
   - Network Element: all host lockouts are cleared from the associated NE.
   - Policy: all host lockouts are cleared for the selected policy.
7. Click OK. The lockouts are cleared.
8. Close the forms.

3

To clear host lockouts from IES or VPRN service SAPs, do the following:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Select a service that has SAPs from which you want to clear host lockouts and click Properties. The <*service*> (Edit) form appears.
3. On the navigation tree, expand Sites→<*site_n*>→Subscriber Interfaces→<*subscriber_interface_n*>→<*group_interface_n*>→Service Access Points.object, and then expand a group interface object to navigate to the SAP .

4. Right-click on the SAP object from which you want to clear host lockouts and choose Properties from the menu. The Service Access Point (Edit) form opens.
5. Click on the Subscriber Management tab, and then click on the Locked Out Hosts tab.
6. Click on the Search button to display a list of locked out hosts.
7. Click Clear Locked Out Hosts. You can either clear all host lockouts on the SAP, or you can clear lockouts selectively by specifying ID criteria.
8. Click OK. The lockouts are cleared.
9. Close the forms.

End of steps

64.23 To configure a RADIUS script policy

64.23.1 Purpose

A RADIUS script policy references python scripts for RADIUS AAA packet manipulation in a subscriber management application.

64.23.2 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→RADIUS Script Policy or choose an existing RADIUS script policy and click Properties. The RADIUS Script Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

End of steps

64.24 To configure an HTTP redirect policy

64.24.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→HTTP Redirect Policy or select an existing HTTP redirect policy and click Properties. The HTTP Redirect Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Forward Entries tab to configure forward entries.

1. Click Create or select an existing forward entry and click Properties. The HTTP Redirect Forward Entry form appears.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.25 To configure an MLD policy

64.25.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→MLD Policy or select an existing MLD policy and click Properties. The MLD Policy (Create|Edit) form opens.

3

Configure the required parameters.

If any (or all) of the three query interval parameters are configured, the Query Interval parameter value must be greater than the Last Listener Query Interval and Response Query Interval parameter values.

4

Select an import policy.

5

Select a redirection policy.

6

If you set the Rate Modification parameter to Scheduler, select a rate modification scheduler.

7

Click on the Static Groups/Sources tab.

1. Click Create or select an existing static multicast group and click Properties. The MLD Policy Static (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.26 To configure a diameter application policy

64.26.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Diameter Application Policy or select an existing diameter application policy and click Properties. The Diameter Application Policy (Create|Edit) form opens.

Modifying an active diameter application policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Select a diameter peer policy.

5

If you set the Application parameter to Gx in Step 3, click on the Gx tab.

1. Select applicable diameter attributes.

   Selecting certain diameter attributes requires additional configuration:
   - NAS Port Type - Port Type Value
   - NAS Port - NAS Port Bit Specification
   - Calling Station ID - Calling Station ID Type
   - User Equipment Info - User Equipment Info Type
   - NAS Port ID - Port Prefix Type, Port Prefix String, Port Suffix Type
2. Configure the remaining parameters, as required.

   Enable the Show Format Examples check box to display a series of MAC address syntax examples for the MAC Format parameter.

6

If you set the Application parameter to Gy in Step 3, click on the Gy tab and configure the required parameters.

7

If you set the Application parameter to NASREQ in Step 3, click on the NASREQ tab.

1. Select applicable diameter attributes.

   Selecting certain diameter attributes requires additional configuration:
   - NAS Port Type - Port Type Value
   - NAS Port - NAS Port Bit Specification
   - Calling Station ID - Calling Station ID Type
   - NAS Port ID - Port Prefix Type, Port Prefix String, Port Suffix Type
2. Configure the Name Format parameter.

   If the Name Format parameter is set to DHCP Client Vendor Options, configure the MAC Format parameter. Enable the Show Format Examples check box to display a series of MAC address syntax examples for the MAC Format parameter.
3. Configure the Domain parameter (if required) and the Password parameter.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.27 To configure a distributed subscriber management traffic policer

64.27.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Distributed Subscriber Management Traffic Policer or select an existing policer item and click Properties. The Distributed Subscriber Management Traffic Policer (Create|Edit) form opens.

Modifying an active distributed subscriber management traffic policer is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policer before you proceed.

3

Configure the required parameters.

The Action parameter is configurable only if the Type parameter is set to Single Bucket Bandwidth.

The CBS parameter is configurable only if the Type parameter is set to Dual Bucket Bandwidth.

4

On the Rates and Adaptation Rules panels, configure the required parameters.

The CIR parameter is configurable only if the Type parameter is set to Dual Bucket Bandwidth.

5

Click OK to save the policer and close the form, or click Apply to save the policer. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policer to NEs.

END OF STEPS

64.28 To configure a distributed subscriber management IP filter policy

64.28.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Distributed Subscriber Management IP Filter Policy or select an existing policy item and click Properties. TheDistributed Subscriber Management IP Filter Policy (Create|Edit) form opens.

Modifying an active distributed subscriber management IP filter policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Click on the IPv4 Filter Entries tab to configure IPv4 filter entries.

5

Click Create to create a new filter entry or select an existing filter entry and click Properties. The DSM IPv4 Filter Entry (Create|Edit) form appears.

1. Configure the required parameters.

   The Port parameter is configurable if the Protocol parameter is set to TCP or UDP.
2. Save your changes and close the form.

6

Click on the IPv6 Filter Entries tab to configure IPv6 filter entries.

7

Click Create to create a new filter entry or select an existing filter entry and click Properties. The DSM IPv6 Filter Entry (Create|Edit) form appears.

1. Configure the required parameters.

   The Port parameter is configurable if the Protocol parameter is set to TCP or UDP.
2. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policer to NEs.

END OF STEPS

64.29 To configure a UPnP policy

64.29.1 Steps

1

Choose Policies→ISA Policies→UPnP Policies from the NFM-P main menu. The UPnP Policies form opens.

2

Click Create or select an existing UPnP policy and click Properties. The UPnP Policy (Create|Edit) form opens.

Modifying an active UPnP policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.30 To configure a PIM policy

64.30.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→PIM Policy or select an existing PIM policy and click Properties. The PIM Policy (Create|Edit) form opens.

Modifying an active PIM policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the Displayed Name and Description parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.31 To configure an SHCV policy

64.31.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→SHCV Policy or select an existing SHCV policy and click Properties. The SHCV Policy (Create|Edit) form opens.

Modifying an active SHCV policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Click on the Triggers tab to configure SHCV triggers.

1. Select a trigger entry in the list and click Properties. The SHCV Trigger Entry form appears.
2. Configure the required parameters and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.32 To configure a RIP policy

64.32.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→RIP Policy or select an existing RIP policy and click Properties. The RIP Policy (Create|Edit) form opens.

Modifying an active RIP policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.33 To configure a service chaining EVPN policy

64.33.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Service Chaining EVPN Policy or select an existing EVPN policy and click Properties. The Service Chaining EVPN Policy (Create|Edit) form opens.

Modifying an active EVPN policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the EVPN Service ID, Import Mode, and Prefix Route Resolution parameters.

These parameters cannot be changed after policy creation.

If Import Mode is set to Bridged, the Prefix Route Resolution parameter is not displayed on the policy form, but is set to Recursive by default.

The EVPN Service ID parameter value must be unique among other EVPN policies, and among services that have a service ID.

4

Depending on the setting of the Import Mode parameter, do one of the following:

- If the Import Mode parameter is set to Bridged, configure the parameters on the Route Distinguisher, Route Target, and VXLAN panels, or accept default values.
- If the Import Mode parameter is set to None or Routed, configure the parameters on the

Route Distinguisher, Route Target, GW Address Range, VXLAN, and IP Advertise Routes panels or accept default values.

The Route Distinguisher and Route Target parameters have a unique mandatory syntax which is described in the tooltips for the fields.

In order for an EVPN policy to be set to Admin State: In Service, the Route Distinguisher, Route Target, and VXLAN parameters must be configured with non-default values.

In order for the IP Advertise Routes - Admin State parameter to be set to Enabled, the GW Address Range - Start and End parameters must be configured.

5

If the Import Mode parameter is set to None or Routed, configure NAT pools for the EVPN policy. Associated NAT pools must be of type L2-aware, and must be configured on the base routing instance or VPRN routing instance.

1. On the NAT Pools panel, click Create or select an existing NAT pool entry and click Properties. The NAT Pool (Create|Edit) form opens.
2. Do one of the following:
   - Set the Router Type parameter to Base and select a pool.
   - Set the Router Type parameter to VPRN, select a VPRN service, and then select a pool.
3. Save your changes and close the form.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.34 To configure a service chaining VAS filter policy

64.34.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→Service Chaining VAS Filter Policy or select an existing VAS filter policy and click Properties. The Service Chaining VAS Filter Policy (Create|Edit) form opens.

Modifying an active VAS filter policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

On the General tab, configure the Displayed Name and Description parameters.

4

Click on the VAS Filter Entries tab to configure filter entries for the policy.

1. Click Create or select an existing filter entry and click Properties. The VAS Filter Entry (Create|Edit) form opens.
2. Configure the required parameters.

   You must configure at least one of the Protocol, Foreign IP Address, or Foreign IP Address Prefix (with Mask and Port) parameters.
3. Click on the Actions tab to configure filter actions.

   In order to configure filter actions, the filter entry Admin State must Out of Service.
4. Click Create or select an existing action entry and click Properties. The VAS Filter Action (Create|Edit) form opens.
5. Configure the required parameters.

   Only one Downstream action and one Upstream filter action can be created per filter entry.
6. Save your changes and close the forms.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.35 To configure a BRG profile

64.35.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→BRG Profile or select an existing BRG profile and click Properties. The BRG Profile (Create|Edit) form opens.

Modifying an active BRG Profile is potentially service-affecting. Ensure that you consider the implications of reconfiguring the profile before you proceed.

3

Configure the required parameters.

4

Select a RADIUS server policy and specify a RADIUS authentication password.

5

Click on the DHCP Pool tab to configure a DHCP pool for the BRG profile.

Configure the required parameters.

6

Click on the Custom Options tab to configure DHCP options for the BRG profile. The options are included in DHCP reply messages.

1. Click Create or select an existing custom option entry and click Properties. The DHCP Pool Custom Option (Create|Edit) form opens.
2. Configure the Number and Type parameters.
3. Depending on the value of the Type parameter, configure the following:
   - IP addresses (for IP address custom option type)
   - String value (for ASCII or Hex string custom option type)
4. Save your changes and close the form.

7

Click on the RADIUS Proxy Servers tab to associate proxy servers with the BRG profile.

1. Click Create or select an existing RADIUS proxy server entry and click Properties. The RADIUS Proxy (Create|Edit) form opens.
2. Configure the Router Instance parameter.
3. Depending on the value of the Router Instance parameter (Base or VPRN), select a VPRN service (where required) and a RADIUS proxy server.
4. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.36 To configure a trace profile

64.36.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create or select an existing trace profile and click Properties. The Trace Profile (Create|Edit) form opens.

3

Configure the required parameters and select the applications you want to trace.

4

Click OK to save the profile and close the form, or click Apply to save the profile. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the profile to NEs.

END OF STEPS

64.37 To configure a MAP-T domain policy

64.37.1 Steps

1

Choose Policies→ISA Policies→MAP-T Domains from the NFM-P main menu. The MAP-T Domains form opens.

2

Click Create or select an existing MAP-T domain policy and click Properties. The MAP-T Domain (Create|Edit) form opens.

Modifying an active MAP-T domain policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

Configure the required parameters.

4

Click on the Mapping Rules tab to configure a mapping rules for the MAP-T domain policy.

1. Click Create or select an existing mapping rule and click Properties. The Mapping Rule (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

64.38 To configure an APN policy

64.38.1 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form opens.

2

Click Create→APN Policy or select an existing APN policy and click Properties. The APN Policy (Create|Edit) form opens.

Modifying an active APN policy is potentially service-affecting. Ensure that you consider the implications of reconfiguring the policy before you proceed.

3

On the General tab, configure the Displayed Name parameter.

4

Click on the APN tab to configure APN entries for the APN policy.

1. Click Create or select an existing APN entry and click Properties. The APN (Create|Edit) form opens.
2. Configure the Displayed Name parameter.
3. Configure the APN entry in ONE of the following ways:
   - For RADIUS authentication; select a RADIUS authentication policy
   - For Diameter authentication; select a Diameter authentication policy
   - For LUDB authentication; select a Local user database
   - For default service; select a service ID and interface ID

   Configure AMBR QoS mapping downlink and uplink information, if required.
4. Save your changes and close the form.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

65 Remote network monitoring policies

65.1 Remote network monitoring policies

65.1.1 Overview

You can create remote network monitoring (RMON) policies that allow you to perform the following tasks.

- Use the policy framework to configure RMON that allows NFM-P users to deploy configurations to multiple devices.
- Map RMON events to information alarms in the NFM-P GUI.

Remote network monitoring allows you to monitor the behavior of functions that are not typically supported by the NFM-P. For example, you can configure events on GNEs to monitor system temperature or CPU usage. The NFM-P can raise informational alarms for the exceeded thresholds. You can also configure custom values to supersede the hard-coded threshold values associated with an NE function.

The NFM-P maps RMON SNMP traps to alarms in the GUI.

Policy distribution and event mapping are supported for the following NE types:

- 7210 SAS
- 7450 ESS
- 7705 SAR
- 7750 SR
- 7950 XRS

Support for 7210 SAS NEs varies depending on the chassis type and release; see the NE documentation for support information.

65.2 To configure a remote network monitoring policy

65.2.1 Steps

1

Choose Tools→Remote Network Monitoring (RMON) from the NFM-P main menu. The Remote Network Monitoring (RMON) form opens.

2

Click Create. The RMON (Create) form opens with the General tab displayed.

3

Configure the required parameters.

4

Create events to associate with the rising and falling thresholds on the remote NE.

1. Click on the Events tab.
2. Click Create. The Event RMON (Create) form opens.
3. Configure the required parameters.
4. Click OK. The Event RMON (Create) form closes and the RMON (Create) form reappears.
5. Repeat 2 to 4 to create another event, if required.

5

Configure the alarm information to associate with an event you created in Step 4 .

1. Click on the Alarms tab.
2. Click Create. The Alarm RMON (Create) form opens.
3. Configure the required parameters.

6

Configure the properties for the rising threshold crossing alarm.

1. Click Select in the Rising TCA Properties panel. The Select Rising Event - Alarm form opens.
2. Select an event and click OK. The Select Rising Event - Alarm form closes and the Alarm RMON form reappears.
3. Click Enable and configure the Rising Threshold parameter for the event.

7

Configure the properties for the falling threshold crossing alarm.

1. Click Select in the Falling TCA Properties panel. The Select Falling Event - Alarm form opens.
2. Select an event and click OK. The Select Falling Event - Alarm form closes and the Alarm RMON form reappears.
3. Click Enable and configure the Falling Threshold parameter for the event.

8

Click OK. The Alarm RMON form closes and the RMON (Create) form reappears.

9

Repeat Step 5 to Step 8 to create another alarm, if required.

10

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

11

Close the form.

END OF STEPS

66 NAT policies

66.1 NAT policies

66.1.1 Overview

A Network Address Translation (NAT) policy, defines general NAT properties and associates a NAT address pool with an ISA-NAT group on the same NE. A NAT policy is associated with a subscriber profile for L2-aware NAT in an IES or VPRN service, or for large-scale NAT in a VPRN service. You can add an IPFIX export policy to a NAT policy. An IPFIX export policy defines how IP traffic flow information is formatted and transferred from an exporter to a collector to facilitate services such as measurement, accounting and billing. You can specify the router instance type, IPFIX Collector address, MTU, and refresh timeouts.

66.1.2 Procedural reference

See Chapter 30, “NAT” for information about configuring and deploying NAT in a network.

67 PCP policies

67.1 PCP policies

67.1.1 Overview

PCP operates between subscribers and NAT directly. This protocol provides the subscriber with limited but direct control over the NAT function. PCP configuration is supported on the 7750 SR.

PCP allows a customer to configure static port forwards, obtain information about existing port forwards and to obtain the outside IP address from software that runs on the home network or the CPE.

A PCP policy allows you to configure PCP parameters that affect the function of a NAT PCP server. The policy is applied to multiple NAT PCP servers so that they behave consistently.

67.2 Workflow to configure and apply PCP policies

67.2.1 Stages

1

Create a PCP policy. See 67.3 “To configure a PCP policy” (p. 1885).

2

Create a NAT PCP server. See 67.4 “To configure a NAT PCP server on a base routing instance” (p. 1886) for information on creating a NAT PCP server on a base routing instance. See 67.5 “To create a NAT PCP server on a VPRN routing instance” (p. 1887) for information on creating a NAT PCP server on a VPRN routing instance.

3

Associate an interface with the NAT PCP server. See 67.6 “To associate an interface with a NAT PCP server” (p. 1887).

67.3 To configure a PCP policy

67.3.1 Steps

1

Choose Policies→ISA Policies→NAT from the NFM-P main menu. The NAT Policies form opens.

2

Choose Port Control Protocol Policy (Network Address Translation) in the object drop-down.

3

Click Create or select an existing PCP policy and click Properties. The PCP Policy (Create|Edit) form opens.

4

Configure the required parameters.

5

Save your changes and close the forms.

End of steps

67.4 To configure a NAT PCP server on a base routing instance

67.4.1 Steps

1

Choose Routing from the navigation tree view selector.

2

Expand *NE→routing_instance*. Right-click on the routing instance and choose Properties. The Routing Instance (Edit) form opens.

3

Click on the PCP Servers tab.

4

Click Create or select an existing PCP server and click Properties. The Port Control Protocol Server (Create|Edit) form opens.

5

Configure the required parameters.

6

Select a PCP policy.

7

Select a routing instance to act as the forwarding inside virtual router.

8

Save your changes and close the forms.

End of steps

67.5 To create a NAT PCP server on a VPRN routing instance

67.5.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

In the navigation tree, expand the Sites icon, right-click on the Routing Instance icon, and choose Properties. The Routing Instance (Edit) form opens.

4

Click on the PCP Servers tab.

5

Click Create or select an existing PCP server and click Properties. The Port Control Protocol Server (Create|Edit) form opens.

6

Configure the required parameters.

7

Select a PCP policy.

8

Select a routing instance to act as the forwarding inside virtual router.

9

Save your changes and close the forms.

END OF STEPS

67.6 To associate an interface with a NAT PCP server

67.6.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

In the navigation tree, expand the Sites icon, right-click on the Routing Instance icon, and choose Properties. The Routing Instance (Edit) form opens.

4

Click on the PCP Servers tab.

5

Choose a server from the list and click Properties. The Port Control Protocol Server (Create) form opens.

6

Click on the Interfaces tab.

7

Click Create. The Port Control Protocol Server Interface form opens.

8

Select an interface for the PCP server.

9

Save your changes and close the forms.

END OF STEPS

68 7705 SAR Security policies

68.1 Security policies

68.1.1 Overview

You can configure the following security policies on the 7705 SAR and 7705 SAR-H:

Release 6.1 R1 or later:

- security profile policy
- security policy
- security zone policy

Release 7.0 R2 or later:

- security log profile policy
- security log policy

Release 7.0 R4 or later:

- security host group policy
- security application group policy
- security policer group policy

Release 20.0 R3 or later:

- security bypass policy

You can view the security policies associated with a 7705 SAR device on the Security tab on the Network Element (Edit) form.

68.2 Zone creation

68.2.1 General Information

On a 7705 SAR, NAT configuration is based on zones. A zone is a group of L2 or L3 interfaces with common criteria that can be bundled together. The same action can be applied to the bundle. For example, in an MC application, all of the SAPs on the access interface used to aggregate the MC can be placed in a single zone and the uplink public interface can be placed in a second zone. All traffic routed between the two zones have NAT applied, based on the security policies created on the first zone. This simplifies the configuration and management of NAT functionality.

You can configure 7705 SAR security policies based on traffic direction entering or leaving the zone.

A zone can be configured on L2 access interfaces, L3 access interfaces, and SDPs on a base routing instance or an EPIPE, IES, MVPLS, VPLS, or VPRN service. Zone configuration is only allowed per service. There is a one-to-one mapping between a service and a zone.

68.3 Security policies and NAT pools

68.3.1 General Information

A 7705 SAR security policy defines how NAT is applied and can vary from subscriber to subscriber. A security policy is applied to zones at creation. Security policies are all type NAPT, denoting a network address and port translation mechanism. Within a security policy, a specific set of matching criteria are configured. If there is a match on a packet, an action is applied. If the action is NAT, the packet has NAT applied to the configured NAT pool IP address and port range.

68.4 Dynamic source NAT

68.4.1 General Information

Source NAT is used to create connections from a private network to a public network. If an arriving IP packet on a 7705 SAR matches the security policy rules, an internal mapping is created between the private source IP address and source port to a public source IP address and source port. The public IP address and port can be configured in the NAT Pool on the security policy.

NAT automatically creates a reverse mapping for incoming traffic from a public domain to a private domain for the same connection. The reverse mapping is based on an outside destination IP address and destination port to an inside destination IP address and destination port.

The configurable outside NAT pool for the source IP address and source port can either be a range of addresses and ports or a unique IP address and port.

The 7705 SAR supports a single public IP address - all inside source IP addresses can be mapped to a single outside IP address and a range of ports by assigning the interface name to the NAT pool configuration. All local interfaces on a 7705 SAR can be assigned to the NAT pool - local L3 interfaces, loopback interfaces, and system interfaces.

68.5 Static destination NAT

68.5.1 General Information

Static destination NAT or static port forwarding is a one-to-one mapping of an outside destination IP address and destination port to an inside destination IP address and port. Static port forwarding can be used when certain applications are hidden behind a single IP address. Each application can be accessed through the public IP address and the destination port for that application.

68.6 Security pairing

68.6.1 General information

A security pair is a minimum configuration requirement on the 7705 SAR to enable Firewall/NAT functionality.

There are 2 steps to creating a security pair:

- Configuring security policies from Policies>Security in the NFM-P main menu.

- Assigning a security zone policy to a network interface. You can also assign a security zone policy to an L3 access interface on an IES or VPRN service site; an L2 access interface or a spoke or mesh SDP on an EPIPE, VPLS, or MVPLS service; and a tunnel interface on a VPRN service.

To enable Firewall or NAT, these two configuration areas must exist and pair together on a 7705 SAR device. The final result is a security policy that is applied to the Zone to enable Firewall/NAT functionality.

68.6.2 Security bypass

When you create a zone for an EPIPE, VPLS, or MVPLS service, you can assign a security bypass policy to L2 services at the site level. The security bypass policy defines protocols and traffic that is permitted to cross the zone at that site, regardless of other firewall configurations. The bypass policy does not affect traffic moving within the security zone, and each bypass policy counts as a global filter entry.

68.7 To configure a security profile policy for a 7705 SAR

68.7.1 Purpose

A security profile policy is used to configure fragmentation, application assurance, and UDP/TCP timers. A security profile policy dictates if fragmented packets are allowed in the network or if the fragmented packets should be discarded. You can configure the timers for different states of a UDP/TCP connection. For example, in a TCP three way hand shake, each state can have its own timer. If the connection does not change state within the allowed time, the connection is closed. Idle timers time out when there are no packets on the session for the period of the configured idle timer. Strict timers time out after the session's last transition state, i.e. the timer starts and counts down from the time that the session was created. The strict timer never renews if a packet arrives on the session.

68.7.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Profile, or choose an existing security profile policy and click Properties. The Security Profile, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

Configure the Allow Fragments parameter.

5

Configure the Application Layer Gateway parameter.

6

Enable the Application Assurance Inspection checkbox, if required.

The Application Assurance parameters appear in the Application Assurance panel when you enable the Application Assurance Inspection checkbox.

7

Configure the ICMP parameters in the ICMP Timeouts panel.

8

Configure the required parameters in the TCP Timeouts, UDP Timeouts, and Other Timeouts panels.

9

Select a policer group policy in the Fwd Policer Group panel and the Rev Policer Group panel, if required.

10

Configure the required parameters in the Application Assurance panel.

You must enable the IP Options Inspection checkbox in the Application Assurance panel in order for the IP Options parameter to appear.

11

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.8 To configure a security policy for a 7705 SAR

68.8.1 Purpose

A security policy is set of rules that a packet must match in order for an action to be performed on the packet.

You cannot delete a security policy if it is associated with a security profile policy or a security zone policy.

68.8.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Policy, or choose an existing security policy and click Properties. The Security Policy, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters and click Apply. Additional tabs are enabled.

4

Assign a security policy entry to the security policy.

1. Click on the Security Policy Entry tab and click Create.
2. Configure the required general parameters.
3. Configure the required parameters in the Limit panel.
4. Select a security profile policy in the Profile panel.

   To create a security profile policy, see 68.7 “To configure a security profile policy for a 7705 SAR” (p. 1891) .
5. Configure the Logging parameter in the Logging panel.
6. If you configure the Logging parameter as To Log, select a Security Log policy. To create a security log policy, see 68.11 “To configure a security log policy for a 7705 SAR” (p. 1897) .

5

Configure the match criteria.

1. Click on the Match Criteria tab.
2. Configure the Flow Direction parameter.
3. Configure the required parameters in the Criteria panel.

   If you specified ICMP (1) as the Protocol, configure the parameters in the ICMP Properties panel.

   If you specified UDPTCP (*), TCP, or UDP as the Protocol, configure the parameters in the Port panel.

   If you enabled the Local parameter, the NAT panel appears. Configure the Destination IP Address and Destination Port parameters.
4. Select an Application Group.
5. Configure the Source IP Operator and Destination IP Operator parameters in the IP Address panel.
6. Select a Source Host Group.
7. Select a Destination Host Group.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.9 To configure a security zone policy for a 7705 SAR

68.9.1 Purpose

Perform this procedure to configure a security zone policy for a 7705 SAR, Release 6.1 R1 or later. You can assign a security zone policy to a network interface. You can also assign a security zone policy to an L3 access interface on an IES or VPRN service site; an L2 access interface or spoke SDP on an EPIPE, VPLS, or MVPLS service; and a tunnel interface on a VPRN service.

68.9.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Zone, or choose an existing security zone policy and click Properties. The Security Zone, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

Configure the Zone Type parameter in the Zone Instance panel.

a. If you specified a Zone Type of BASE, go to Step 5 .

b. If you specified a Zone Type that is a service type, select a service in the appropriate service panel (for example, an IES service in the IES panel).

5

Configure the required parameters in the Inbound Concurrent Sessions panel and Outbound Concurrent Sessions panel.

6

If you chose VPRN in Step 4 , configure the Auto-bind parameter in the MP-BGP Auto-Bind panel.

7

Select a security log policy in the Logging panel. You can click Create to create a security log policy; see 68.11 “To configure a security log policy for a 7705 SAR” (p. 1897) .

8

Select a security policy in the Security Policy panel.

Only one security policy can be assigned to each security zone. The same security policy can be assigned to multiple security zones. See 68.8 “To configure a security policy for a 7705 SAR” (p. 1892) for more information about configuring a security policy.

9

Save your changes. If you specified a Zone Type of VPLS, MVPLS, or EPIPE, go to Step 16.

10

Assign a NAT pool to the security zone.

1. Click on the NAT Pool tab and click Create. The NAT Pool, Global Policy, (Create) form opens with the General tab displayed.
2. Configure the required parameters.

11

To assign a NAT pool entry to the NAT pool, click on the NAT Pool Entry tab and click Create. The Nat Pool Entry, Global Policy (Create) form opens.

12

Configure the required parameters.

13

Perform one of the following:

a. If you specified a Zone Type of BASE in Step 4 , select a Source Network Interface.

b. If you specified a Zone Type of IES in Step 4 , select a Source IES L3 Access Interface.

c. If you specified a Zone Type of VPRN in Step 4 , select a Source VPRN L3 Access Interface.

14

Configure the IP Operator parameter in the IP Address panel.

Note: The configuration of the network interface and IP address are mutually exclusive.

15

Configure the Port Operator parameter in the Port panel.

16

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.10 To configure a security log profile policy for a 7705 SAR

68.10.1 Steps

1

Choose Policies→Log Configuration from the NFM-P main menu. The Manage Log Configuration Policies form opens.

2

Click Create and choose Log Profile or choose an existing log profile policy and click Properties. The Log Profile, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

Configure the parameters in the Throttle Events panel, and click Apply.

5

Click on the Log Event tab.

6

Choose an event type and click Properties.

7

Configure the Event Control parameter and click OK. Repeat Step 6 and Step 7 to configure more event types.

8

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.11 To configure a security log policy for a 7705 SAR

68.11.1 Steps

1

Choose Policies→Log Configuration from the NFM-P main menu. The Manage Log Configuration Policies form opens.

2

Click Create or choose an existing security log policy and click Properties. The Security Log, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

If you configure the Log Destination as SysLog, select a Syslog ID.

5

Select a security log profile policy in the Log Profile panel. See 68.10 “To configure a security log profile policy for a 7705 SAR” (p. 1896) to configure a Security Log Profile policy.

6

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.12 To configure a security host group policy for a 7705 SAR

68.12.1 Purpose

A host group policy is configured with single source/destination IP addresses or a range of IP addresses. For a single pre-assigned set of IP addresses that are allowed in a network, a list can be created once and assigned to every edge router. The host group policy can be assigned to a security policy.

68.12.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Host Group, or choose an existing security host group policy and click Properties. The Host Group, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters and click Apply.

4

Click on the Host Group Entries tab and click Create. The Host Group Entry, Global Policy form opens.

5

Configure the required general parameters and the parameters in the IP Address panel.

6

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.13 To configure a security application group policy for a 7705 SAR

68.13.1 Purpose

An application group policy maintains a group configuration based on common criteria. For example, an application group policy can contain UDP/TCP ports for a specific application or the ICMP type and code which must be matched in the policy. An application group policy can be assigned to a security policy.

68.13.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Application Group, or choose an existing security application group policy and click Properties. The Application Group, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters and click Apply.

4

Click on the Application Group Entries tab and click Create. The Application Group Entry, Global Policy form opens.

5

Configure the required general parameters and the Protocol parameter.

If you choose ICMP for the Protocol parameter, configure the parameters in the ICMP Properties panel. If you choose TCP or UDP for the Protocol parameter, configure the parameters in the Port Properties panel.

6

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.14 To configure a security policer group policy for a 7705 SAR

68.14.1 Purpose

A policer group policy ensures that each connection does not use more than the assigned bandwidth allocation. Any packet above the policed rate is dropped and discarded with a log event. A policer group policy can be assigned to a security profile policy.

68.14.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Policer Group, or choose an existing security policer group policy and click Properties. The Policer Group, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

Configure the Maximum Ingress Bandwidth (Mbps) and Committed Burst Size (bytes) parameters in the Burst Size panel and click Apply.

5

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

68.15 To configure a security bypass policy for a 7705 SAR

68.15.1 Purpose

A security bypass policy is a set of match criteria that define what traffic is allowed to pass across a security zone at the L2 service where the policy is applied. If a security bypass policy is not applied, then traffic is filtered based on the security zone configuration; if the policy is applied, then traffic that matches the criteria is allowed to pass regardless of other security zone configurations.

68.15.2 Steps

1

Choose Policies→Security from the NFM-P main menu. The Security Policies form opens.

2

Click Create and choose Security Bypass Policy, or choose an existing policy and click Properties. The Security Bypass, Global Policy (Create|Edit) form opens with the General tab displayed.

3

Configure the required general parameters.

4

Click on the Security Bypass Entry tab and create security bypass entries as required. Traffic matching the criteria defined in the entries is allowed to pass through the site where the policy is applied. Perform the following to create a security bypass entry:

1. Click Create. The Security Bypass Entry form opens.
2. Configure the required general parameters.
3. Click on the Match Criteria tab and configure the match criteria for the entry.
4. Save your changes.

5

Save your changes. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

69 PDN profile policies

69.1 PDN profile policies

69.1.1 Overview

You can create a PDN (Packet Data Network) profile policy to define the cellular interface used to connect mobile data users to the network from 7705 SAR-Hm nodes.

The PDN profile policy defines the APN (Access Point Name) used to define the type of network connection used, the Authorization Type (either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol), the user name, and password.

69.2 To configure a PDN Profile policy for a 7705 SAR-Hm

69.2.1 Steps

1

Choose Policies→PDN Profile from the NFM-P main menu. The PDN Profile Policy form opens.

2

Click Create or select an existing PDN Profile policy and click Properties. The PDN Profile Policy, Global Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

As required, click on the appropriate tabs and click Search to determine to which network objects the PDN Profile Policy is applied.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the policy to NEs.

END OF STEPS

Part VI: Service management

Overview

Purpose

This part provides information about service management using the NFM-P.

Contents

70 Service management and QoS

70.1 Service management and QoS

70.1.1 Overview

The NFM-P supports the configuration of the following network services:

- virtual leased line (VLL) service
- virtual private LAN service (VPLS) / management virtual private LAN service (MVPLS) / hierarchical virtual private LAN service (HVPLS)
- internet enhanced service (IES)
- virtual private routed network (VPRN) service
- virtual LAN (VLAN) service
- composite service
- mirror service

An NFM-P operator can create, configure and delete services on sites and routers that are within their span of control.

An administrator can configure NFM-P system preferences related to service creation and deletion; see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

The benefits of the NFM-P service model include:

- service creation using configuration forms
- service management using configuration forms and a navigation tree
- service priority setting to manage service deletion: low priority services can be deleted by simply ticking a box in the confirmation dialog. Medium priority services can be deleted by confirming the priority. High priority services can only be deleted by a user with administrator rights.
- linking of services to create composite services that support complex customer applications
- policies that specify the NE traffic classification, policing, shaping, time of day restrictions, and marking
- traffic management capabilities to customize the delivery of different services according to SLAs
- integrated alarm management functions

Customer traffic enters a service through one or more access interfaces. In a local service, all access interfaces are on one NE. In a distributed service, multiple NEs are deployed at the PE. Customer traffic is transported across an IP/MPLS core network in unidirectional service tunnels that use GRE or MPLS LSPs. Many services can use the same tunnel.

Figure 70-1 Distributed VLL service

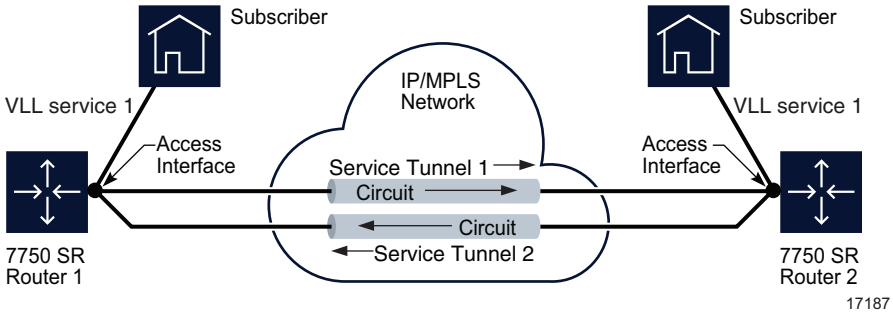

Figure 70-2 Distributed VPLS

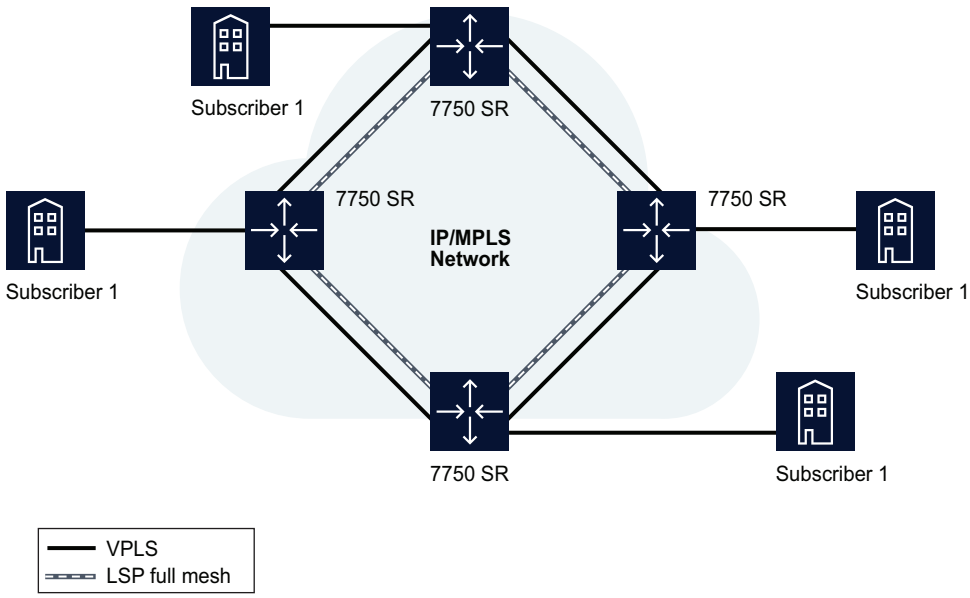

For IES, the managed devices are deployed at the provider edge and customer traffic enters the service using access interfaces. IES is a routed connectivity service where the customer communicates with an IP router interface to send and receive Internet traffic.

Figure 70-3 Sample IES

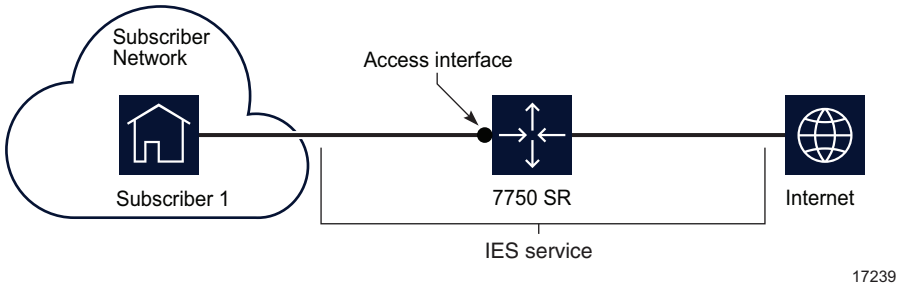

For VPRN services, the managed devices can be deployed as PE or provider core routers. Data and routing information are forwarded across an IP/MPLS service provider core network.

Figure 70-4 Sample VPRN service

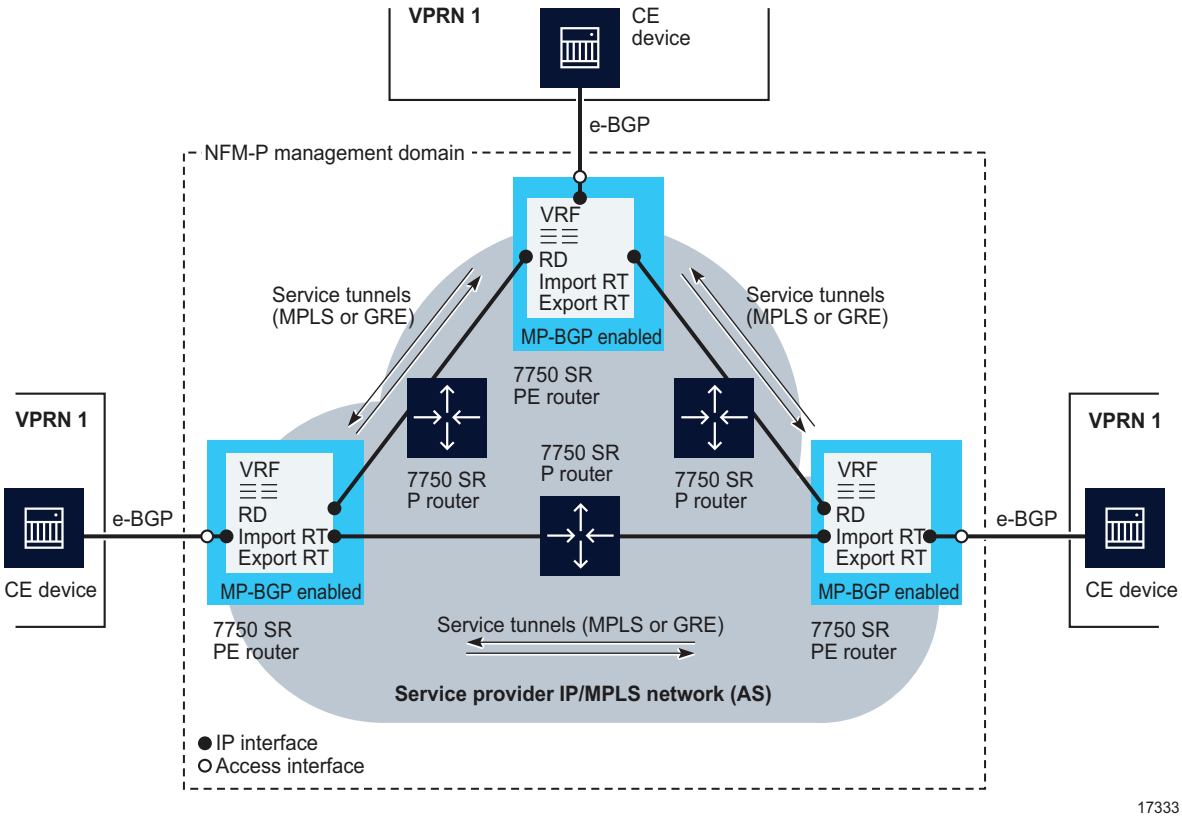

VLAN ring groups are used to send traffic across an Ethernet ring using copper or fiber optic connections from the source traffic device, for example, a 7450 ESS, to all devices in the ring. Any breaks in the physical links between devices are rerouted.

Figure 70-5 Sample VLAN configuration for L2 VPNs

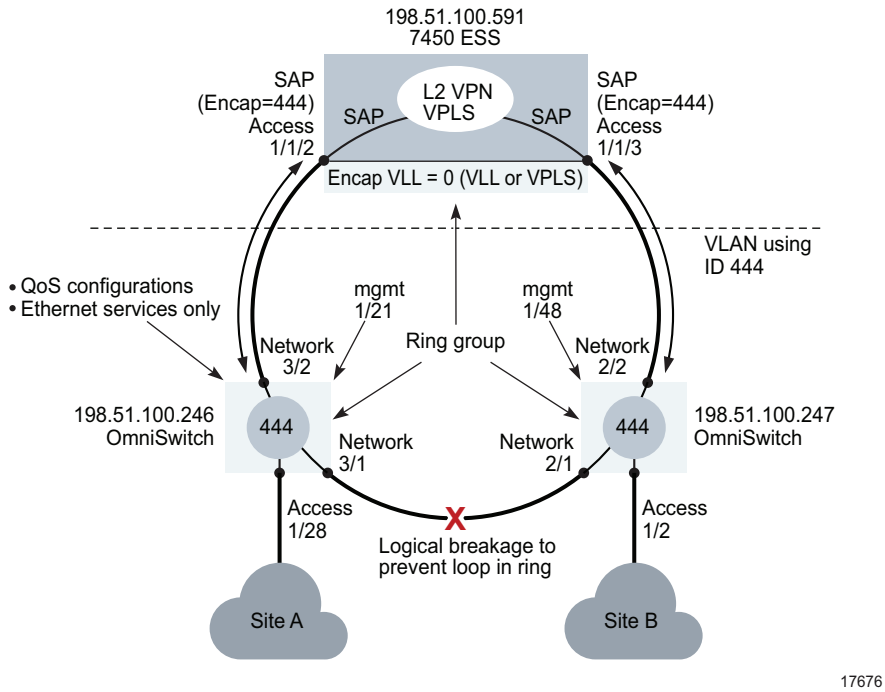

A composite service allows the interconnection of different service types to form a service delivery network that is tailored to a specific application. For example, VPRN, VPLS, and VLAN services can be joined to create a routed multicast TV distribution mechanism that spans a wide geographical area.

A mirror service is a unidirectional service function that duplicates a specified traffic stream and sends the duplicate stream to a monitoring device for troubleshooting or surveillance purposes. With pseudo-wire redundancy support, an ICB can be enabled in a mirror service spoke and a remote source, which can provide bidirectional service that enables support for active and standby PE redundancy.

QoS provides the ability to rate limit across multiple queues from one or more access interfaces for a customer, and to differentiate service levels for different types of traffic. For higher priority traffic such as VoIP or video, you can specify reserved bandwidth. Lower priority applications, such as data traffic, may not have reserved bandwidth but can burst to use all the available bandwidth.

The main elements of QoS are:

- QoS markings

Customer traffic may be marked with QoS markings, such as DSCP, EXP, and dot1p, that are mapped to forwarding classes.

All forwarding classes support profile marking of packets as in-profile or out-of-profile. In-profile packets have a high enqueuing priority. Out-of-profile packets have a low enqueuing priority. Profile marking of packets can occur at two points: when packets are classified into forwarding classes at access ingress and when packets are classified at service egress. Profile marking is only done on the internal header and not in an actual encapsulation.

- forwarding classes

  Provide network elements with a method to weigh the relative importance of packets, only in relation to other forwarding classes. A forwarding class is also referred to as a Class of Service.

- queues

  Location for buffering packets that are to be forwarded before they are scheduled.

- schedulers

  Hardware scheduling (or single-tier scheduling) exists by default on a device and consists of a high-priority and a low-priority scheduler.

  Scheduler policies (or multi-tier scheduling) provide a more complex, hierarchical structure of virtual schedulers that override the default hardware behavior for more flexible scheduling capabilities.

- slope policies

  Define the WRED slope characteristics of hardware buffer space that is used by the ingress and egress queues

See Chapter 49, “Policies overview” for more information about policies on the NFM-P. See the *7750 SR OS Services Guide* for more detailed information about QoS.

70.2 Access interfaces

70.2.1 Overview

Each customer service is configured with at least one access interface point called a SAP. The access interface identifies the point of customer interface for a service on the managed device.

A Layer 2 or Layer 3 access interface is uniquely identified using these parameters:

- physical Ethernet port or POS port and channel
- encapsulation type (if applicable)
- encapsulation id (if applicable)

Depending on the encapsulation type, a physical port or channel can have more than one access interface associated with it. Using encapsulation or a SONET/SDH channel, devices can support multiple services for a customer or for multiple customers.

Access interfaces can only be created on ports or channels that are designated as access in the physical port configuration. Access interfaces cannot be created on ports designated as core-facing network ports because these ports have a different set of features enabled in software.

Access interfaces can participate in policies. Time of day suites can also be associated with access interfaces to apply a set of time-based policies, filters, and schedulers. Configuration of access interfaces can be performed during service configuration or modification.

When you configure an access interface, consider the following:

- An access interface is owned by and associated with the service in which it is created.
- An access interface is a local entity and is locally unique to a specific device. The same access interface ID value can be used on another device.
- There are no default access interfaces. All access interfaces must be created.
- The default administrative state for an access interface at creation time is administratively enabled.
- If a port or channel is shut down (administratively or operationally), access interfaces on that port/channel are operationally out of service.

70.3 Automatic SDP (service tunnel) binding for services

70.3.1 Overview

You configure automatic service tunnel (SDP) binding for services when you configure a service using configuration forms.

The NFM-P defines its own internal rules on how automatic mesh SDP bindings are performed.

1. The NFM-P tries to find the least-used SDP (service tunnel) with the lowest load factor (the lower the number of bindings, the lower the load factor) from the source to destination NE when the service tunnel meets the following conditions:
   - the service tunnel operational state is up
   - the operational MTU is greater than or equal to the MTU value of the service site
   - T-LDP is set for SDP bindings of VPLS, IES, or VLL services
   - the selected transport method, either GRE, LDP, or RSVP
2. When no service tunnel operational state is up, NFM-P tries to find the least-used service tunnel with the lowest load factor in an operationally down state.
3. For mirror services, binding are created from the source sites to the destination sites.
4. When a service tunnel cannot be found for a service site, a change in the site leads to another search. For example, if the service MTU is 1500 and the highest path MTU in all SDPs from that site is 1472, no SDP binding can be successful. If the service MTU is lowered to 1472 or less, a successful SDP binding results.
5. When a service tunnel is not found, and the transport method selected is GRE, NFM-P attempts to create a GRE service tunnel with a path MTU equal to the service site service MTU with T-LDP signaling turned on.

See 16.5.9 “MTU size and port configuration” (p. 573) in 16.1.2 “Working with port and channel objects” (p. 565) for more information on MTU size considerations.

70.4 Multi-segment tunnel selection

70.4.1 Overview

Multi-segment tunnel selection functionality enhances tunnel selection functionality to allow the system to select intermediate NEs in a path, as well as the service tunnels from a selected intermediate NE to its next hop. Multi-segment tunnel selection functionality is used to select existing SDPs. It cannot create new SDPs.

Tunnels are selected based on the following criteria:

- If a metric value is specified for any of the SDPs along the path, the path with the lowest sum of the metric property for all SDPs along the path is selected.
- If no metric is specified for any of the SDPs along the path, the path with the fewest number of SDPs (hops) is selected.
- If two paths or SDPs are equal based on the selection criteria mentioned above, the SDP/path with the lowest load factor (number of services using that SDP) is selected.

Multi-segment tunnel selection functionality must be enabled on the Services tab on the NFM-P System Preferences form, as well as at the tunnel selection profile level. The table below describes the objects to which a tunnel selection profile with multi-segment tunnel selection can be assigned, and the service tunnel configurations that result.

Table 70-1 Multi-segment tunnel selection scenarios

Tunnel selection profile applied to...	Resulting tunnel configuration
Spoke SDP binding on a VLL service	This scenario applies to a VLL service with two terminating sites. One of the sites has a spoke SDP binding configured with Auto-Select Transport Tunnel enabled, and a tunnel selection profile with Multi-segment Tunnel Selection enabled. If a valid path is found between the two terminating sites, the spoke SDP bindings and switching sites are created.
VLL service	This scenario applies to a VLL service with two terminating sites. The service is configured with Automatic SDP Binding/PBB Tunnel Creation enabled, and a tunnel selection profile with Multi-segment Tunnel Selection enabled. If a valid path is found between the two terminating sites, the spoke SDP bindings and switching sites are created.

Table 70-1 Multi-segment tunnel selection scenarios (continued)

Tunnel selection profile applied to...	Resulting tunnel configuration
Spoke termination point	This scenario applies to a spoke SDP binding between any two supporting objects. These include (M-)VPLS sites, (M-)VPLS I-sites, (M-)VPLS B-sites, VLL sites, IES L3 access interfaces, VPRN L3 access interfaces, VLL endpoints, and VPLS endpoints. The spoke SDP binding is configured with Auto-Select Transport Tunnel enabled, and a tunnel selection profile with Multi-segment Tunnel Selection enabled. If a termination point is specified, the spoke SDP bindings and return spoke SDP bindings are created. Switching sites are created as follows: • If the spoke SDP binding is created on a VLL service, the switching sites are created on the VLL service. • If the specified termination point is on a VLL service, the switching sites are created on the VLL service. • If neither end of the tunnel is on a VLL service, a new EPIPE service is created. The service component ID is auto-selected by the system. The subscriber ID is the same as the service on which the spoke is created. All other parameters are default values.

70.4.2 Redundant path selection

If Redundant Path Selection is enabled on a tunnel selection profile, a redundant path is created. Redundant Path Selection can only be enabled if Multi-Segment Tunnel Selection is also enabled. The selected redundant path does not use any of the same switching sites as the primary path. Redundant path selection is not supported when a termination point is selected on a spoke SDP binding. The table below describes the objects to which a tunnel selection profile with redundant path selection can be assigned, and the service tunnel configurations that result.

Table 70-2 Redundant path selection scenarios

Tunnel selection profile applied to...	Resulting redundant path configuration
Spoke SDP binding on a VLL service	This scenario applies to a VLL service with two terminating sites. One of the sites has a spoke SDP binding configured with Auto-Select Transport Tunnel enabled, and a tunnel selection profile with Multi-segment Tunnel Selection and Redundant Path Selection enabled. If a valid path is found between the two terminating sites, the spoke SDP bindings and switching sites for the redundant path are created. In this case, an endpoint and a return endpoint must be specified.

Table 70-2 Redundant path selection scenarios (continued)

Tunnel selection profile applied to...	Resulting redundant path configuration
VLL service	This scenario applies to a VLL service with two terminating sites. The service is configured with Automatic SDP Binding/PBB Tunnel Creation enabled, and a tunnel selection profile with Multi-segment Tunnel Selection and Redundant Path Selection enabled. If a valid path is found between the two terminating sites, the spoke SDP bindings and switching sites for the redundant path are created. From a VLL Service, endpoints on the terminating sites that do not have SAPs or spokes are used for the auto-created spokes. If multiple endpoints exist on a terminating site, the endpoint without any SAPs is used. If multiple viable endpoints are available, an error is raised by the NFM-P and the operation fails. If there are no endpoints available, an endpoint is automatically created with default parameter values, with the following exceptions: • The Name parameter value is the endpoint Service ID. • The Description parameter value is auto-generated by the system.

70.4.3 Service tunnel required bandwidth

Service tunnel bandwidth requirement can be specified as an additional constraint in the path search process. When required bandwidth is specified at the service or service site level, the selection of service tunnels is based on bandwidth availability. If all other tunnel selection criteria are equal, the tunnel with higher available bandwidth is selected.

Note: Service tunnel required bandwidth functionality requires that the CPAM feature set is correctly configured. See the *NSP NFM-P Control Plane Assurance Manager User Guide* for information about the CPAM feature set. Additionally, the Multi-Segment Tunnel Selection and Service Bandwidth Management system options must be enabled.

All physical links between service sites must be configured (either manually or through LLDP) before bandwidth can be allocated to service tunnels.

This functionality applies to RSVP tunnels on VLL, VPLS, and MVPLS services, and G.8032 Ethernet rings on VPLS services. Bandwidth requirement criteria are specified at the service level, and are inherited by all associated service sites by default. Optionally, bandwidth requirement criteria can be specified on individual service sites. Bandwidth requirements are specified on the service or site configuration form, Bandwidth tab - Required Bandwidth sub-tab.

Bandwidth can be user-specified on a per-CoS basis, or it can be calculated automatically. When bandwidth values are calculated automatically, the calculation is based on the ingress CIR configurations, taking the sum of the bandwidths for the endpoints (SAPs) connected to each site. In the event that two sites with different bandwidth sums are connected by a service path, the lower of the two bandwidth sums is booked for the service path.

There are two options that determine how service bandwidth is booked in the network:

- For Epipe services only, the user has the option to enable bandwidth booking on both the active and redundant service paths (if a redundant path is available). Otherwise, only the active path is booked.

- For all VLL, VPLS, and MVPLS services, the user has the option to enable bandwidth booking on all operationally up LSP paths in use by the service. Otherwise, only the active path is booked.

70.5 Automatic PBB tunnel binding

70.5.1 Overview

This feature is applicable only to VLL Epipes and I-VPLS services using a PBB tunnel.

The configuration of the service CAC functionality allows you to manage bandwidth at the service level and at the link level, which in turn enables the NFM-P to automatically select the best tunnel based on the number of active links and on the amount of available bandwidth on the service. The NFM-P also considers the shortest available path when it allocates bandwidth.

Although the NFM-P can select tunnels automatically—see 70.3 “Automatic SDP (service tunnel) binding for services” (p. 1910) —if service CAC is not configured, bandwidth will not be considered by the NFM-P when selecting the tunnel.

Bandwidth availability for the tunnel is calculated at the service admission request time and is based on the links currently used by the tunnel forwarding path. The booking of reserved bandwidth by the service is done directly on the link and directly on the tunnel, for a bandwidth-reserved tunnel.

The link monitors the bandwidth used by all of the services using that link. Since the bandwidth reserved on each tunnel that is used by a service can be different, the link monitors bandwidth usage through the tunnels rather than through a service.

You can perform a system-wide audit to ensure that the bandwidth on each link is properly calculated. The audit visits all tunnels in the network and adjusts the available bandwidth in all links currently being used by the forwarding tunnel path. The adjustments are made based on total bandwidth requests per CoS for all EVPLs currently using the tunnels.

You can perform this audit by clicking on the CAC Audit button on the Manage Services form, or by choosing CAC Audit on the More Actions button. The audit also occurs automatically on system startup and when service CAC is switched from disabled to enabled.

70.6 Lightweight SAPs

70.6.1 Overview

SAPs that are associated with residential split horizon groups on VPLS sites are called lightweight SAPs. An RSHG uses dual-pass queue optimization and does not support downstream broadcast or multicast traffic. Lightweight SAPs have fewer internal configuration settings than regular SAPs. Therefore, you can create more lightweight SAPs on a NE. The SAP Lightweight property is automatically set at creation time, when the SAP is associated with an RSHG, and cannot be modified later.

Users can:

- assign SAPs to residential split horizon groups
- list lightweight SAPs using the Lightweight filter property

Note: To get the most recent state of a lightweight SAP, it is recommended that users perform a resynchronization by clicking on the Resync button in the SAP configuration form. If this button is not visible, click on the More Actions button and choose Resync.

70.7 QoS policies

70.7.1 Overview

Policies group and manage the various QoS elements used to determine how traffic is routed.

- access ingress policies

  Specify how QoS marking is interpreted, how customer traffic is mapped into queues, and how queues are classified.
- access egress policies

  Specify how customer traffic is mapped into queues, specify queue classification, queue parameters, and QoS marking.
- network policies

  Specify QoS marking to forwarding class mapping on ingress and QoS marking to forwarding class mapping on egress.
- network queue policies

  Specify CIR, PIR, and burst sizes for each queue. Forwarding class to queue mapping is not configurable.
- scheduler policies

  Specify custom settings and a hierarchical structure of virtual schedulers to replace the default hardware schedulers on the device.
- port scheduler policies

  Specify bandwidth allocation at the egress port level.
- HSMDA scheduler policies

  Specify schedulers to define egress port and ingress scheduler behavior on an HSMDA.
- slope policies

  Specify WRED settings to customize how in-profile and out-of-profile traffic is processed in hardware buffers, applied to daughter cards or ports.
- HSMDA slope policies

  Specify settings for controlling how the depth of HSMDA queues is managed.
- ATM QoS policies

  Specify ATM QoS settings to customize ATM traffic parameters including service category and shaping.

The figure below shows where QoS policies are applied at the service level.

Figure 70-6 Service-level view of policies

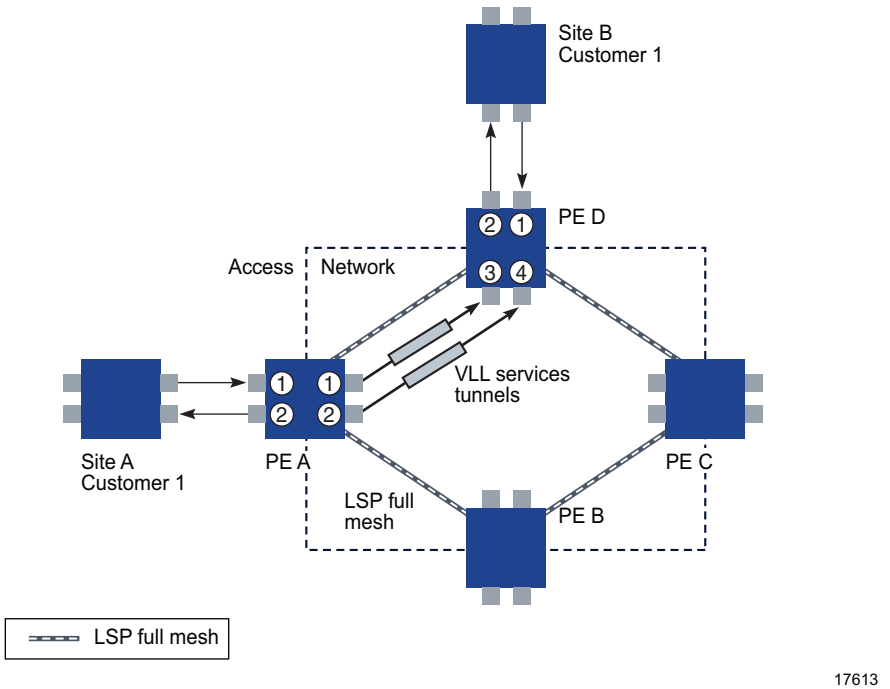

The figure below shows where policies are applied on a device with respect to access and network ingress and egress traffic.

Figure 70-7 Types of traffic on a device and applied policies

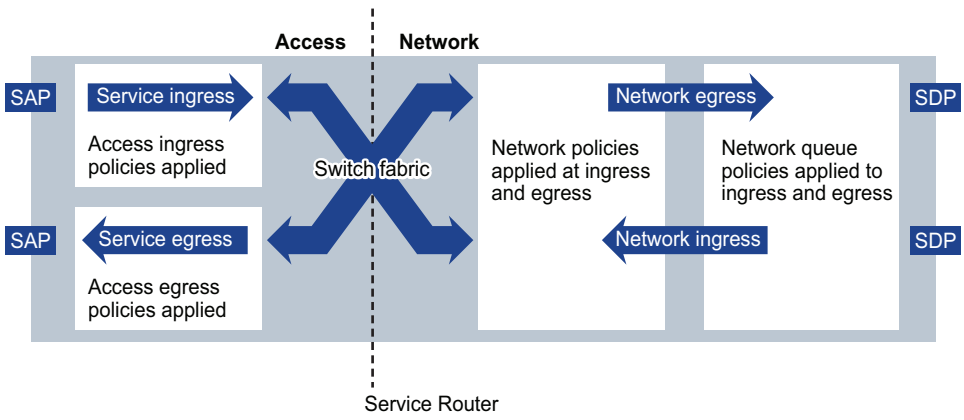

70.7.2 SAP access ingress policies

Access ingress policies are applied to access interfaces and specify QoS characteristics on ingress. Participation in access ingress policies is defined when access interfaces are configured or modified.

Access ingress policies include:

- mapping of QoS marking, such as dot1p, DSCP, and precedence, and IP/MAC address information to forwarding classes
- forwarding class definitions and mapping to queues
- queue definitions and mapping to schedulers

Figure 70-8 Access ingress policy elements

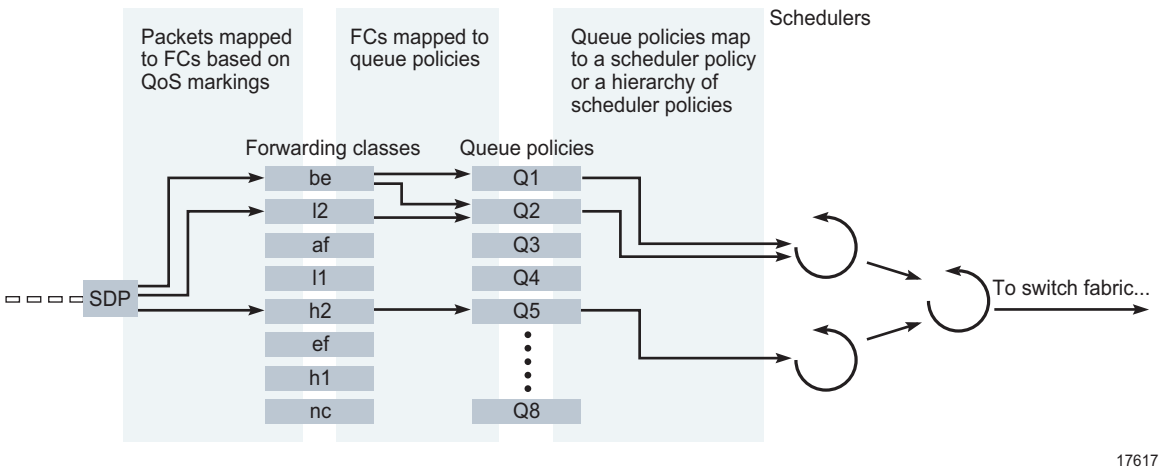

See 50.2 “SAP access ingress policies” (p. 1507) in “QoS policy types” (p. 1507) for more information about access ingress policies.

70.7.3 SAP access egress policies

Access egress policies are applied to access egress interfaces and specify QoS characteristics on egress. Participation in access egress policies is defined when access interfaces are configured or modified.

Access egress policies include:

- forwarding class definitions and mapping to queues
- queue definitions and mapping to schedulers

Figure 70-9 Access egress policy elements

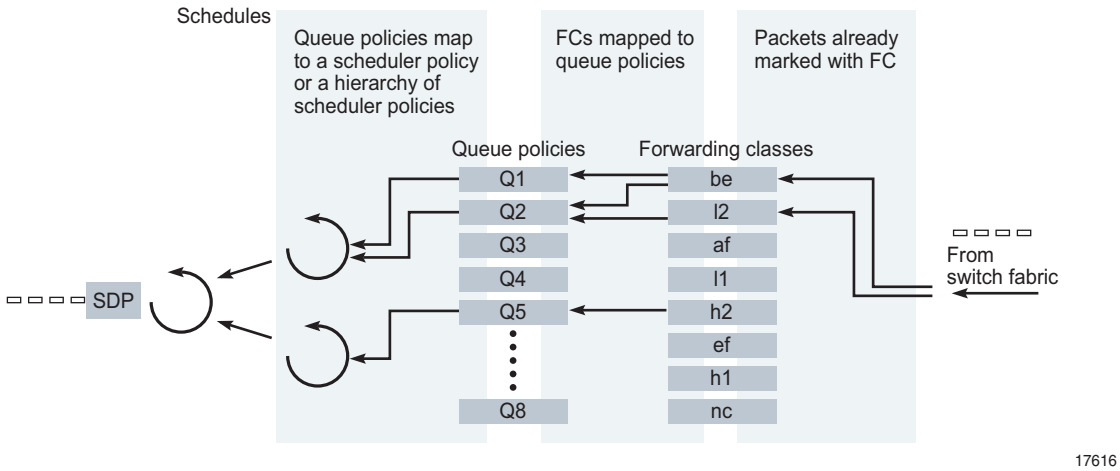

In Figure 70-9, “Access egress policy elements” (p. 1918) , packets are marked with an FC, either by:

- ingress policy if the packet was received on the same device
- in the tunnel transport encapsulation if received using a service tunnel

See 50.3 “SAP access egress policies” (p. 1513) in “QoS policy types” (p. 1507) for more information.

70.7.4 Network policies

Network policies define egress QoS markings and ingress QoS interpretation for traffic on core network IP interfaces.

A network policy defines:

- DSCP name mapping to forwarding classes
- LSP EXP value mappings to forwarding classes
- whether QoS remarking is enabled

The figure below shows the sequence of how the elements of network and network queue policies are applied at ingress.

Figure 70-10 Network and network queue policy elements on ingress

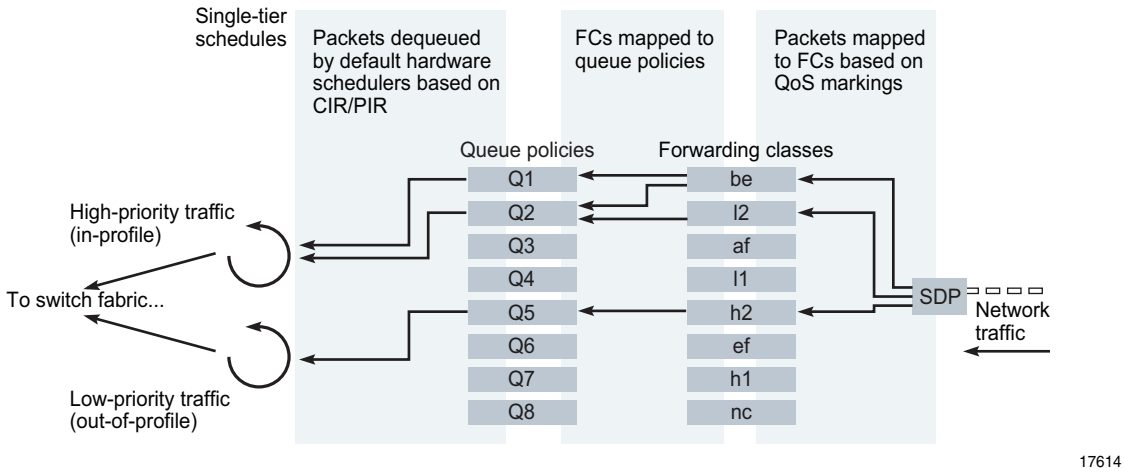

The figure below shows the sequence of how the elements of network and network queue policies are applied at egress.

Figure 70-11 Network and network queue policy elements on egress

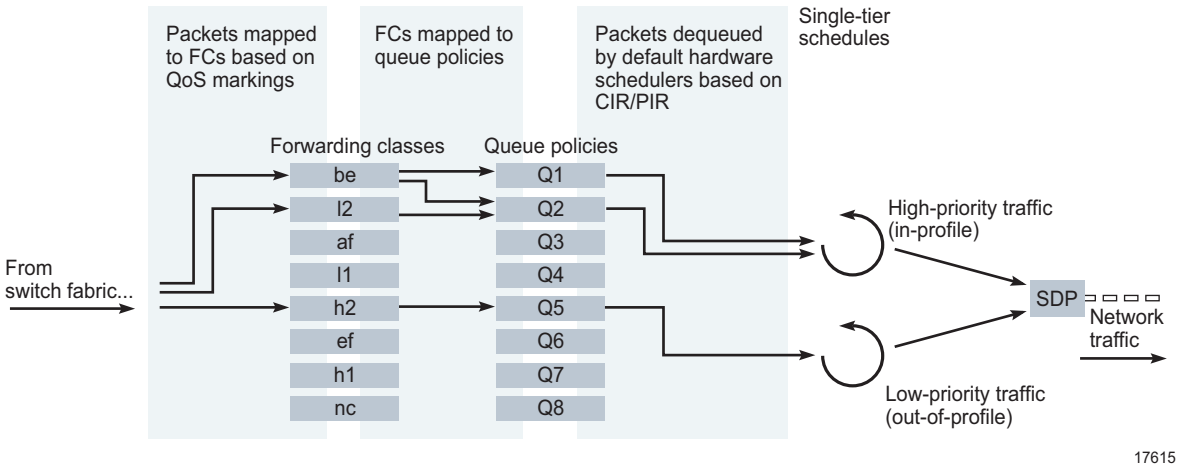

See 50.8 “Network policies” (p. 1516) in “QoS policy types” (p. 1507) for more information.

70.7.5 Network queue policies

Network queue policies define the network forwarding class queue characteristics for core egress network ports and for ingress on MDAs.

See Figure 70-10, “Network and network queue policy elements on ingress” (p. 1919) and Figure 70-11, “Network and network queue policy elements on egress” (p. 1919) for the sequence of how the different elements of network and network queue policies are applied on ingress in and on egress.

The queued packets are serviced by single-tier schedulers on the device and are forwarded to a single destination switch fabric port or a network interface.

See 50.9 “Network queue policies” (p. 1516) in “QoS policy types” (p. 1507) for more information about network queue policies.

70.7.6 Scheduling

Scheduling defines the order and method for how packets which are enqueued in different queues, are dequeued. Ingress schedulers control the data transfer between the queues and the switch fabric. Egress schedulers control the data transfer between the egress queues and the switch fabric. Packets are not actually forwarded to schedulers, but are forwarded from the queues directly to ingress and egress interfaces. Participation in scheduler policies is defined when access interfaces are configured or modified.

There are two types of scheduling:

- Single-tier scheduling is the hardware-based default method for scheduling queues on the device. There are no configurable parameters for single-tier schedulers. When a scheduler policy is not specified for an access interface, rate limiting is specified by the values specified in the queue and scheduling is performed by the default hardware scheduler on the device. Single-tier scheduling consists of a pair of scheduling priority loops in the 7750 SR and bases scheduling on the CIR and the PIR set in the queue policy. One loop is for scheduling high-priority (in-profile) traffic, and the other loop is for low-priority (out-of-profile) traffic.
- Virtual hierarchical (multi-tier) scheduling can provide more flexible scheduling for access ingress and egress interfaces, and determine how queues are scheduled. They are defined using a Scheduler policy, and can be configured to override default hardware scheduling. You can create up to three tiers of virtual schedulers.

  Aggregation schedulers are used to share a scheduler policy across a number of ports or daughter cards. This can be useful when a number of ports or cards are dedicated to the same customer. See 50.13 “Scheduler policies” (p. 1521) in “QoS policy types” (p. 1507) for more information about configuring aggregation schedulers.

70.7.7 Port scheduler policies

Port scheduler policies define the bandwidth allocation based on the available bandwidth at the egress port level. A port scheduler policy manages a bandwidth allocation algorithm that represents a virtual multi-tier scheduling hierarchy.

The port scheduler allocates bandwidth to each service or subscriber that is associated with an egress port. Egress queues on the service may have a child association with a scheduler policy on the SAP or multi-service site. All queues must compete for bandwidth from an egress port.

There are two methods of bandwidth allocation on the egress access port:

- **direct association of port scheduler on a SAP or multi-service site with service or subscriber queue**

A service or subscriber queue is associated with a scheduler on the L2 access interface or multi-service site, and the service-level scheduler policy is associated with a port level scheduler.

- **direct association of port scheduler with service or subscriber queue**
  A service or subscriber queue is associated with a port scheduler. The port scheduler hierarchy allocates bandwidth at each priority level to each service or subscriber queue.

See 50.14 “Port scheduler policies” (p. 1523) in “QoS policy types” (p. 1507) for more information about configuring port scheduler policies.

70.7.8 HSMDA scheduler policies

The port-based scheduler manages forwarding for each egress port on the HSMDA. Each port-based scheduler maintains up to eight forwarding levels. Eight scheduling classes contain each of the queues that are assigned to the port scheduler. Membership in a scheduler is defined by the queue identifier.

The port-based scheduler supports a port-based shaper that is used to create a sub-rate condition on the port. Each scheduling level can be configured with a shaping rate to limit the amount of bandwidth allowed for that level.

The port-based scheduler allows two weighted groups (group 1 and group 2) to be created. Each group can be populated with three consecutive scheduling classes.

See 50.15 “HSMDA scheduler policies” (p. 1523) in “QoS policy types” (p. 1507) for more information about configuring slope policies.

70.7.9 Slope policies

Slope policies manage how shared buffers are utilized on the SR. When traffic is queued, the WRED slope parameters in the slope policy determine how the traffic is buffered for de-queuing, as shown in the figure below.

Figure 70-12 Slope policy characteristics

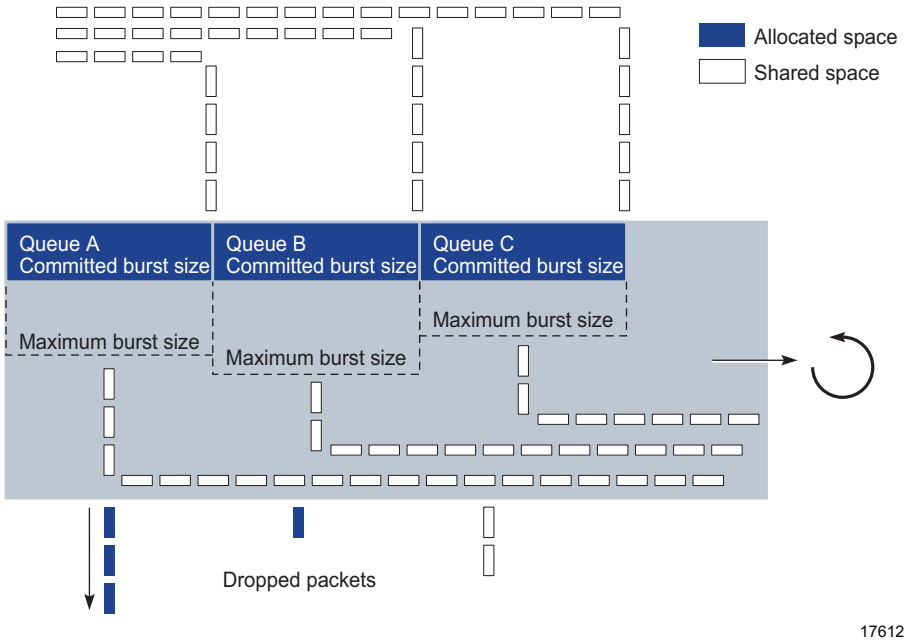

All queues are in contention for shared buffer space when they exceed their CBS, and can use their MBS, when space is available in the shared buffer space. The WRED parameters determine whether a packet is discarded or not, and, as a result, determine whether the packet is dequeued. When the shared buffer space exceeds or approaches the maximum percentage defined by the WRED configuration, packets are discarded.

By using two independent slope policy configurations, one for in-profile traffic and one for out-of-profile traffic, you can configure in-profile traffic to receive preferential treatment over out-of-profile traffic.

See 50.11 “Slope policies” (p. 1519) in “QoS policy types” (p. 1507) for more information about configuring slope policies.

70.7.10 HSMDA slope policies

HSMDA slope policies control the HSMDA queues. Each queue supports an index for an HSMDA slope policy table. Each policy in the table consists of two RED slopes (one high priority and another for low priority) to manage queue congestion. HSMDA RED slopes operate on the instantaneous depth of the queue.

A packet that attempts to enter a queue triggers a check to see whether the packet is allowed based on queue congestion conditions. The packet contains a congestion-priority flag that indicates whether the HSMDA is to use the high or low slope. The slope policy containing the slope is derived from the policy index in the queue configuration parameters on the HSMDA.

The RED slope discards are based on the current queue depth before a packet is allowed into the queue. Therefore, a queue may consume buffers that are greater than the configured MBS value

based on the size of the packet. After the packet maximum is reached, packets that are associated with the queue are discarded. When the schedule removes packets from the queue, the queue depth decreases, which eventually lowers the depth of the threshold.

See 50.12 “HSMDA WRED slope policies” (p. 1520) in “QoS policy types” (p. 1507) for more information about configuring HSMDA slope policies.

70.8 The triple play service delivery architecture

70.8.1 Overview

The triple play service delivery architecture (TPSDA) is based on three major components, as shown in the figure below:

- broadband service aggregator
- broadband service router
- service and policy activation

Figure 70-13 TPSDA components

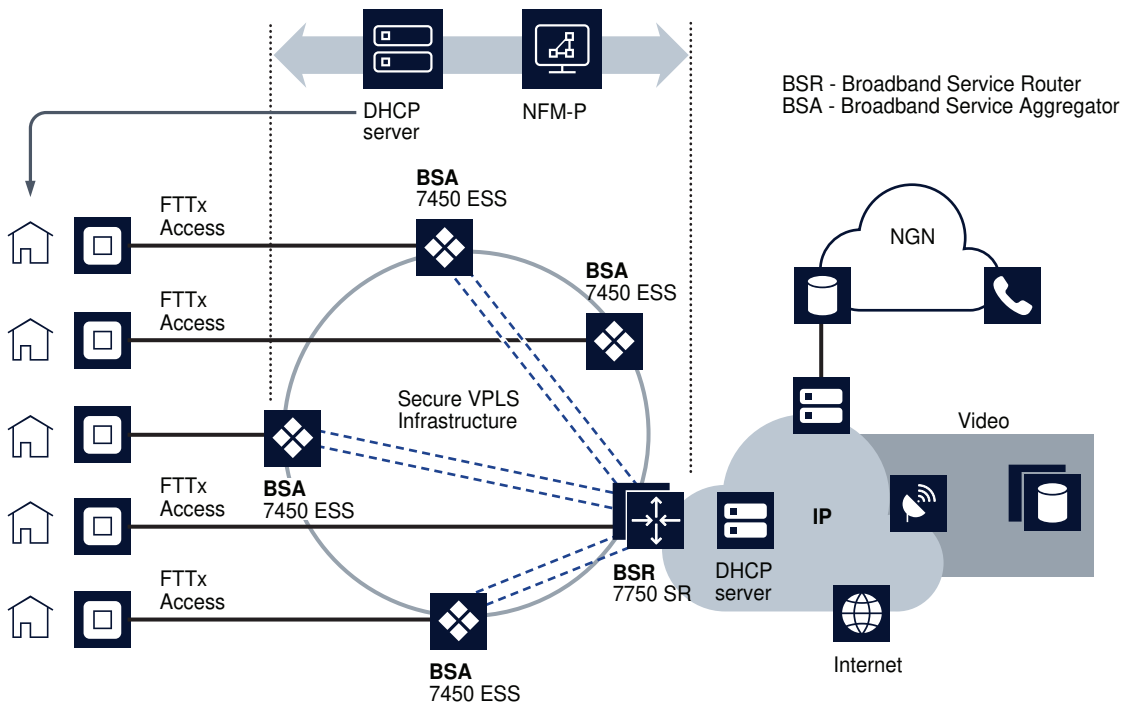

The table below lists the Nokia TPSDA product components.

Table 70-3 Product components of the TPSDA

Product	Role	Notes
NFM-P	Provides network, service, and policy management across the TPSDA architecture, including a unified interface for element management and simple service activation and monitoring.	—
7750 SR	BSR	Support per-service and per-content type differentiation of QoS levels and supports distribution of multicast traffic.
7450 ESS	BSA	Aggregate traffic from DSLAMs and other FTTx access devices that are connected to end-user residential gateways.

DSLAMs or other access NEs are connected to Ethernet access ports on the broadband service aggregator. Typically, a single VLAN for each customer is set up between the access NE and the BSA. This a configuration enables the application of consistent per-customer policies, such as QoS, filtering, and accounting, to be applied on the BSA.

Scaling of traffic and services is done by dividing L2 and L3 functions between the BSA and the BSR. The BSA aggregates traffic over Gigabit Ethernet ports and performs per-customer service queuing, scheduling, accounting and filtering, as described later in this chapter. The BSR terminates L2 access and routes over IP/MPLS with support for all protocols, including multicasting. Time of day QoS policies can be applied using NFM-P policy management.

Interconnectivity between BSAs and BSRs is provided by VPLS. VPLS instances can be automatically established using hub-and-spoke or ring topologies. Both can be configured and sites added to the VPLS using the NFM-P. Regardless of the fiber plant layout, VPLS enables a full mesh between all sites that are receiving and distributing customer traffic in the TPSDA, ensuring efficient transport and protection from NE or fiber failures.

VPLS also provides mechanisms for traffic security, including residential split horizon groups in which direct user bridging is prohibited; ARP and broadcast suppression; DHCP-populated MAC and IP address filtering to prevent denial or service and theft of service using DHCP snooping, and RADIUS or TACACS+ authentication.

70.9 Service differentiation and QoS

70.9.1 Overview

This TPSDA approach provides a model based on call admission for video and VoIP, with the need to guarantee delay, jitter, and loss characteristics after the service connection is accepted. The architecture also meets the different QoS needs of HSI, namely per-user bandwidth controls, including shaping and policing functions that have little or no value for video and VoIP service delivery.

In conjunction with the architecture's support for content differentiation, this approach enables differentiated service pricing within high-priority data packages, also known as HSI. The distribution of QoS policy and enforcement across BSA and BSR allows the service provider to implement

meaningful per-user service level controls. Sophisticated and granular QoS in the BSA allows the service provider to deliver truly differentiated IP services differentiation based on the user as well as on the content.

In the BSR to BSA downstream direction, IP services rely on IP layer classification of traffic from the network to queue traffic appropriately towards the BSA. Under extreme loading (only expected to occur under network fault conditions), lower priority data services and/or HSI traffic are compromised to protect video and voice traffic. Classification of HSI traffic based on source network address or IEEE 802.1p marking allows the QoS information to be propagated to upstream or downstream devices.

The BSR performs service distribution routing based on guarantees required to deliver the service and associated content, rather than on individual end users. The BSR only needs to classify content based on its forwarding class for a specific BSA to ensure that traffic for each service receives the appropriate treatment towards the BSA.

Figure 70-14 TPSDA downstream QoS configurations

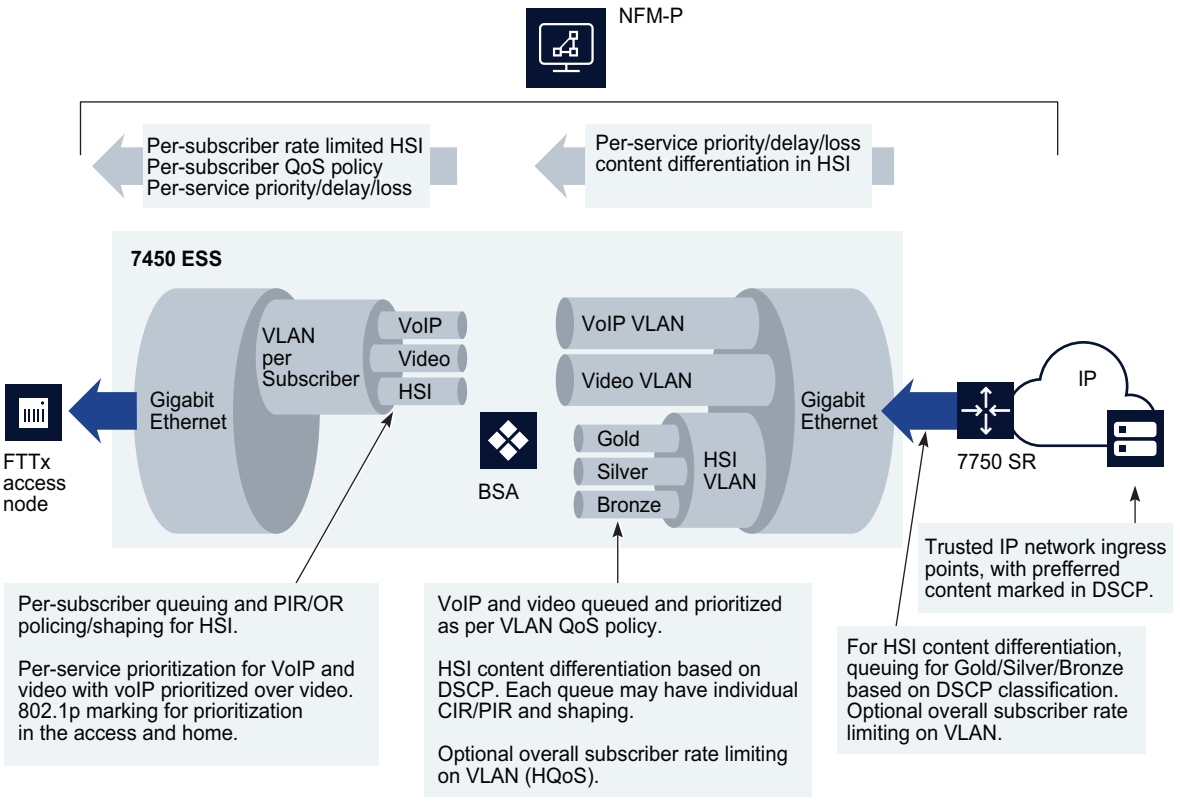

In the BSA-to-BSR upstream direction, traffic levels are substantially lower. Class-based queuing is used on the BSA network interface to ensure that video traffic is forwarded with minimal delay and that preferred data or HSI high-priority data traffic services receive better treatment than for best-

effort Internet traffic. The IP edge device (BSR) therefore does not need to enforce per-user policies for hundreds of thousands of users. This function is distributed to the BSAs, and the per-user policies can be implemented on the interfaces directly facing the access NEs.

The BSA is capable of scheduling and queuing functions on a per-service, per-user basis, in addition to performing wire-speed packet classification and filtering based on both L2 and L3 interfaces. In addition to per-service rate limiting for Internet services, service traffic for each user can be rate-limited as an aggregate using a bundled service policy created using the NFM-P. These functions allow different users to receive different service levels independently and simultaneously.

Figure 70-15 TPSDA upstream QoS configurations

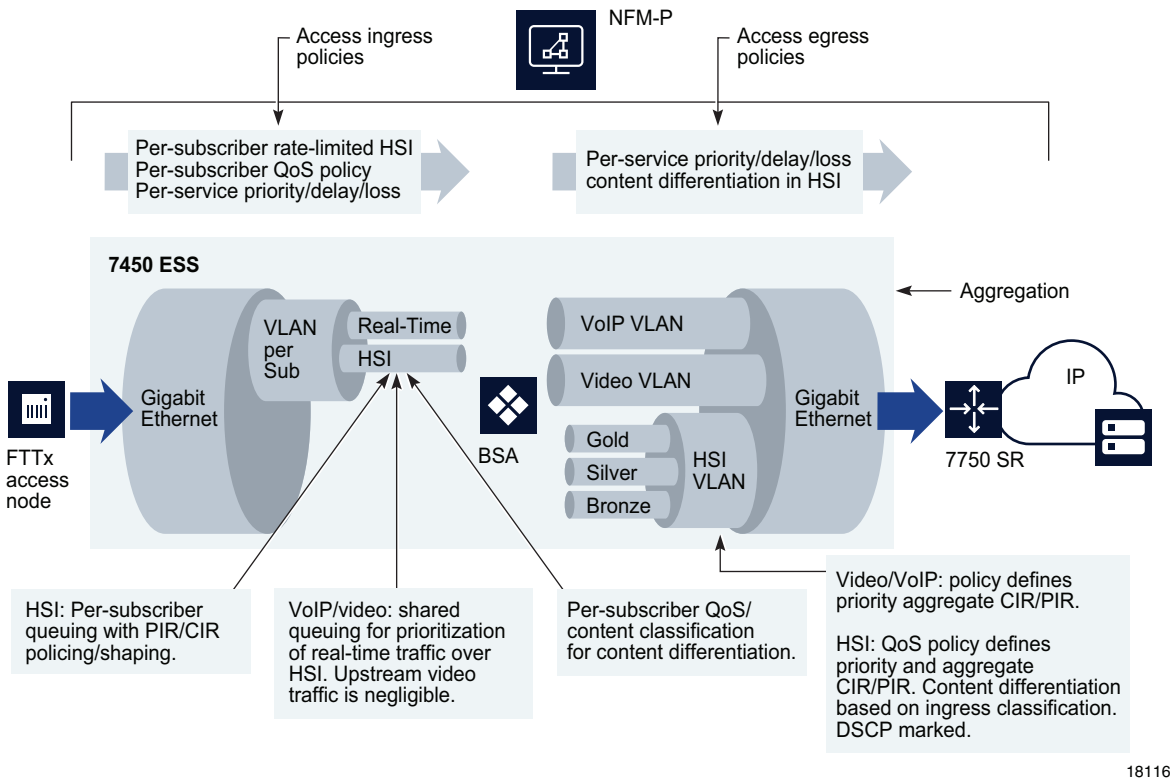

When a residential host device such as a residential gateway or a set-top box in the customer's home starts up, it requests network information, including the required IP address from a DHCP server. The figure below shows IP address assignment. See Table 70-4, “TPSDA features” (p. 1927) for information about DHCP configuration options in the TPSDA.

Figure 70-16 DHCP IP address assignment in the TPSDA

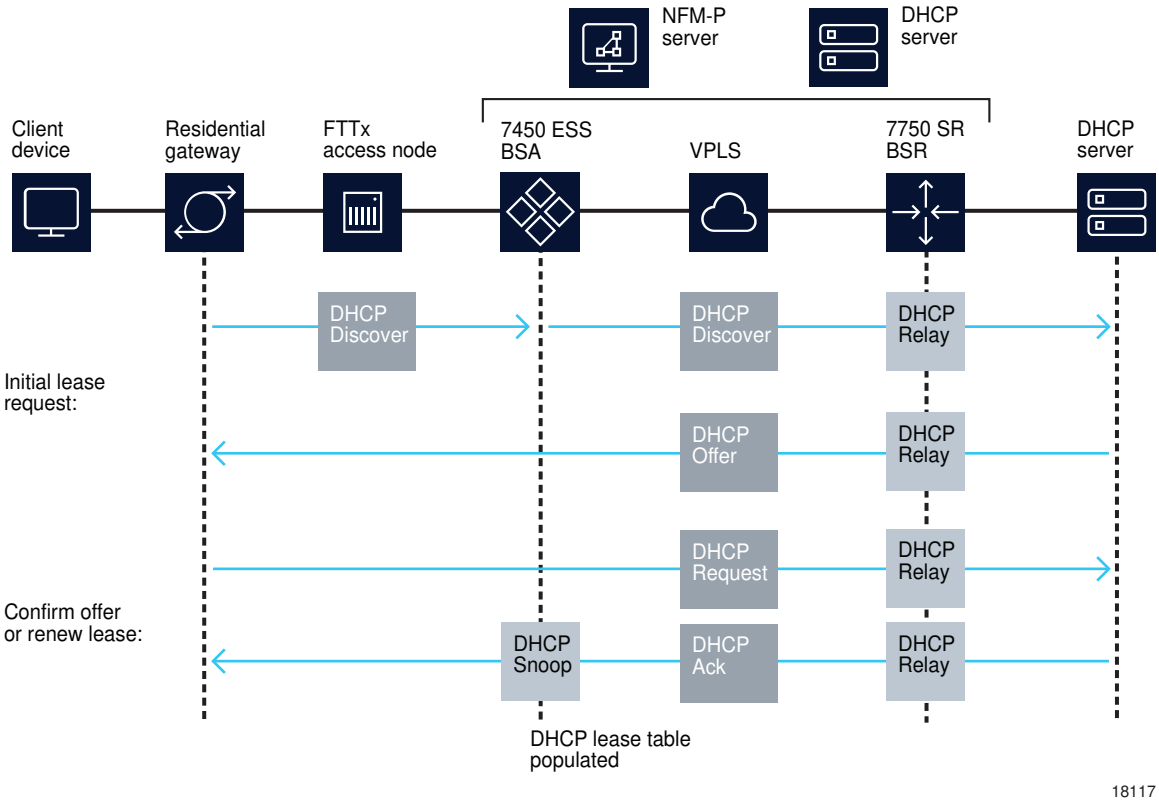

The table below lists the TPSDA features that you can configure using the NFM-P.

Table 70-4 TPSDA features

Feature and use	Notes	Reference
Split horizon groups		
For the TPSDA, there can be no user-to-user communication in the BSA; instead, all communication is done through the BSR. This residential bridging is done using split horizon groups, which ensures that traffic from different SAPs in the same service are not forwarded to other SAPs or spokes.	Traffic arriving on a spoke service tunnel or SAP within the split horizon group is not copied to other SAPs or spoke service tunnels. Traffic is copied to SAPs and spokes in other split horizon groups existing within the same service, such as a VPLS.	See Chapter 77, "VPLS management" for configuration of split horizon groups.
DHCP		

Table 70-4 TPSDA features (continued)

Feature and use	Notes	Reference
For the TPSDA, host devices, such as a residential gateway, SIP phone, or set-top box, use DHCP to obtain IP address and other network configuration information. The client device sends a DHCP discover message to request an IP address. The sequence of events is shown in Figure 70-16, "DHCP IP address assignment in the TPSDA" (p. 1927) .	Information added to the DHCP discover requests may include information added by the FTTx access NE or the BSA, for example, the shelf, slot, port, VPI, VCI, or other identifier of the host. You can use the NFM-P to configure DHCP relay on the first IP interface in the upstream direction. The BSA or BSR relays the message to a DHCP server. The gateway (residential gateway) IP address indicates to the DHCP server the subnet an IP address should be allocated to for the host.	See the appropriate service configuration chapter, as DHCP, option 82, and DHCP relay are configured at the service level.
DHCP relay DHCP discover messages are broadcast packets that typically do not leave the IP subnet. DHCP relay agents intercept the requests and forward them as unicast messages to a DHCP server. DHCP request messages from subscriber hosts are usually sent from the FTTx access NE, with information appended to uniquely identify the residential gateway, either by MAC address of the residential gateway or by an option 82 string identifier that indicates the device, port type, rack, shelf, slot, port, and VLAN ID or VPI/VCI of the circuit connected to the residential gateway.	The DHCP relay agent sets the GIADDR in the packet to the IP address of the ingress interface (SAP). You must configure the BSA and BSR devices as DHCP relay agents when the DHCP requests are going to be forwarded to a DHCP server. The maximum DHCP relay packet size is 1500 bytes.	See the appropriate L3 service (IES and VPRN) or L2 service (VPLS) configuration chapter.
DHCP lease state table The BSA maintains the identities of hosts that are allowed network access for each SAP on each service.	The lease state information is collected by snooping DHCP acknowledge messages on the SAP, using DHCP snooping. Entries in the DHCP lease state table remain valid for the duration of the IP address lease.	See Chapter 16, "Port and channel object configuration" .
DHCP snooping The BSA can copy DHCP packets and inspect them to help secure the system. For example, if malicious user A sends an IP packet requesting a video stream intended for user B, return packets sent to user B could jam B's connection.	Use the NFM-P to configure DHCP snooping for the following purposes: • To insert Option 82 information when the system is not configured for DHCP relay by enabling DHCP snooping on the SAP closest to the host. • To build a DHCP lease state table by enabling DHCP snooping on the service tunnel nearest the network egress and on the SAP closest to the host. • To efficiently associate dynamic hosts with subscriber instances and associated network resources in a triple play service configuration	See the appropriate service chapters for configuration of the lease populate and snooping parameters. See Chapter 74, "Residential subscriber management" for information about using DHCP snooping for subscriber identification purposes.

Table 70-4 TPSDA features (continued)

Feature and use	Notes	Reference
Option 82 The DHCP relay option allows managed devices to append information to the DHCP request that identifies where the DHCP request originated. You can also independently insert Option 82 information when DHCP snooping is enabled on a VPLS SAP. The Option 82 information can be: • The DHCP Option 82 string circuit ID value associated with the 7330 ISAM FTTN, or other ISAM family of NEs. The *device port_type rack/shelf/slot/port: VPI:VCI* identifier on the 7330 ISAM FTTN indicates that this is the connection to the residential gateway. • The DHCP Option 82 string remote ID value associated with the 7330 ISAM FTTN, or other ISAM family of NEs.	Using Option 82, you can identify: • the circuit ID (service tunnel binding) that is unique to the device relaying the circuit • the remote ID (MAC address) of the host at the far end of the circuit • the subscriber to which a host belongs for the purpose of assigning network resources The maximum DHCP relay packet size is 1500 bytes. If adding Option 82 information to the packet causes the packet to exceed 1500 bytes, the DHCP relay request is forwarded without including the Option 82 information.	See the appropriate service configuration chapter. For DHCP option 82 information inserted to identify subscribers, see Chapter 74, “Residential subscriber management” .

70.10 BTV multicast

70.10.1 Overview

This section describes how the NFM-P can be used to configure and manage the delivery of BTV multicast traffic streams. See 70.11 “BTV multicast configuration examples” (p. 1931) for examples.

Optimizing for broadcast TV means implementing multicast packet replication throughout the network. Multicasting improves the efficiency of the network by reducing the bandwidth and fiber needed to deliver broadcast channels to the end user. A multicasting device can receive a single copy of a broadcast channel and replicate it to any downstream devices that require it, thus substantially reducing the required network resources. This efficiency becomes increasingly important closer to the end user.

70.10.2 Multicast routing overview

Multicast routers direct traffic to several receivers without incurring any additional overhead for the source router or the receivers. In contrast, unicast networks can suffer from increased bandwidth requirements as the number of receivers increases.

In multicast routing, receivers query the source router to request a specific data stream. The multicast group is not restricted by physical location, so long as the hosts can be reached through the Internet. Routers in the network must use IGMP to send and receive multicast data streams.

A multicast-enabled device, such as a switch or router, distributes a data stream to multiple receivers. Multicast packets are replicated in the network by routers that are enabled with PIM, which results in the efficient delivery of data to multiple receivers using less bandwidth.

1. A switch or router distributes a data stream to multiple receivers, such as multicast-enabled PE switches or routers.

2. The multicast-enabled switch or router replicates the data stream, when required, and transmits a copy to each downstream switch or router in the multicast tree.
3. Each client receives the data stream it has subscribed to from the downstream switch or router.

The NEs involved in delivering BTV multicast streams are first preconfigured through CLI for discovery and management by NFM-P. After discovery, routing protocols are applied to the NEs using NFM-P. Routing, QoS, and network queue policies are then created. On some devices, multicast package and ACL filter policies are created. These policies are applied to NEs during service creation through NFM-P.

70.10.3 Content delivery

BTV source traffic consists of one IP multicast stream per broadcast channel. As a multicast stream enters the core network, it is directed by PIM to the RP, which replicates the multicast traffic to all DRs that have requested the specific multicast stream. DRs distribute the multicast stream directly to set-top receivers or through an M-VPLS to BTV VLAN rings to which customer set-top receivers are connected. Multicast streams are forwarded only to those set-top receivers that have requested them through IGMP and are entitled to them as subscriber hosts.

70.10.4 Content management

Multicast package policies on some devices define the available multicast addresses (BTV channels) for end users in a BTV network. Typically, a root package policy which includes all BTV channels associated with an NFM-P customer service is created. Subsets of the root policy are then created as BTV content packages to which end customers can subscribe. ACL filter policies on CE devices ensure that only the channels to which an end customer has subscribed are delivered to the customer set-top receiver.

70.10.5 PIM

PIM uses RPF to correctly forward multicast packets down a distribution tree, using the independent multicast and unicast routing tables created by the 7450 ESS in mixed mode or the 7750 SR. The unicast routing table is populated by the unicast routing protocols, such as OSPF, BGP, IS-IS, or static routes, which can also be configured to submit routes to the multicast routing table.

Depending on the configuration of the PIM routing instance, RPF can use the unicast routing table, the multicast table populated by the unicast routing protocols, or both to determine the upstream sources of multicast streams. PIM forwards a multicast packet only if it is received on an upstream interface that is associated with a source address of an upstream router. This RPF check assures that there are no loops in the distribution tree.

PIM uses a multicast domain to group receiver hosts on a router called the rendezvous point (RP). A bootstrap router (BSR) elects an RP from available candidates. The BSR manages RP information, disseminates it to all PIM routers in the multicast domain, and elects a new RP in the case of RP unavailability.

A receiver host becomes a member of a multicast domain by sending an IGMP join request for a multicast stream to a PIM designated router (DR). If the router does not currently receive the multicast stream, PIM updates the DR routing table with the receiver host IP address and requests the multicast stream from the RP. The RP adds the router to the distribution tree. Packets sent to

the multicast IP address are propagated down the distribution tree to the receiver host. DRs use the RP as the source for a multicast stream unless a source with a lower path cost is available.

PIM stops sending a multicast stream to a router when it determines that there are no active receiver hosts for the multicast stream in that branch of the distribution tree.

70.10.6 MVR on VPLS

PIM is not supported on 7450 ESSs. When receiver hosts are connected to a PIM DR by way of a 7450 ESS, MVR must be configured on the switch. MVR allows multicast traffic to be forwarded downstream from the DR to the receiver host over an MVR VPLS.

70.10.7 IGMP

IGMP is a multicast protocol which service providers can use to establish multicast group memberships on a LAN. Within the LAN, end users use IGMP to communicate with a local multicast router, which then uses PIM to distribute the IGMP messages to other local and remote multicast routers. Multicast routers send regular membership queries to IGMP hosts which respond with membership reports. Multicast routers can use these reports to determine which hosts are interested in receiving particular multicast messages.

IGMP operates above the network layer on IPv4 networks.

70.10.8 MLD and MLD-snooping

The Multicast Listener Discovery protocol is essentially the IPv6 version of IGMP. It is used by IPv6 routers to discover the presence of multicast listeners (that is, NEs that wish to receive multicast packets) on their directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring NEs.

MLD version 2, or MLDv2, can interoperate with MLD version 1, or MLDv1. MLDv2 adds the ability for an NE to report interest in listening to packets with a specific multicast address only from specific source addresses or from all sources except for specific source addresses.

Most NFM-P NEs support MLD-snooping but not MLD. The 7450 ESS, 7750 SR, and 7950 XRS routers allow the enabling of MLD snooping for VPLS services. All variants of the 7705 SAR support MLD but not MLD-snooping.

70.11 BTV multicast configuration examples

70.11.1 Overview

Figure 70-17, “BTV multicast delivery examples” (p. 1933) shows a simple BTV network and three methods of content delivery, examples A, B, and C. The sequence of specific configuration steps for each example follows general device, network and multicast configuration information common to all examples.

Note: In the examples on the following pages, references to “IGMP-snooping” may generally be read as “IGMP- or MLD-snooping”, provided the routers employed in the configurations support the MLD protocol.

70.11.2 Device preconfiguration

A network device requires CLI preconfiguration before it can be managed by the NFM-P.

The primary CLI preconfiguration actions for a device are:

- Assigning a system ID to the device
- Enabling and configuring SNMP on the device
- Enabling Telnet access on the device

See Chapter 8, “Device commissioning and management” and the specific device documentation for more information about enabling device functionality before using the NFM-P.

Figure 70-17 BTV multicast delivery examples

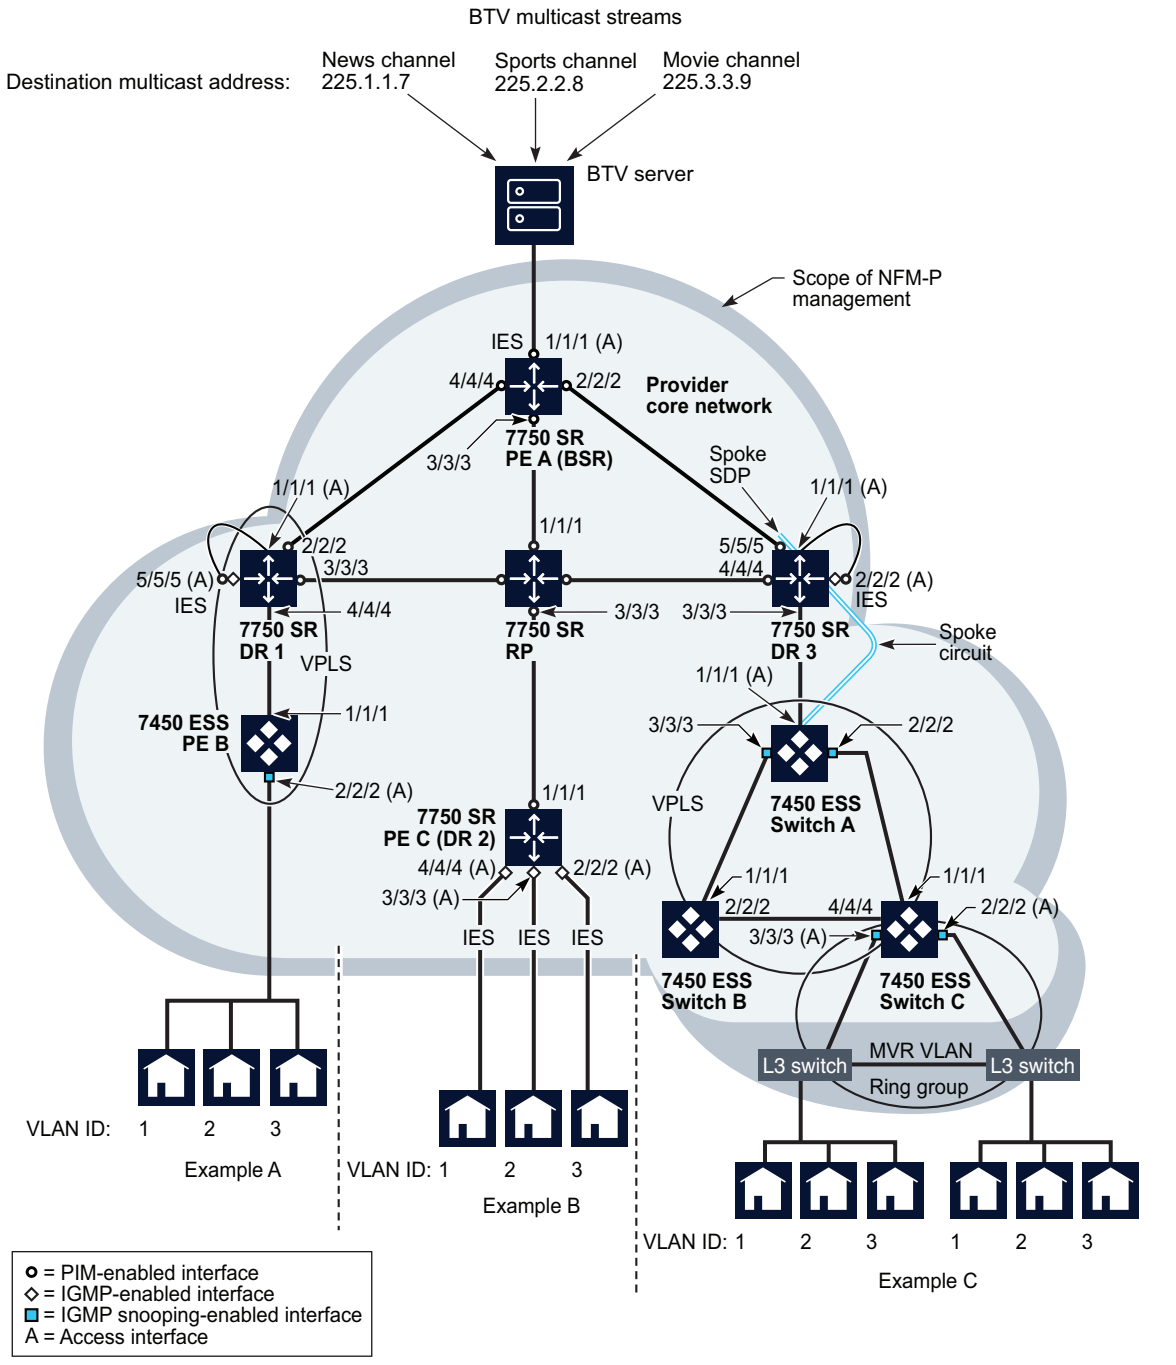

After CLI preconfiguration, further actions are required:

- Discover devices, including mediation configuration with CLI user names and passwords.
- Use NFM-P to set the discovered device in a managed state.

See Chapter 9, “Device discovery” for more information about device discovery and management using NFM-P.

70.11.3 Network preconfiguration

The core network shown in Figure 70-17, “BTV multicast delivery examples” (p. 1933) represents a fully meshed group of devices. For simplicity, only the devices relevant to the BTV multicast examples are shown.

Network preconfiguration consists of the following sequence of actions:

1

Configure network devices for in-band or out-of-band management. See Chapter 8, “Device commissioning and management” for more information.

2

Configure a system interface on each device to serve as the identifier for the device. See Chapter 27, “NE routing and forwarding” for more information.

3

Configure network interfaces on each router to establish a full mesh of interconnectivity between devices. In Figure 70-17, “BTV multicast delivery examples” (p. 1933) the interfaces to be configured are PE A, ports 1/1/1, 2/2/2, 3/3/3, and 4/4/4; DR 1, port 2/2/2; RP, ports 1/1/1, 2/2/2, and 4/4/4; and DR 3, ports 1/1/1 and 4/4/4. See Chapter 27, “NE routing and forwarding” for more information.

4

Cable the network-interface ports between routers in the core network to establish the physical connectivity shown in Figure 70-17, “BTV multicast delivery examples” (p. 1933) .

5

Use CLI ping commands to check IP connectivity between devices. See the device documentation for more information.

6

Enable IGPs such as RIP, OSPF, or IS-IS on devices according to network size and complexity. See Chapter 28, “Routing protocol configuration” for information about enabling routing protocols.

7

Enable an inter-AS routing protocol such as BGP or OSPF to PE routers, if required. See Chapter 28, “Routing protocol configuration” for more information.

8

Create routing policies as required. Create one multicast group for each BTV multicast destination address during policy creation. See Chapter 27, “NE routing and forwarding” for more information.

9

Configure routing protocols and apply routing policies as required. See Chapter 28, “Routing protocol configuration” for more information.

10

Configure LDP and MPLS, if required. See Chapter 28, “Routing protocol configuration” for information about configuring LDP. See Chapter 31, “MPLS” for information about configuring MPLS.

- Enable MPLS and LDP on the routing instance of each device that is participating in the MPLS network.
- Assign a Layer 3 interface to the MPLS instance on each MPLS-enabled device.
- Create a mesh of MPLS paths.
- Create a mesh of LSPs.
- Use the NFM-P to create MPLS administrative groups, and assign the groups to MPLS interfaces and LSP paths as required.

70.11.4 Multicast configuration common to all examples

The network connections shown between PE A and DR 1 and between PE A and DR 3 represent redundant multicast routes used by PIM in the event of an RP failure. PIM dynamically adjusts to BSR or RP failure by electing a replacement BSR or RP or by using a previously defined backup BSR or RP. PIM chooses the most appropriate source for a multicast stream based on path cost and source availability and bypasses the RP if a better source for a multicast stream is found. For simplicity, Figure 70-17, “BTV multicast delivery examples” (p. 1933) does not show routes to PE C (DR 2) from DR 1 or DR 3. As shown, PE C (DR 2) is isolated from multicast traffic in the event of RP failure.

Network multicast configuration common to all three examples involves the following sequence of actions:

1

Enable IGMP on routers DR 1, PE C (DR 2), and DR 3. See 28.103 “To enable IGMP on a routing instance” (p. 1016) for more information.

2

Configure IGMP on routers DR 1, PE C (DR 2), and DR 3. See 28.104 “To configure an IGMP site on a router” (p. 1017) for more information.

3

Enable PIM on routers PE A, RP, DR 1, PE C (DR 2), and DR 3. See 28.97 “To enable PIM on a routing instance” (p. 998) for more information.

4

Configure PIM on routers PE A, RP, DR 1, PE C (DR 2), and DR 3. See 28.98 “To configure PIM on a routing instance” (p. 998) for more information.

- Specify PE A as the candidate bootstrap router.
- Specify RP as the candidate rendezvous point. You can also specify it as a static RP for a multicast domain, if there are multiple BTV domains, and configure a second router as a redundant RP.
- Specify IES as the Apply to parameter value on routers PE A, DR 1, PE C (DR 2), and DR 3.

5

Create PIM interfaces at PE A, ports 1/1/1, 2/2/2, 3/3/3, and 4/4/4; RP, ports 1/1/1, 2/2/2, 3/3/3 and 4/4/4; DR 1, ports 2/2/2 and 3/3/3; DR 2, port 1/1/1 and DR 3, ports 4/4/4 and 5/5/5. See 28.101 “To create a PIM interface on a base routing instance or VPRN routing instance” (p. 1013) for more information.

6

Create QoS, scheduling, and accounting policies for the ingress BTV traffic. See Chapter 49, “Policies overview” for more information.

7

Create an IES from PE A, port 1/1/1, to the BTV multicast provider's network. See 78.3 “To create an IES” (p. 2430) and 78.5 “To configure an IES site” (p. 2431) for more information. Enable PIM on the IES SAP during IES creation.

70.11.5 Example A configuration

In Example A, IGMP join requests from residential hosts ingress a VPLS SDP. IGMP snooping on the VPLS registers the join requests on the local switch. The switch sends the requests over the VPLS, which is physically cross-connected to an IGMP- and PIM-enabled IES SAP on the DR. PIM on the DR requests the desired multicast stream, if not present, from the RP. The requested stream then traverses the VPLS and is sent to end users.

1

Configure PE B, port 1/1/1 and DR 1, port 4/4/4 as network ports. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

2

Configure PE B, port 2/2/2 and DR 1, port 1/1/1 as access ports. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

3

Cable DR 1, port 3/3/3, and PE B, port 1/1/1 to establish physical connectivity.

4

Configure DR 1, port 5/5/5 as an access port. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

5

Create an IES on DR 1, port 5/5/5. See 78.3 “To create an IES” (p. 2430) and 78.5 “To configure an IES site” (p. 2431) for more information.

- Enable IGMP on the IES SAP during IES creation.
- Enable PIM on the IES SAP during IES creation.

6

Connect a cable between ports 1/1/1 and 5/5/5 on DR 1 as a service cross connect.

7

Create QoS, scheduling, filter, and accounting policies to apply to egress BTV traffic during service creation. See Chapter 49, “Policies overview” for information about policy creation.

8

Create a distributed VPLS with endpoints at PE B, port 2/2/2 and DR 1, port 1/1/1. See 77.5 “To create a VPLS” (p. 2249) and 77.33 “To configure a VPLS site” (p. 2294) for more information.

- Enable IGMP snooping on the VPLS SDP at PE B, port 2/2/2.
- Apply previously defined QoS, scheduling, filter, and accounting policies to the VPLS SDP at PE B, port 2/2/2.

70.11.6 Example B configuration

In Example B, an IGMP join request ingresses an IES SAP on the DR. PIM on the DR requests the desired multicast stream, if not present, from the RP. The requested stream is then delivered over an IES to an end user.

1

Configure PE C (DR 2), ports 2/2/2, 3/3/3, and 4/4/4 as access ports. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

2

Create QoS, scheduling, filter, and accounting policies to apply to egress BTV traffic during IES creation. See Chapter 49, “Policies overview” for information about policy creation.

3

Create IES services on PE C (DR 2), ports 2/2/2, 3/3/3, and 4/4/4 that terminate on the CE set-top devices. See 78.3 “To create an IES” (p. 2430) for more information.

- Enable IGMP on each IES SAP during IES creation.
- Apply previously defined QoS, scheduling, filter, and accounting policies to each IES, as required.

70.11.7 Example C configuration

In Example C, IGMP join requests from residential hosts pass over an MVR VLAN to an VPLS. IGMP snooping on the VPLS registers the join requests on the local switch, which passes them over the VPLS to a spoke SDP on the DR. The spoke SDP' port is physically cross-connected to an IGMP- and PIM-enabled IES SAP on the DR. PIM on the DR requests the desired multicast stream, if not present, from the RP, then sends the stream over the VPLS and MVR VLAN to the end users.

1

Create a BTV MVR VLAN of L3 switches in a ring group with endpoints on Switch A, ports 2/2/2 and 3/3/3. See 75.4 “Sample BTV VLAN configuration” (p. 2069) and 75.9 “To create an OmniSwitch BTV VLAN service” (p. 2078) for more information. For redundancy, the MVR VLAN can be configured with endpoints on different switches. A VLL between the two switches acts as an unbreakable connection.

2

Enable and configure IGMP snooping on the bridge instances for the L3 switches included in the MVR VLAN.

3

Create QoS, scheduling, filter, and accounting policies to apply to egress BTV traffic during service creation. See Chapter 49, “Policies overview” for information about policy creation.

4

Configure the following as network ports:

- Switch A, ports 1/1/1, 2/2/2, and 3/3/3
- Switch B, ports 1/1/1 and 2/2/2
- Switch C, ports 1/1/1, 2/2/2, 3/3/3 and 4/4/4

  See 16.24 “To configure Ethernet ports” (p. 599) for more information.

5

Create a distributed VPLS consisting of Switch A, Switch B, and Switch C. See 77.5 “To create a VPLS” (p. 2249) and 77.33 “To configure a VPLS site” (p. 2294) for more information.

- Apply previously defined QoS, scheduling, filter, and accounting policies to the VPLS SDPs.
- Enable and configure IGMP snooping on the VPLS SDPs that are part of the MVR VLAN.
- Ensure that the encapsulation value of the VPLS SDPs that are part of the MVR VLAN matches the MVR VLAN ID.
- Create a split horizon group during VPLS creation to allow later addition of a spoke circuit to the VPLS.
- Configure STP on the VPLS, as required.

6

Configure DR 3, port 2/2/2 as an access port. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

7

Create an IES on DR 3, port 2/2/2. See 78.3 “To create an IES” (p. 2430) and 78.5 “To configure an IES site” (p. 2431) for more information.

- Enable IGMP on the IES SAP during IES creation.
- Enable PIM on the IES SAP during IES creation.

8

Connect a cable between ports 1/1/1 and 2/2/2 on DR 3 as a service cross connect.

9

Create a VPLS spoke SDP at DR 3, port 1/1/1. See 77.6 “To create an HVPLS” (p. 2249) for more information.

70.11.8 Example D configuration

Figure 70-18, “BTV multicast delivery using MVR on VPLS example” (p. 1940) shows an example of BTV multicast delivery using MVR on VPLS. See Figure 70-17, “BTV multicast delivery examples” (p. 1933) , and 70.11.3 “Network preconfiguration” (p. 1934) and 70.11.4 “Multicast configuration common to all examples” (p. 1935) for common network configuration information.

Figure 70-18 BTV multicast delivery using MVR on VPLS example

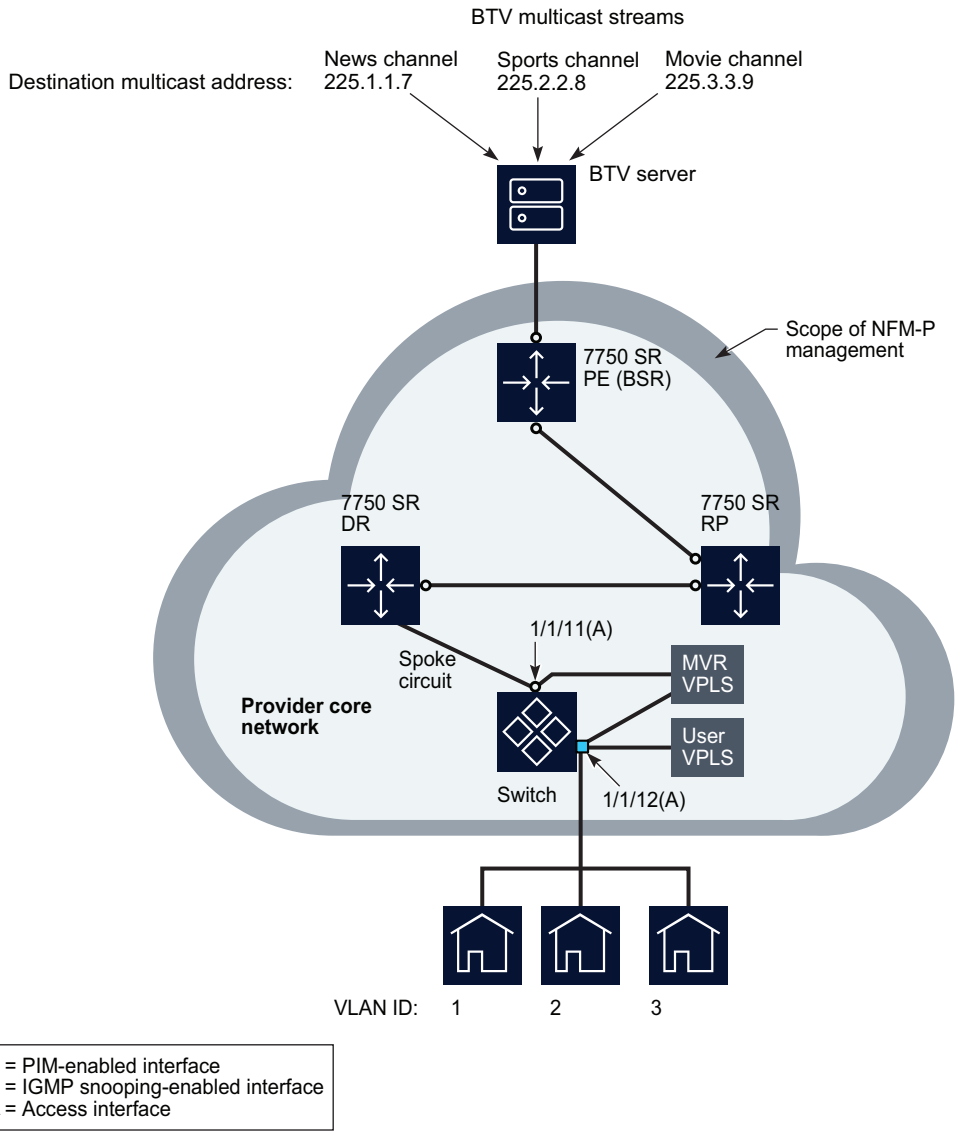

70.11.9 BTV multicast delivery using MVR on VPLS

In Example D, IGMP join requests from residential hosts are sent to a user VPLS on the 7450 ESS. IGMP snooping on the user VPLS registers the join requests on the switch, which sends them to the 7750 SR DR. PIM on the DR requests the desired multicast stream, if not present, from the RP, then sends the stream over the MVR VPLS to the user VPLS, from which the multicast stream is sent to the end users.

1

Create a multicast package policy to apply to the 7450 ESS that belongs to the MVR VPLS. See Chapter 52, “Multicast policies” for more information.

2

Configure the following ports as access ports. See 16.24 “To configure Ethernet ports” (p. 599) for more information.

- 7450 ESS, port 1/1/11
- 7450 ESS, port 1/1/12

3

Create an MVR VPLS on the 7450 ESS with SAPs 1/1/11 and 1/1/12. Apply the previously defined multicast package policy to the MVR VPLS.

4

Create a user VPLS on SAP 1/1/12 of the 7450 ESS.

- Associate the user VPLS with the previously created MVR VPLS to identify the MVR VPLS as the source of the multicast traffic.
- Enable and configure IGMP snooping on the site.

5

Create a spoke circuit between the 7450 ESS (endpoint 1/1/11) and the 7750 SR DR.

70.12 Sample network QoS configuration

70.12.1 Overview

Figure 70-19 Example of a service configuration using QoS

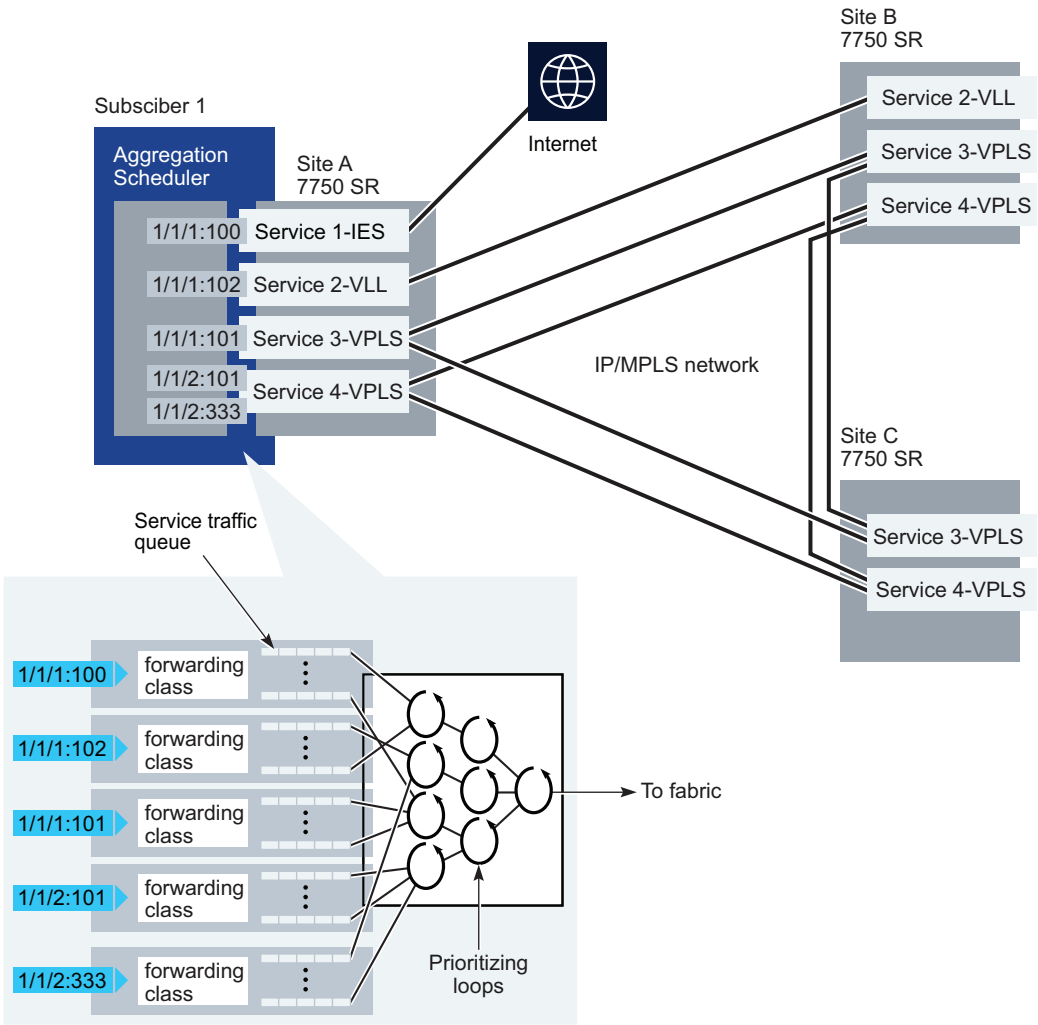

70.12.2 Workflow to configure network QoS

Note: In this configuration, the following services are provisioned:

- Service 1: IES for Internet access, which requires a CIR of 10 Mb/s and a PIR of 100 Mb/s
- Service 2: VLL service for FTP connectivity between Site A and Site B, which requires a CIR of 10 Mb/s and a PIR of 20 Mb/s

- Service 3: VPLS for video-conference service over sites A, B, and C, which requires a CIR of 20 Mb/s and a PIR of 50 Mb/s
- Service 4: VPLS for voice traffic, which requires a CIR of 10 Mb/s and a PIR of 20 Mb/s

The cumulative rate at site A needs to be limited to 70 Mb/s.

The following high-level steps are required to create the Figure 70-19, “Example of a service configuration using QoS” (p. 1942) configuration with rate limiting using QoS at Site A. Similar steps are required to configure QoS for Subscriber 1 on Sites B and C:

1

Configure a scheduler policy.

2

Create Subscriber 1.

3

Create the aggregation scheduler for Subscriber 1 on site A and assign ingress and egress scheduler policies to the aggregation scheduler.

4

Create IES, VLL, and VPLS for Subscriber 1.

- Specify sites for the services.
- Specify access interfaces for the sites.
- Specify the aggregation scheduler policy for the access interfaces.
- Bind the services to tunnels for transport through the IP/MPLS network.

70.13 Sample SAP QoS configuration

70.13.1 Overview

You can use NFM-P to configure and enforce traffic rate limiting, based on the priority of the traffic entering the ingress SAP of a service. This configuration limits bandwidth, to ensure that SLAs are met and higher priority traffic is processed first.

The figure below shows traffic of different priorities from a customer site to an ingress SAP.

Figure 70-20 Example showing ingress traffic to a service SAP

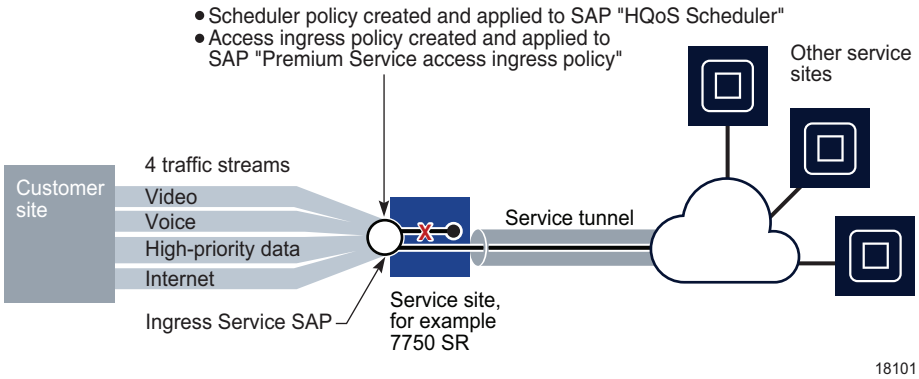

Based on the example, use the NFM-P client GUI to perform the following actions:

- Configure a parent scheduler that handles scheduling for a child scheduler.

  You do not have to configure a parent tier 1 scheduler for a child tier 2 scheduler. You can create a tier 2 scheduler on its own, with no parent. This example is meant to show hierarchical QoS.
  - Add a root tier 1 scheduler.
  - Add one child tier 2 scheduler with the root scheduler as a parent.

Create an access ingress policy named Premium service access ingress policy.

- Create 4 queues within the access ingress policy, one for each type of traffic from the customer site.
- Associate each queue with a forwarding class
- Classify the incoming customer site traffic to a forwarding class. For this sample, IP precedence is used to classify traffic. You could classify traffic other ways, for example, based on filters for IP address or DSCP marking.
- Associate each queue with the appropriate scheduler.

Apply the schedulers and access ingress policy to the appropriate L2 or L3 interface for the customer service, for example, a VPLS L2 interface.

Figure 70-21 Traffic handling based on ingress access policy and scheduler queue

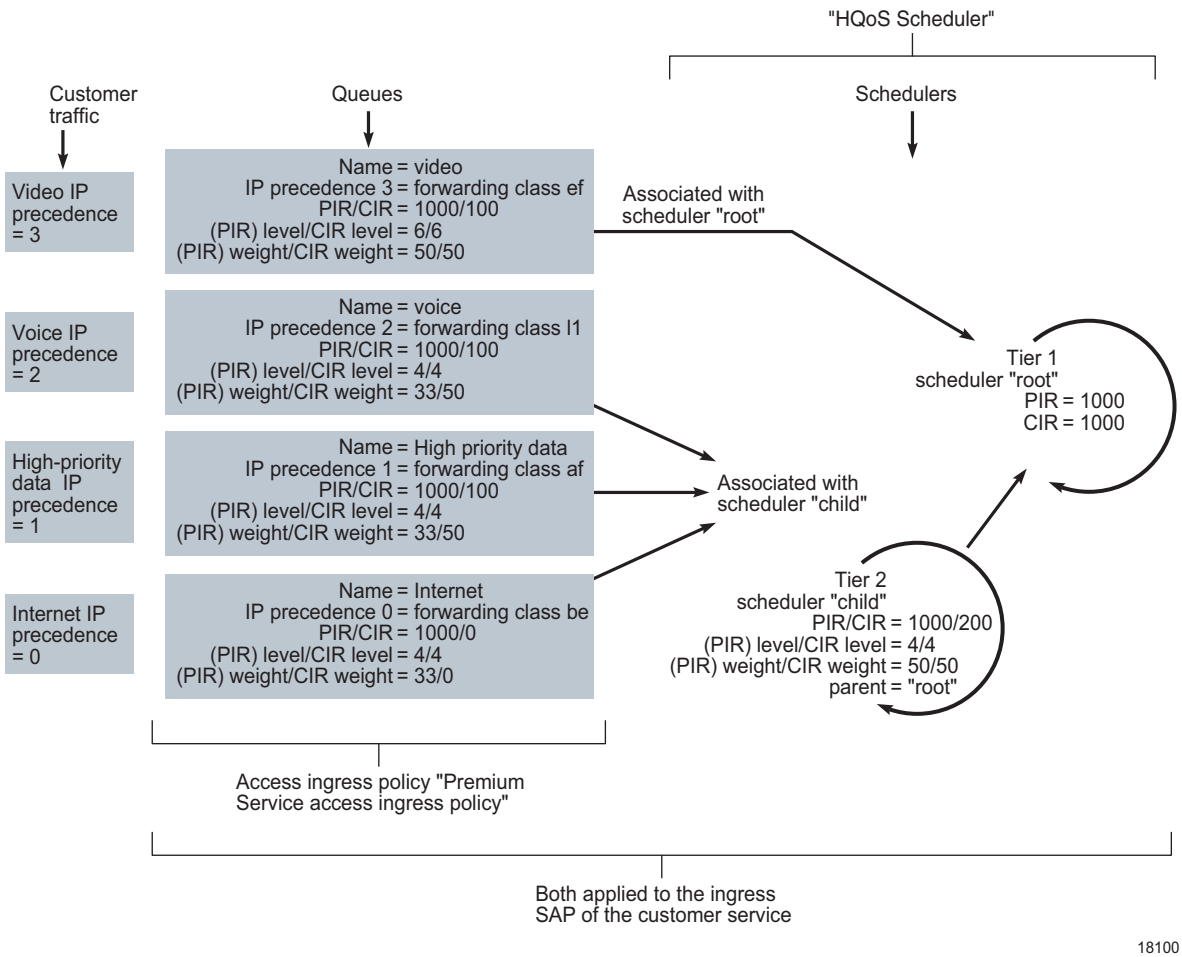

70.13.2 Sample configuration steps for SAP QoS

Note: In this sample:

- higher-priority video traffic, with IP precedence bit 3 set, goes into queue 4 with a PIR of 1000 Kb/s and a CIR of 100 Kb/s. This traffic is handled by the tier 1 scheduler. The levels and weights associated with the video queue (forwarding class of ef, level of 6, PIR/CIR weights of 50/50) ensure this traffic gets all required bandwidth.
- all other traffic goes into its appropriate queues, based on the mapping of the IP precedence bit with the forwarding class.
- The voice, high-priority data, and Internet queues are serviced by the tier 2 scheduler.
- Because the voice and high-priority data queues have higher PIR/CIR weight than the Internet queue, when there is contention for bandwidth, the voice and high-priority data queues are processed first.

Preconfiguration

1

Before you begin, you must have:

- the service ingress SAP configured as access
- all necessary cabling and network routing protocol configurations complete to handle routing packets to and from the CPE equipment

Configuration

2

Choose Policies→QoS→SROS QoS→Scheduler→Scheduler from the NFM-P main menu. The manage scheduler policies form opens.

3

Click Create. The scheduler policy create form opens.

4

Configure the parameters. Set the displayed name to High QoS.

5

Click on the Schedulers tab.

6

Click Create. The scheduler entry form opens.

7

Create a root, tier 1 scheduler entry.

1. Configure the parameters.
   - Displayed name to root
   - tier to 1
   - Summed CIR to false
   - PIR (kbps) to 1000
   - CIR (kbps) to 1000
2. Click OK. The root scheduler appears in the list of scheduler entries.

8

Create a child, tier 2 scheduler entry.

1. Click Create.
2. Configure the parameters.
   - Displayed name to child
   - tier to 2
   - Summed CIR to false
   - PIR (kbps) to 1000
   - CIR (kbps) to 200
   - Parent scheduler to root, using the Select button to choose root from the list.
   - Level (PIR level) to 4
   - Weight (PIR weight) to 50
   - CIR Level to 4
   - CIR Weight to 50
3. Click OK. The child scheduler appears in the list of scheduler entries.

9

Click Apply.

10

Close the Scheduler form.

11

Choose Policies→QoS→SROS QoS→Access Ingress→SAP Access Ingress from the NFM-P main menu. The SAP Access Ingress Policies form opens.

12

Click Create. The SAP Access Ingress Policy, Global Policy [Create] form opens.

13

Configure the parameters. Set the displayed name to Premium service access ingress policy.

14

Click on the Queues tab. Create four queues. For this sample, default queue 1 is modified, and three new queues are added. The default multicast queue 11 is unchanged.

15

Select queue 1 from the list.

1. Click on the Properties button. The queue 1 edit form opens.
2. Configure the parameters.
   - Displayed Name to Internet
   - Scheduler to child by using the Select button and choosing child from the list
   - Level (PIR level) to 4
   - CIR Level to 4
   - Weight (PIR weight) to 33
   - CIR Weight to 0
3. Click on the CIR/PIR tab.
4. Configure the parameters.
   - Rate Type to kbps
   - Policed to false
   - deselect both MAX buttons
   - Cir (kbps) to 0
   - Pir (kbps) to 1000
5. Click OK.

16

Add queue 2.

1. Click Create. The queue create form opens.
2. Configure the parameters.
   - ID to 2
   - Displayed Name to High-priority traffic
   - Scheduler to child by using the Select button and choosing child from the list
   - Level (PIR level) to 4
   - CIR Level to 4
   - Weight (PIR weight) to 33
   - CIR Weight to 50
3. Click on the CIR/PIR tab.
4. Configure the parameters.
   - Rate Type to kbps
   - Policed to false
   - deselect both MAX buttons
   - Cir (kbps) to 100
   - Pir (kbps) to 1000
5. Click OK. The queue is added to the list.

17

Add queue 3.

1. Click Create. The queue create form opens.
2. Configure the parameters.
   - ID to 3
   - Displayed Name to Voice
   - Scheduler to child by using the Select button and choosing child from the list
   - Level (PIR level) to 4
   - CIR Level to 4
   - Weight (PIR weight) to 33
   - CIR Weight to 50
3. Click on the CIR/PIR tab.
4. Configure the parameters.
   - Rate Type to kbps
   - Policed to false
   - deselect both MAX buttons
   - Cir (kbps) to 100
   - Pir (kbps) to 1000
5. Click OK. The queue is added to the list.

18

Add queue 4.

1. Click Create. The queue create form opens.
2. Configure the parameters.
   - ID to 4
   - Displayed Name to Video
   - Scheduler to root by using the Select button and choosing root from the list
   - Level (PIR level) to 6
   - CIR Level to 6
   - Weight (PIR weight) to 50
   - CIR Weight to 50
3. Click on the CIR/PIR tab.
4. Configure the parameters.
   - Rate Type to kbps
   - Policed to false
   - deselect both MAX buttons
   - Cir (kbps) to 100
   - Pir (kbps) to 1000
5. Click OK. The queue is added to the list.

19

Click Apply to save the changes. Confirm the action.

20

Associate each queue with a forwarding class.

1. Click on the Forwarding Classes tab.
2. Click Create. The forwarding class create form opens.
3. Configure the parameters.
   - Forwarding class to be
   - Queue ID to 1. This is the best-effort Internet traffic queue.
4. Click OK. Confirm the action.
5. Click Create. The forwarding class create form opens.
6. Configure the parameters.
   - Forwarding class to af
   - Queue ID to 2. This is the high-priority data traffic queue.
7. Click OK. Confirm the action.
8. Click Create. The forwarding class create form opens.
9. Configure the parameters.
   - Forwarding class to l1
   - Queue ID to 3. This is the voice traffic queue.
10. Click OK. Confirm the action.
11. Click Create. The forwarding class create form opens.
12. Configure the parameters.
    - Forwarding class to ef
    - Queue ID to 4. This is the highest priority, video traffic queue.
13. Click OK. Confirm the action.

21

Click Apply. Confirm the action.

22

Associate the IP precedence bits of the incoming customer traffic with the forwarding class. The forwarding class is already associated with a queue. For this sample, IP precedence bits are used to associate different types of traffic with the forwarding class. You could classify traffic other ways, for example, based on filters for IP address or DSCP marking.

1. Click on the Precedence tab.
2. Click Create. The precedence create form opens.
3. Configure the parameters to associate IP precedence 0 (Internet traffic from the customer site) with forwarding class be (the be forwarding class is associated with queue 1)
   - Precedence is 0

- Forwarding Class is be

4. Click OK. Confirm the action. The association between precedence 0 and the be forwarding class is added to the list.
5. Click Create. The precedence create form opens.
6. Configure the parameters to associate IP precedence 1 (high-priority data traffic from the customer site) with forwarding class af (the af forwarding class is associated with queue 2)
   - Precedence is 1
   - Forwarding Class is af
7. Click OK. Confirm the action. The association between precedence 1 and the af forwarding class is added to the list.
8. Click Create. The precedence create form opens.
9. Configure the parameters to associate IP precedence 2 (voice traffic from the customer site) with forwarding class l1 (the l1 forwarding class is associated with queue 2):
   - Precedence is 2
   - Forwarding Class is l1
10. Click OK. Confirm the action. The association between precedence 2 and the l1 forwarding class is added to the list.
11. Click Create. The precedence create form opens.
12. Configure the parameters to associate IP precedence 3 (video traffic from the customer site) with forwarding class ef (the ef forwarding class is associated with queue 3)
    - Precedence is 3
    - Forwarding Class is ef
13. Click OK. Confirm the action. The association between precedence 3 and the ef forwarding class is added to the list.

23

Click Apply. Confirm the action.

24

Associate a service SAP with the created High QoS scheduler and the created Premium service access ingress policy. There are many ways to associate policies with L2 or L3 interfaces used as service SAPs, for example, from the service creation form or the port properties form. This sample modifies an existing L2 interface for an existing VPLS.

1. Choose Manage→Service→Services from the NFM-P main menu. The manage services form opens.
2. Set the filters and click on the Search button. A list of filtered services appears.
3. Select the service and click on the Properties button. The service form opens.
4. Click on the L2 Access Interfaces tab.
5. Choose an interface and click on the Properties button. The L2 interface edit form opens.
6. Click on the QoS tab.

7. Configure the parameter. Use the Select button to set the Ingress Policy to Premium service access ingress policy. The policy ID and displayed name appear.
8. Click on the Schedulers tab.
9. Configure the parameter. Use the Select button to set the Ingress Scheduler to High QoS. The displayed name appears.
10. Click Apply to save the changes. Confirm the action.

70.14 Sample QoS configuration on the 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS

70.14.1 Stages

Planning and configuration

1

Perform network planning activities:

1. Determine the required types of services or applications; for example, voice, video, and data.
2. Review SLAs.
3. Perform traffic engineering activities at the IP/MPLS core level to ensure that resources are available.

2

Configure the IP and MAC ACL filters, as required. See 51.5 “To configure an ACL IP filter policy” (p. 1671) and 51.4 “To configure an ACL MAC filter policy” (p. 1668) for more information.

3

Configure the slope policies. See 50.49 “To configure a WRED slope policy” (p. 1587) and 50.51 “To configure an HSMDA WRED slope policy” (p. 1591) for more information.

4

Configure the scheduler policies. Scheduler policies can be shared between ingress and egress policies, depending on your specific requirements. See 50.55 “To configure a scheduler policy” (p. 1596) for more information.

5

Configure the port scheduler policies. See 50.57 “To configure a port scheduler policy” (p. 1599) for more information.

6

As required, create the aggregation schedulers. See 50.56 “To create an Aggregation Scheduler” (p. 1597) for more information.

7

Configure the SAP access ingress policies. See 50.28 “To configure a SAP access ingress policy” (p. 1538) for more information.

1. Configure the forwarding classes.
2. Configure the queues.
3. Map QoS markings on ingress packets to the forwarding classes.
4. Map forwarding classes to the queue definitions.
5. Map queue definitions to the scheduler policies.

8

Configure the SAP access egress policies. See 50.30 “To configure a SAP access egress policy” (p. 1550) for more information.

1. Configure the queues.
2. Map forwarding classes to the queue definitions.
3. Map queue definitions to the schedulers policies.

9

Configure the network policies. See 50.41 “To configure a QoS network policy” (p. 1567) for more information.

1. Configure for ingress:
   - Map QoS markings on ingress packets to the forwarding classes.
   - Map forwarding classes to the queue definitions.
2. Configure for egress:
   - As required, configure remarking.
   - As required, map QoS markings to the forwarding classes.

10

Configure the network queue policies. See 50.46 “To configure a network queue policy” (p. 1580) for more information.

11

Configure the time of day policies.

1. Configure the time ranges. See 53.3 “To configure a time range policy” (p. 1732) for more information.

2. Create time of day suites. See 53.4 “To configure a time-of-day suite policy” (p. 1733) for more information.

Application of policies and schedulers to equipment and interfaces

12

Associate the slope policies with ports or daughter cards. See 16.24 “To configure Ethernet ports” (p. 599) and 15.78 “To configure an MDA” (p. 536) for more information.

13

Associate the network queue policies with MDAs. See 16.24 “To configure Ethernet ports” (p. 599) and 16.72 “To configure an L3 interface on a DS3/E3 channel on a channelized ASAP MDA” (p. 674) for more information.

Configuration of customers and services

14

Configure the customers. See the appropriate service chapter for more information.

15

Create the customer services, and assign policies during the configuration. See the appropriate service chapter for more information.

16

Assign the aggregation schedulers at the interface level, if required. See the appropriate service chapter for more information.

70.15 Sample QoS configuration on the 7210 SAS

70.15.1 Overview

Quality of service configuration for 7210 SAS devices is based on the NFM-P policy framework, and is similar to support for other Nokia devices. However, some configurations are specific to 7210 SAS devices. The NFM-P provides separate QoS policies for 7210 SAS NEs. See 50.23 “7210 SAS QoS policies” (p. 1528) for more information.

Some 7210 SAS QoS settings are not configured within the policy framework. Instead, they are configured as device or port properties, or during SAP configuration. See 70.15.2 “Stages” (p. 1955) for more information.

For some 7210 SAS QoS policies and functions, device system resources must be appropriately allocated. See 6.5.13 “System resource profile” (p. 220) and the NE documentation for more information about system resource allocation.

On supporting 7210 SAS NEs, you can configure QoS for self-generated traffic on routing instances and VPRN sites; see27.1.13 “Self-generated traffic and QoS” (p. 822).

See the 7210 SAS OS Quality of Service Guides for more information about 7210 SAS QoS and QoS policies.

70.15.2 Stages

The following workflow provides an overview of the tasks required to configure QoS on 7210 SAS NEs. The required configurations vary, depending on the chassis type and network requirements. See the 7210 SAS OS Quality of Service Guides for more information about 7210 SAS QoS.

1

Based on SLAs and network capacity, determine the bandwidth and line rate requirements for the NEs and services for which you are configuring QoS.

2

Configure the required NE properties:

a. Two WRED slopes; see 12.48 “To configure two WRED slopes on a 7210 SAS” (p. 379)

b. Frame-based accounting; see 12.49 “To configure frame-based accounting for QoS policies on a 7210 SAS” (p. 380)

c. System resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) or 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382)

d. Port-based scheduling; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384)

e. Buffer management (port decommissioning); see 12.54 “To configure buffer management for the 7210 SAS” (p. 385)

3

Configure the required port properties:

a. Port egress rate limiting; see 16.24 “To configure Ethernet ports” (p. 599)

b. DEI classification; see 16.24 “To configure Ethernet ports” (p. 599)

c. Egress scheduler mode (for the 7210 SAS-X); see 16.24 “To configure Ethernet ports” (p. 599)

d. SAP-based remarking (for the 7210 SAS-X); see 16.24 “To configure Ethernet ports” (p. 599)
.

4

Configure and distribute the required QoS policies:

a. 7210, 7250 and 1830 SAP access ingress; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)

b. 7210 and 1830 port access egress; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)

c. 7210 SAP access egress; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)

d. 7210 and 1830 network; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571)

e. 7210 and 1830 network queue; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)

f. 7210 and 1830 slope; see 50.50 “To configure a 7210 and 1830 slope policy” (p. 1589)

g. 7210 and 7250 queue management; see 50.52 “To configure a 7210 and 7250 Queue Management policy” (p. 1592)

h. 7210, 7250 and 1830 port scheduler; see 50.60 “To configure a 7210, 7250 and 1830 Port Scheduler policy” (p. 1601)

i. 7210/7250 Dot1p classification; see 50.84 “To configure a 7210/7250 Dot1p classification policy” (p. 1637)

j. 7210/7250 DSCP classification; see 50.85 “To configure a 7210/7250 DSCP classification policy” (p. 1638)

k. 7210/7250 MPLS LSP-EXP classification; see 50.86 “To configure a 7210/7250 MPLS LSP-EXP classification policy” (p. 1639)

l. 7210 remarking; see 50.80 “To configure a 7210 remarking policy” (p. 1630)

m. 7210 MPLS LSP-Exp map; see 50.83 “To configure a 7210 MPLS LSP-EXP Mapping policy” (p. 1635)

n. 7210 multipoint bandwidth management; see 52.16 “To configure a 7210 multipoint bandwidth management policy” (p. 1725)

5

Assign QoS policies to network resources.

a.

Assign 7210/7250 Dot1p and DSCP classification policies to:

- SAP access ingress policies; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
- network policies; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571)
- Ethernet ports (7210/7250 DSCP classification policies only), for table-based color-aware ingress classification on supporting NEs; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636).
- L3 interfaces configured for RVPLS (7210/7250 DSCP classification policies only), for table-based color-aware ingress classification on supporting NEs; see 78.30 “To bind an IES L3 access interface to a VPLS site or VPLS I-site” (p. 2475) and 79.98 “To bind a VPRN L3 access interface to a VPLS site or VPLS I-site” (p. 2673)

b. Assign 7210/7250 MPLS LSP-EXP classification policies to network policies; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571).

c.

Assign slope policies to queues configured in:

- SAP access ingress policies; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
- SAP access egress policies; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)
- network queue policies; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)
- network policies; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571)

d.

Assign queue management policies to queues configured in:

- SAP access ingress policies; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544)
- port access egress policies; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)
- SAP access egress policies; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)
- network queue policies; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583)

e.

Assign remarking policies to:

- port access egress policies; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556)
- SAP access egress policies; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)
- network policies; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571)

f. Assign MPLS LSP-Exp Map policies to network policies of network interface type; see 50.42 “To configure a 7210 and 1830 network policy” (p. 1571) .

g.

Assign SAP access ingress and SAP access egress policies to SAPs; see the relevant procedure, depending on the service type:

- 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
- 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
- 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES

- 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN
- 93.10 “To create an L2 access interface on a destination site” (p. 3174) for mirror services (SAP access egress policies only)

h. Assign network policies to network interfaces; see 27.17 “To create an L3 network interface on a routing instance” (p. 856) and 27.18 “To configure L3 network interfaces” (p. 863) .

i. Assign port access egress, network, network queue, and port scheduler policies to ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) .

j. Assign slope policies to buffer pools on ports; see 16.24 “To configure Ethernet ports” (p. 599) .

k. Assign multipoint bandwidth management policies to 7210 SAS-X NEs; see 52.16 “To configure a 7210 multipoint bandwidth management policy” (p. 1725) .

6

Configure table-based color-aware ingress classification on supporting NEs; see 50.23.3 “Workflow to configure table-based DSCP ingress classification” (p. 1529).

7

Configure hierarchical ingress policing (H-metering) for access ingress meters:

1. Configure a 7210, 7250 and 1830 SAP access ingress policy with all meters set to trTCM (RFC 4115); see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .
2. Assign the policy to L2 or L3 access interfaces and configure the aggregate rate limit parameters for ingress meters; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
   - 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
   - 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
   - 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN

8

Configure hierarchical ingress queues (ingress scheduling) for 7210 SAS-K and 7210 SAS-X NEs:

1. Enable the Queues function in the system resource profile for the 7210 SAS-X NE; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

2. Configure ingress queues in a 7210, 7250 and 1830 SAP access ingress policy; see 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .
3. Assign the policy to access interfaces and configure the aggregate shaper rate parameters; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
   - 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
   - 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
   - 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN

9

Configure egress policing (egress meters) for SAPs on the 7210 SAS-D, 7210 SAS-M, 7210 SAS-R, 7210 SAS-S, 7210 SAS-Sx, and 7210 SAS-T:

1. Enable the SAP Egress Aggregate Meter parameter in the global system resource profile for the NE; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

   For 7210 SAS-R NEs, enable the Egress SAP Aggregate Meter parameter in the system resource profile policy assigned to the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).
2. For 7210 SAS-Mxp and 7210 SAS-R NEs, enable port-based scheduling; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384).
3. Configure egress meters for access interfaces; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
   - 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
   - 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
   - 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN

10

Configure egress scheduling for 7210 SAS-K NEs.

1. If required, configure port-level egress rate limiting; see 16.24 “To configure Ethernet ports” (p. 599) .
2. For SAP-based access egress scheduling, configure queues in a 7210 SAP access egress policy; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558) . Assign the policy to an L2 access interface; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL

- 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS

3. For network egress scheduling, configure queues in a 7210 and 1830 network queue policy; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583) . Assign the policy to ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) .
4. Configure the aggregate rate limit values for SAPs; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
   - 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS

11

Configure egress scheduling for 7210 SAS-Mxp, 7210 SAS-R, and 7210 SAS-X NEs.

1. If required, configure port-level egress rate limiting; see 16.24 “To configure Ethernet ports” (p. 599) .
2. For SAP-based access egress scheduling, configure queues in a 7210 SAP access egress policy; see 50.32 “To configure a 7210 SAP access egress policy” (p. 1558) . Assign the policy to L2 or L3 access interfaces; see the relevant procedure, depending on the service type:
   - 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
   - 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS
   - 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479) for IES
   - 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665) for VPRN
3. For port-based access egress scheduling on the 7210 SAS-Mxp and 7210 SAS-R, configure the QoS function in the system resource profile for the device; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384) . Configure queues in a 7210 port access egress policy; see 50.31 “To configure a 7210 and 1830 port access egress policy” (p. 1556) .
4. For network egress scheduling, configure queues in a 7210 and 1830 network queue policy; see 50.47 “To configure a 7210 and 1830 network queue policy” (p. 1583) . Assign the policy to ports; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636) .
5. For access ports on the 7210 SAS-X, configure the Egress Scheduler Mode parameter for FC-based or SAP-based scheduling; see 16.24 “To configure Ethernet ports” (p. 599) .
6. Configure the aggregate rate limit values for SAPs; see the relevant procedure, depending on the service type:
   - 76.44 “To configure scheduling on a VLL L2 access interface” (p. 2186) for VLL
   - 77.71 “To configure scheduling on an L2 access interface” (p. 2345) for VPLS

- 78.34 “To configure scheduling on an IES L3 access interface” (p. 2481) for IES
- 79.91 “To configure scheduling on a VPRN L3 access interface” (p. 2666) for VPRN

12

Configure overrides for access ingress meters, as required; see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) .

13

Configure overrides for access ingress queues, as required; see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

14

Configure overrides for port access egress queues, as required; see50.100 “To configure QoS policy overrides on port access egress queues for a 7210 SAS” (p. 1660).

15

For the 7210 SAS-K, preserve ingress dot1p values for egress traffic on L2 interfaces; see the relevant procedure, depending on the service type:

- 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for VLL
- 77.70 “To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site” (p. 2343) for VPLS

70.16 Sample QoS configuration on an OmniSwitch

70.16.1 Stages

1

Enable QoS on the OmniSwitch and configure global settings such as global port parameters, default disposition for flows, and timeouts. The parameters that you need to configure globally depend on the types of policies that you need to configure.

Typically, you do not need to change any of the global defaults. See 28.128 “To configure bridging on an OmniSwitch” (p. 1043) for information.

2

Configure QoS port parameters, which includes setting QoS parameters on a per port basis. Typically you do not need to change the port defaults. See Step 12 of 16.56 “To configure OmniSwitch Ethernet ports” (p. 647) for information.

3

Configure QoS policies. See Chapter 50, “QoS policies” for more information.

71 Queue groups

71.1 Overview

71.1.1 General information

Queue groups are objects created on access or network Ethernet ports that allow SAP or IP interface forwarding classes to be redirected from standard queue mapping to a shared queue. This provides some flexibility to maintain QoS and helps save resources on the node.

Access ingress ports support a single queue group for each ingress port. Access egress and network egress ports support the creation of multiple queue groups.

71.1.2 Queue Group Template policies

Queue Group Template policies allow you to define the queuing and parenting structure for queue groups on Ethernet ports. The policy defines the number and types of queues within the port queue group, and provides the default queue parameters.

Before you create a queue group on an Ethernet port, you must first create a Queue Group Template.

See 50.19 “Queue Group policies” (p. 1526) in “QoS policy types” (p. 1507) for more information.

71.1.3 Port queue groups

The port queue group contains the queue groups that are created based on the queue IDs defined within the associated Ingress/Egress Queue Group Template policies. Port queue groups are supported on Ethernet ports and can be created on ports within a LAG. Port queue groups are not supported on HSMDA Ethernet ports and VSM MDAs.

Network egress queue groups are not supported on the following IOM-1 cards:

- IOM-10 G
- IOM-20 G
- IOM-20 G-B

You can create a port queue group on an Ethernet port after creating an Ingress/Egress Queue Group Template policy.

Note: You must use the same name for the port queue group and Ingress/Egress Queue Group Template policy.

71.1.4 Port queue group LAGs

When a port queue group is created on a LAG, the group is individually instantiated on each link in the LAG. The queue parameters for a queue within the queue group are used for each port queue.

You can create, modify, or delete a port access ingress, access egress, or network egress queue group on the primary port of the LAG. (The primary port is the port with the lowest port ID.) The NE automatically replicates the create, modify, or delete action for the queue group on all other ports within the LAG.

Note: The NFM-P does not allow you to create, modify, or delete an Access Ingress, Access Egress, or Network Egress queue groups on non-primary ports.

When you add a port to a LAG, the port must use the same access ingress, access egress, or network egress queue groups as the existing ports on the LAG. To ensure this requirement for the port, the NFM-P implements the following sequential comparison:

- number of queue groups
- queue group names
- queue group Instance IDs
- number of queue overrides
- individual parameters

Note: Nokia recommends that you add all required ports to the LAGs before the configuration of the port queue group.

In the services context, the following requirements must be observed:

- For any service SAP that associates its connection to a LAG, the forwarding plane must exist on the card for the LAG's primary port to be selectable.
- Whenever you add a port from a different card to a LAG, the forwarding plane redirect name and associated Instance ID parameters you configure must be identical to those that have already been configured for the service.

71.1.5 Access SAP forwarding class-based redirection

Typically, each SAP has dedicated ingress and egress queues that are only used by that specific SAP. Individual SAP queuing requires a more complex provisioning model to configure the ingress and egress SLAs of the SAP. The configuration requires service awareness at the aggregation locations in the network. There are cases where individual SAP queuing is not preferred. In these cases, you can use a shared queue or an individual port queue model. You can configure a shared queue by creating access ingress and access egress queue groups, and mapping the forwarding classes of the SAP to the queues within the queue group.

You can configure forwarding class redirection on a SAP to a queue group queue ID using the access ingress/egress QoS policy. In each policy, the forwarding class to queue ID mapping can specify a queue group name.

See 71.2 "Workflow to configure access SAP forwarding class-based redirection" (p. 1967) for more information.

71.1.6 Network IP interface forwarding class-based redirection

You can create queue groups on egress network ports to provide network IP interface queue redirection. A single set of egress port-based forwarding class queues are available by default, and

all IP interfaces on the port share the queues. The creation of a network queue group allows one or more IP interfaces to selectively redirect forwarding classes to the group to override the default behavior.

The redirection of the egress forwarding class on an IP interface to an egress queue group queue ID is provisioned using the Network policy. The actual queue group name can be specified when the Network Policy is applied to the IP interfaces.

You can configure dedicated queues for each IP interface using network egress queue groups.

See 71.3 “Workflow to configure network IP interface forwarding class-based redirection” (p. 1968) for more information.

71.1.7 Queue group statistics

The packets sent to the queue of a SAP are statistically tracked by a set of counters associated with the queue group queue, not the SAP counters. The tracking occurs when a forwarding class is redirected to an ingress or egress port queue group queue.

On a network interface, the counter sets are created for each egress IP interface, not for each egress queue. The same counter set is used when a forwarding class for an egress IP interface is redirected from the default egress port queue to a queue group queue.

See “Statistics collection in the NFM-P” in the *NSP NFM-P Statistics Management Guide* for information about managing statistics collection and to view a list of the MIB counters that are available for collection using the NFM-P.

71.1.8 Queue group configuration validation rules

Table 71-1 Validation requirements for queue group configurations

Component		Validation action
Ingress Queue Group Template policy	Queue deletion	The deletion is blocked if there is an Access Ingress QoS policy with a forwarding class that is associated with the queue ID. The deletion can be blocked by the NFM-P or the NE. You can use the name binding list to verify dependencies.
	Policy deletion	The deletion is blocked if there is an Access Ingress QoS policy with a forwarding class that is associated with the policy.
Egress Queue Group Template policy	Queue deletion	The deletion is blocked if there is an access egress QoS policy with a forwarding class that is associated with the queue ID. The deletion can be blocked by the NFM-P or the NE. You can use the name binding list to verify dependencies.
		The deletion is blocked if there is a network policy with a forwarding class that is associated with the queue ID.
	Policy deletion	The deletion is blocked if there is an access egress QoS policy with a forwarding class that is associated with the policy.
		The deletion is blocked if there is a network IP interface associated with the policy.

Table 71-1 Validation requirements for queue group configurations (continued)

Component		Validation action
Port queue groups	Port Access Ingress Queue Group deletion	The deletion is blocked if there is an access ingress QoS policy, applied to the SAP, with a forwarding class that is associated with the port queue group. The SAP is directly or indirectly bound to the port by a LAG associated with the port queue group.
	Port Access Egress Queue Group	The deletion is blocked if there is an access egress QoS policy, applied to the SAP, with a forwarding class that is associated with the port queue group. The SAP is directly or indirectly bound to the port by a LAG associated with the port queue group.
	Network Egress Queue Group	The deletion is blocked if there is a network IP interface that is directly or indirectly bound to the port by a LAG associated with the port queue group.
Access SAP forwarding class-based redirection	Access ingress QoS policy	The Queue Group name must exist as an Ingress Queue Group Templates policy.
		The Queue ID must exist within the associated Ingress Queue Group Templates policy with appropriate queue type.
		Only one unique Queue Group may be referenced within one Access Ingress QoS policy.
		The current access ingress QoS policy should not be applied to the SAPs on a non-Ethernet port or an Ethernet port where the specified Access Ingress Queue Group does not exist.
		The current access ingress QoS policy should not be applied to a SLA Profile policy.
	Access egress QoS policy	The Queue Group name must exist as an Egress Queue Group Templates policy.
		The Queue ID must exist within the associated Egress Queue Group Templates policy.
		The current access egress QoS policy should not be applied to the SAPs on a non-Ethernet port or an Ethernet port where the specified Access Egress Queue Group does not exist.
		The current access egress QoS policy should not be applied to a SLA Profile policy.
	SAP Access Ingress/Egress QoS policy assignment	When a SAP Access Ingress/Egress QoS policy with a forwarding class redirection to a Queue Group Queue ID is applied to a SAP, the following configurations are verified: • The Queue Group specified in any forwarding class redirection must exist as an Access Ingress/Egress Queue Group on the port associated with the SAP • The Access Ingress/Egress QoS policy with Queue Group specified cannot be applied to SLA Profile policy

Table 71-1 Validation requirements for queue group configurations (continued)

Component		Validation action
Network IP Interface forwarding class-based redirection	Network policy	The specified queue ID must exist within the Egress Queue Group Templates policy for all IP interfaces where the Network policy is applied. If the Network policy is currently applied to any IP interfaces without an explicit Network Egress Queue Group specified, the configuration fails.
		The following configurations are verified when the Network policy is applied to a network IP interface: • The network policy with a redirected queue group cannot be applied to the network IP interface without a port binding. • The redirected queue group name must exist as a Network Egress Queue Group on the port or LAG associated with the IP interface. • The queue ID for the redirected queue group in the associated Network policy must exist within the associated Egress Queue Group Templates policy.

71.2 Workflow to configure access SAP forwarding class-based redirection

71.2.1 Stages

The following workflow describes the steps required to configure access SAP forwarding class-based redirection.

1

Create a global ingress/egress queue group template policy, as required.

a. Create an ingress queue group template policy. See 50.74 “To configure a queue group ingress template policy” (p. 1617) for more information.

b. Create an egress queue group template policy. See 50.75 “To configure a queue group egress template policy” (p. 1619) for more information.

2

Distribute the global ingress/egress queue group template policy to the NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

3

Create an ingress/egress queue group on the Ethernet access port. See 16.37 “To add a queue group to an Ethernet port” (p. 627) for more information.

4

Create or configure a global access ingress/egress QoS policy with the forwarding class mapped to the queue group queue ID. The queue group queue ID must be included in the queue group template policy. Perform the following, as required:

a. Create a SAP access ingress QoS policy. See 50.28 “To configure a SAP access ingress

policy” (p. 2656) for more information.

b. Create a SAP access egress QoS policy. See 50.30 “To configure a SAP access egress policy” (p. 1550) for more information.

5

Distribute the global access ingress/egress QoS policy to the NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

On the SAP configuration form, assign the access ingress/egress QoS policy to the SAP bound to the port on which the access ingress/egress queue group was created in Stage 3 . Perform one of the following, as required:

a. Assign an access ingress/egress QoS policy to an L2 access interface on a VLL. See 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184) for more information.

b. Assign an access ingress/egress QoS policy to an L2 access interface on a VPLS. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) , 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) , and 77.88 “To create a VPLS I-L2 access interface” (p. 2372) for more information.

c. Assign an access ingress/egress QoS policy to an L3 access interface on an IES. See 78.28 “To configure an L3 access interface on an IES site” (p. 2472) for more information.

d. Assign an access ingress/egress QoS policy to an L3 access interface on a VPRN. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) for more information.

71.3 Workflow to configure network IP interface forwarding class-based redirection

71.3.1 Stages

The following workflow describes the steps required to configure a network IP interface forwarding class-based redirection.

1

Create a new global egress queue group template policy. See 50.75 “To configure a queue group egress template policy” (p. 1619) for more information.

2

Distribute the global egress queue group template policy to the NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

3

Create an egress queue group on the Ethernet network port. See 16.37 “To add a queue group to an Ethernet port” (p. 627) for more information.

4

Create or configure a global network policy with the forwarding class mapped to the queue group queue ID. The queue group queue ID must be included in the egress queue group template policy, which is specified when the network policy is applied to the IP interface. See 50.41 “To configure a QoS network policy” (p. 1567) for more information.

5

Distribute the global network policy to the NEs. See 49.6 “To release and distribute a policy” (p. 1476) for more information.

6

On the network interface configuration form, assign the network policy and the queue group template policy to the network interface. The network egress queue group must be created on the port to which the network interface is bound. The network egress queue group must use the same name as the selected queue group template policy. See 27.17 “To create an L3 network interface on a routing instance” (p. 856) for more information.

71.4 Workflow to configure statistics collection for queue groups

71.4.1 Stages

The following workflow describes the steps required to enable the collection of queue groups statistics.

1

Create or modify an accounting policy. See “To configure an accounting policy” in the *NSP NFM-P Statistics Management Guide* for more information.

2

Configure the Type parameter on the Accounting Policy form to collect one of the following statistics options:

- Combined Queue Group
- Queue Group Octets
- Queue Group Packets

3

Enable the Collect Accounting Statistics parameter on the port’s queue group Physical Port (Edit) form. See 16.37 “To add a queue group to an Ethernet port” (p. 627) for more information about adding an access ingress/egress queue group on an Ethernet access port.

72 Virtual ports

72.1 Overview

72.1.1 Purpose

A 7750 SR can act as a Broadband Network Gateway (BNG), fairly distributing bandwidth among the subscriber host sessions by accounting for the packet encapsulation overhead and ATM bandwidth expansion for each type of broadband session. In this way, subscriber packets are less likely to be dropped downstream in the DSLAM DSL port. Furthermore, the BNG shapes the aggregate rate of each subscriber and the aggregate rate of all subscribers destined to a given DSLAM to prevent congestion of the DSLAM.

In the BNG application, when a set of per FC queues are applied to each subscriber host context, the host per FC queue packet rate is overridden by the rate provided in the Radius access-accept message. This rate represents the ATM rate that will be seen on the last mile and includes the encapsulation offset and the per packet expansion due to ATM segmentation into cells at the BSAN.

In order to enforce the aggregate rate of each destination BSAN, a virtual port must be configured. Virtual ports are scheduling nodes that operate like port schedulers, with the exception that multiple virtual ports can be created on the egress context of an access/hybrid Ethernet port. A virtual port and a port scheduler cannot exist simultaneously on a single port.

Virtual ports can be configured on a port that is a member of a LAG. When a virtual port is created, modified, or deleted on the primary port of a LAG, this action is replicated on all other ports within the LAG. These actions can only be performed on the primary port. When a port is added to a LAG, it must have the same virtual ports defined as the existing ports on the LAG. The name of a virtual port is local to the port on which it is applied, but must be the same for all member ports of a LAG.

Virtual ports are supported on Ethernet ports on IOM3/IMM on the 7750 SR, 7750 SR-c4, and 7450 ESS.

Virtual ports are not supported on the 7750 SR-1 and the 7450 ESS-1, HSMDA Ethernet ports, or VSM MDA.

Virtual port statistics can be collected and displayed for all queues forwarding to a virtual port configured on an egress of an Ethernet port. This is only supported on an 7750 SR or 7450 ESS equipped with a 2 x XP MDA IOM-3. The collected data is available on the Statistics tab of the Egress Scheduling Virtual Port properties form.

SLA Profiles

A subscriber host queue with the port-parent option enabled can be scheduled within the context of a port scheduler policy associated with a port or a virtual port. To specify that a subscriber host queue with the port-parent option enabled be scheduled within the context of a virtual port, the Scheduler Type parameter must be set to Virtual port when configuring an SLA Profile.

See 64.5 “To configure an SLA profile” (p. 1845) for more information.

Subscriber Profiles

The subscriber aggregate rate is adjusted to account for the fixed offset and per packet variable expansion of the last mile for the specific session used by the subscriber host. The adjustment is based on the average frame size.

See 64.4 “To configure a subscriber profile” (p. 1840) for more information.

72.1.2 Contents

Virtual ports procedures

72.2 Workflow to configure and manage virtual ports

72.2.1 Overview

The following workflow describes the steps required to configure and manage virtual ports.

72.2.2 Stages

1

Configure an SLA Profile, as described in 64.5 “To configure an SLA profile” (p. 1845) .

2

Configure a Subscriber Profile, as described in 64.4 “To configure a subscriber profile” (p. 1840) .

3

Configure virtual ports, as required.

a. Configure a virtual port using the navigation tree, as described in 72.3 “To configure a virtual port using the navigation tree” (p. 1973) .

b. Configure a virtual port using the Port QoS form, as described in 72.4 “To configure a virtual port using the Port QoS form” (p. 1975) .

c. Copy a virtual port using the Port QoS form, as described in 72.5 “To copy a virtual port” (p. 1976) .

72.3 To configure a virtual port using the navigation tree

72.3.1 Steps

1

On the Equipment navigation tree, expand *NE*→Shelf→Card Slot→Daughter Card Slot, right-click on an access/hybrid Ethernet port and choose Properties. The Port (Edit) form opens.

This applies to ports on the IOM3 and IMM on the 7750 SR, 7750 SR-c4 and 7450 ESS.

2

Click on the Egress Scheduling Virtual Port tab.

3

Click Create. The Egress Scheduling Virtual Port (Create) form opens.

4

Configure the required general parameters.

5

Select a port scheduler policy in the Port Scheduler Policy panel.

You cannot select a port scheduler policy for the virtual port if you have already selected a port scheduler policy for the associated channel. Likewise, selecting a port scheduler policy for the virtual port prevents you from selecting a port scheduler policy for the associated channel.

6

Select a scheduler policy in the Scheduler Policy panel.

7

If you did not select a scheduler policy in Step 6 , configure the parameters in the Aggregate Rate Limit panel.

The Limit Unused Bandwidth parameter is only available on access ports and on SONET channels when an ATM MDA is used.

8

Configure the Status parameter in the Egress Rate Modify panel.

9

Configure the parameters in the Hardware Aggregate Shaper Scheduler panel.

1. Enter the Displayed Name parameter.
2. Select Enable to monitor the hardware aggregate shaper scheduler mentioned in the previous step.

10

Click on the Host Matching tab.

11

Click Create. The Host Matching (Create) form opens.

12

Configure the Destination String parameter.

13

Save your changes and close the forms.

END OF STEPS

72.4 To configure a virtual port using the Port QoS form

72.4.1 Steps

1

Choose Manage→Equipment→Port QoS from the NFM-P main menu. The Manage Port QoS form opens.

2

Click Add →Virtual Port(s). The Egress Scheduling Virtual Port (Create) form opens.

3

Configure the required general parameters.

4

Select a port scheduler policy in the Port Scheduler Policy panel.

5

Select a scheduler policy in the Scheduler Policy panel.

6

If you did not select a scheduler policy in Step 6 , configure the parameters in the Aggregate Rate Limit panel.

The Limit Unused Bandwidth parameter is only available on access ports and on SONET channels when an ATM MDA is used.

7

Configure the Status parameter in the Egress Rate Modify panel.

8

Click on the Host Matching tab.

9

Click Create. The Host Matching (Create) form opens.

10

Configure the Destination String parameter.

11

Save your changes and close the form.

12

Click on the Targeted Physical Ports tab.

13

Click Add. The Select Site form opens.

14

Select one or more sites.

The count displayed on the form reflects the total number of found ports and may not align with the number of applicable ports displayed in the list.

15

Save your changes and close the forms.

END OF STEPS

72.5 To copy a virtual port

72.5.1 Steps

1

Choose Manage→Equipment→Port QoS from the NFM-P main menu. The Port QoS list form opens.

2

Select Egress Scheduling Virtual Port from the object type drop-down menu.

3

Select a virtual port and click Copy. The Egress Scheduling Virtual Port (Create) form opens with the values of the virtual port you selected to copy.

4

Click on the Targeted Physical Ports tab.

5

Click Add. The Select Site form opens.

6

Click Search. The form displays a list of sites.

7

Select one or more sites.

8

Save your changes and close the forms.

END OF STEPS

73 Customer configuration and service management

73.1 Overview

73.1.1 Purpose

This chapter provides overview and procedural information on customer configuration and service management in the NFM-P.

73.1.2 Contents

Customer configuration and service management

73.2 Overview

73.2.1 General information

The NFM-P allows you to create a customer profile and manage the services the customer subscribes to.

An end user is the recipient of application content that is delivered through a service. The service transports the application content and is owned by a service customer. For example, an end user has high-speed Internet access through a service owned by a service provider who is a customer of the network provider.

On the General tab of a customer properties form, you can configure basic customer information. From the other tabs on the form, you can configure and monitor service objects associated with the customer.

Each customer is associated with an ID that is assigned when the customer account is created. Depending on the NE release, you can also configure a customer name. When configuring a service, you can use the ID or the listed name of the customer to associate a customer with a service.

A customer can own more than one service, but an individual service is owned by only one customer. Two or more services can be joined to form a composite service. The individual services that comprise the composite service can be owned by different customers. Each customer that owns a service in a composite service is associated with the composite service.

73.2.2 Managing the size of a service associated with a customer

A customer service can become very large over time if new sites are continually added. As a result, a service containing thousands of sites or instances may become cumbersome to manage. The NFM-P allows you to easily move sites from one service to another in order to reduce the overall size of a particular service.

Sites can be moved for the following service types supported in NFM-P: VPRN, VPLS, MVPLS, IES, and all VLL types. However, the following restrictions apply:

- You can only move a site between like services, that is, from one VPLS to another VPLS, but not between a VPLS and a VPRN, for instance.
- The services that you are moving sites to and from must have the same Service ID.
- The services that you are moving sites to and from must have the same Customer ID.

If you need to reduce the size of an existing service, you can use another existing service that meets these restrictions. Alternatively, you can create a new service of the same type for this purpose. The service and customer identifiers (name or ID number) of the new service must match those of the existing service.

In addition, if the existing service which you are moving sites from is a member of a composite service, then the newly-created service you are moving the sites to will also belong to the same composite service. This action will only take place after moving the sites that form the composite

service connectors. See 73.8 “To move sites from one service to another” (p. 1984) for more information about how to move sites from one service to another.

73.3 Workflow to create a customer profile and manage customer services

73.3.1 Stages

1

Create a customer profile. See 73.4 “To create a customer profile” (p. 1982) .

a. Configure or modify key customer contact and billing information.

b. Assign or associate equipment or resources to the customer, as appropriate.

2

As required, configure the system preferences of services associated with a customer such as specifying the default service priority or the automatic removal of empty service. See the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide* for more information.

3

Create, modify, or delete existing services such as IES, VPLS, VPRN, or VLL that are associated with customers. See the relevant service chapter in this guide.

4

Monitor or troubleshoot customers, based on SLAs between the customer and the service provider.

a. Retrieve customer information and contact the customer when service problems or maintenance windows occur. See 73.5 “To list the services associated with a customer” (p. 1982) .

b. View a service map for a customer. See 73.6 “To view a service map for a customer” (p. 1983) .

c. Perform diagnostics as appropriate to troubleshoot service problems associated with customers. See the relevant service chapter in this guide.

d. As required, view and modify the inventory of all the services, interfaces, circuits, and other information that are associated with a customer. See 73.7 “To modify and manage customer information” (p. 1983) .

e. If required, reduce the service size associated with a customer by moving sites from one service to another. See 73.8 “To move sites from one service to another” (p. 1984) .

5

As required, delete a customer from the NFM-P database. See 73.9 “To delete customers” (p. 1986) .

Customer configuration and service management procedures

73.4 To create a customer profile

73.4.1 Steps

1

Choose Manage→Service→Customers from the NFM-P main menu.

2

Click Create. The Customer (Create) form opens.

3

Configure the required general parameters.

4

Save the changes and close the forms.

END OF STEPS

73.5 To list the services associated with a customer

73.5.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens with Service (Service Management) selected in the filter list.

2

Sort the Customer Name column to filter the list results.

3

Click Search. A list of services matching the filter criterion appears.

END OF STEPS

73.6 To view a service map for a customer

73.6.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Use the filter for the Customer Name column to filter the list results.

3

Click Search. A list of services matching the filter criterion appears.

4

Choose a service from the list and click Topology View. A dialog box appears.

5

Click Yes. The Service Topology - *Service Name* form opens.

6

Scroll the map to view the topology of the service and the devices used for the service.

7

Close the forms.

END OF STEPS

73.7 To modify and manage customer information

73.7.1 Purpose

View and modify the inventory of all the services, interfaces, circuits, and other information that are associated with a customer.

Note: Starting with NE release 15.1 R1, the Customer Name parameter can only be configured during creation of the customer. It cannot be changed after the policy is created. However, the name can be changed if the NE release is 15.0, R4 or later.

CAUTION

Service Disruption

Ensure configuration changes do not affect customer services. Use the Turn Down button to turn down a service before making any changes that may affect customer traffic.

73.7.2 Steps

1

Choose Manage→Service→Customers from the NFM-P main menu. The Manage Customers form opens.

2

Choose a customer and click Properties. The Customer (Edit) form opens.

3

View or modify the information for the customer.

Each tab lists parameters that you can view or modify, or contains functions that you can perform.

- General—lists the customer ID, name, and contact information
- NSP Analytics Parameters—lists the Apdex and MOS thresholds that the customer can set
- Services—lists the services associated with the customer.
- Sites—lists the PE NEs of the customer
- Aggregation—lists the aggregation schedulers of the customer; see Chapter 50, “QoS policies” for information about creating aggregation schedulers
- Spans—lists the span of control over customers

4

Save the changes and close the forms.

END OF STEPS

73.8 To move sites from one service to another

73.8.1 Purpose

Sites can be moved between services. The following service types are supported: VPRN, VPLS, MVPLS, IES, and all VLL types including Wavence VLL.

The following restrictions apply:

- You can only move a site between like services, that is, from one VPLS to another VPLS, but not between a VPLS and a VPRN, for example.
- The services that you are moving sites to and from must have the same Service ID.
- The services that you are moving sites to and from must have the same Customer ID.
- The OLC state of the service site may change based on the destination service configuration.

See 73.2.2 “Managing the size of a service associated with a customer” (p. 1980) in this chapter, for more information.

Note: When you perform a service size reduction, the NFM-P does not create a composite service for services that have only mesh SDP bindings.

See Chapter 85, “Composite service management” for more information about composite services.

Note: The performance of the NFM-P may be affected if you attempt to move a high number of sites or large sites. If you need to move large sites, move fewer sites at one time.

73.8.2 Steps

1

To move more than 25 sites at a time, or to create composite services after moving sites, go to Step 2 . Otherwise, go to Step 5 .

2

Choose Administration→System Preferences from the NFM-P main menu. The System Preferences form opens.

3

In the Services tab, in the Composite Services panel, enable Auto Discover Spoke, CCAG, SCP, RVPLS Connectors. You can also change the Maximum number of sites that can be moved parameter, to up to 50.

Note: The Auto Discover Composite Services parameter can be modified only by a user with administration privileges.

4

Save the changes and close the form.

5

Perform one of the following:

a. Choose Manage→Service→Services from the NFM-P main menu.

The Manage Services form opens.

b. Choose Manage→Service→Composite Service from the NFM-P main menu.

The Manage Composite Services form opens.

6

Click Search to populate the list.

7

Choose the service from which you want to move the sites or the composite service that the service belongs to and click Properties. The *Service*(Edit) |*Composite Service*(Edit) form opens.

8

In the service navigation tree, expand Sites, right-click on the site you want to move and choose Move to Another Service from the contextual menu. A dialog box appears.

9

Click Yes. The Select Services form opens and displays a list of applicable services to which you can move the selected site.

10

Choose the required service and click OK. A progress message is displayed, followed by a confirmation message, after the move is complete.

The NFM-P creates a composite service if:

- The Auto Discover parameter is enabled.
- There are connections between the moved sites and the remaining sites, if the services are not already part of composite services.

11

To view the composite service that is automatically created by the NFM-P after a successful move, click Properties in the Composite Service panel. The Composite Service (Edit) form opens with the General tab displayed.

12

Close the forms.

END OF STEPS

73.9 To delete customers

Note: You cannot delete customers that have associated services.

73.9.1 Steps

1

Choose Manage→Service→Customers from the NFM-P main menu. The Manage Customers form opens.

2

CAUTION

Service Disruption

Removing a customer deletes all information associated with the customer. Ensure that the correct customer has been selected.

Choose the customer that you want to delete from the list and click Delete. A Confirm dialog box appears.

3

Perform one of the following:

a. Click No to avoid deleting the customer. The Confirm dialog box closes.

b. Click Yes. The customer is deleted from the NFM-P database.

END OF STEPS

74 Residential subscriber management

74.1 Residential subscriber management

74.1.1 Overview

The emergence of residential broadband networks and the availability of multiple service offerings, such as triple play applications create a requirement for greater service differentiation and control at the level of the individual service recipient. The NFM-P provides functions for the efficient provisioning of access, QoS, and security features on IES, VPLS, and VPRN for residential subscribers.

In the context of the NFM-P, a residential subscriber, sometimes called a subscriber, is a unique identifier that associates a group of end-user devices with policies. A subscriber can be associated with multiple SAPs on multiple NEs, and a customer can be associated with multiple subscribers.

A subscriber host, sometimes called a host, is an end-user device, such as a computer, VoIP telephone, or set-top box, that connects to the provider network and receives the service traffic. Hosts with the same subscriber identifier share overall HQoS and accounting characteristics as defined in a customer SLA, but may use QoS policies and queues that differ by the type and class of service offering.

A subscriber is an abstract entity created by the NFM-P when the first end user of a service connects to the service. A subscriber instance is created on the NE through which the end user connects. When multiple end users connect to the same service through multiple NEs that act as sites for the service, one subscriber instance is created on each NE. When the NFM-P instantiates a subscriber on an NE, the policies and profiles associated with the subscriber are automatically distributed to the NE and associated with the new subscriber instance.

Residential subscriber management, also called enhanced subscriber management, supports service delivery models in a routed or bridged configuration, such as one VLAN per host, one VLAN per application, one VLAN for all applications, and one VLAN per service provider per application.

74.1.2 Configuration requirements

The following must be true before you can enable residential subscriber management on a SAP or deploy a profile to an NE.

- For dynamic hosts on the SAP:
  - DHCP snooping is enabled on the upstream SAP or SDP.
  - For IES, DHCP relay is enabled on the downstream SAP.
- The NE has sufficient resources for creating the SLA profile instances, queues, and schedulers.
- For new or existing static hosts on the SAP:
  - Anti-spoofing on the SAP is configured with at least IP matching criteria.
  - A subscriber identification string is assigned.
  - A subscriber profile is assigned.
  - An SLA profile is assigned.

– An IP address is assigned.

74.1.3 Residential subscriber management sample configuration

The figure below shows how residential subscriber management assigns resources to dynamic and static hosts. Many different configurations are possible; the sample portrays the residential subscriber management mechanism rather than a particular service delivery model. For simplicity, the sample consists of one subscriber instance on one SAP, and two types of service offering. Each service offering is available in two classes.

Figure 74-1 Residential subscriber management sample configuration

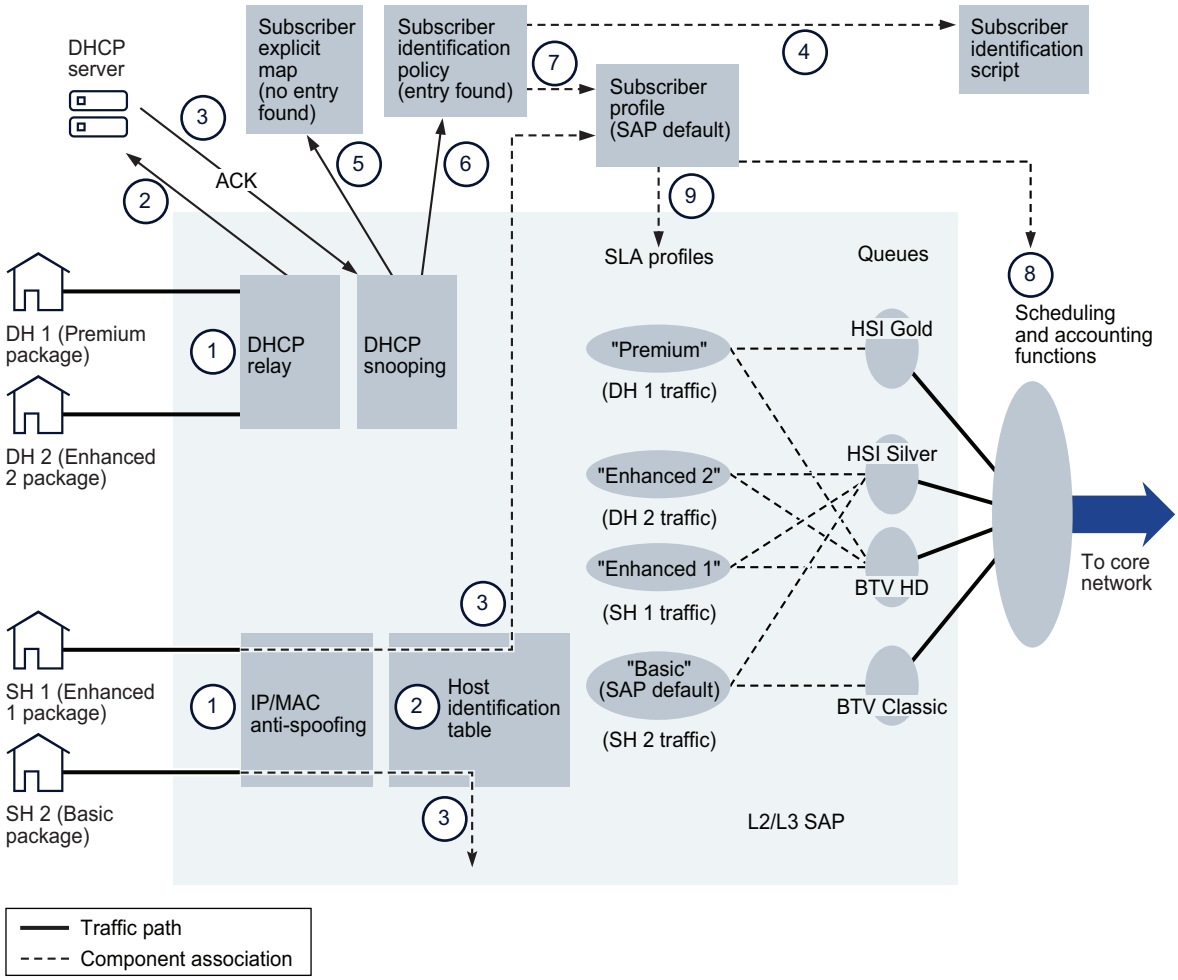

The customer provides a subscriber identification string, the names of the four end-user service packages, and the bandwidth and class requirements for each package. Each package is a combination of HSI and BTV. An NFM-P operator creates four queues and four SLA profiles with

which the queues are associated, as shown in the table below. Nokia recommends that the NFM-P operator give the SLA profiles the same names as the packages to prevent confusion.

Table 74-1 Sample configuration service packages

Subscriber host	Package name	SLA profile name	Associated queues
DH 1	Premium	Premium	HSI Gold, BTV HD
DH 2	Enhanced 2	Enhanced 2	HSI Silver, BTV HD
SH 1	Enhanced 1	Enhanced 1	HSI Gold, BTV Classic
SH 2	Basic	Basic	HSI Silver, BTV Classic

The customer offers two classes of BTV and HSI services:

- BTV HD: high-definition broadcast television
- BTV Classic: regular-definition broadcast television
- HSI Gold: high-bandwidth broadband Internet access
- HSI Silver: regular broadband Internet access

The following steps, which correspond to the numeric labels in the upper part of Figure 74-1, “Residential subscriber management sample configuration” (p. 1990), define the sequence of events for dynamic subscriber hosts that attempt to join the network.

1. Dynamic Host 1 (DH 2) and Dynamic Host 2 (DH 2) each send a DHCP request to the SAP.
2. DHCP relay on the SAP forwards the DHCP requests to the DHCP server.
3. The DHCP server authorizes the requests and responds with a DHCP ACK message for each subscriber host.
4. The subscriber identification policy uses a script to obtain the subscriber identification string, an optional subscriber profile string, and an optional SLA profile string from the Option 82 information in each ACK message. DH 1 and DH 2 provide the same subscriber profile identification string but different SLA profile strings.
5. The NE checks the subscriber identification string values against the entries in the subscriber explicit map and finds no matching entries for the hosts.
6. The NE checks the subscriber profile string and the SLA profile string values for each host against the subscriber identification policy.
7. The NE assigns the same subscriber profile to DH 1 and DH 2 based on the subscriber profile string provided by each host.
8. DH 1 is the first host to join the network, so the scheduling and accounting functions associated with the assigned subscriber profile are instantiated on the SAP.
9. The NE matches the SLA profile string provided by each host to an SLA profile. The NE assigns the Premium SLA profile to DH 1 and the Enhanced 2 SLA profile to DH 2 based on the provided SLA profile strings. These are the first hosts of the subscriber to join the network, so

the appropriate queues are instantiated on the SAP based on the SLA profile specifications, and host traffic subsequently flows.

The following steps, which correspond to the numeric labels in the lower part of Figure 74-1, "Residential subscriber management sample configuration" (p. 1990), define the sequence of operations for static subscriber hosts that join the network.

1. Static Host 1 (SH 1) turns on the computer, and Static Host 2 (SH 2) turns on the television.
2. The host devices request network access; IP- matching (and optional MAC-matching) anti-spoofing on the SAP checks the static host table on the NE and validates both requests.
3. The NE assigns resources to each static host based on subscriber profile and SLA profile designations, and host traffic subsequently flows.
   - The static host table entry for SH1 names a subscriber profile and an SLA profile. Although SH 1 is the first host to use this SLA profile, the queues defined in the profile are already instantiated because of the application of the SLA profiles for DH 1 and DH 2.
   - For SH 2, there is no explicit association between the host and a subscriber profile or an SLA profile, so the NE assigns the SAP default subscriber and SLA profiles, which define the most basic service package the customer offers to end users.

74.1.4 Migrating to the TPSDA model

The NFM-P facilitates the migration of hosts from SAP-based aggregation to the TPSDA model. You can use residential subscriber management functions to:

- Associate a subscriber with a SAP by using the SAP identifier as the subscriber identification string. The existing hosts on a SAP can then be automatically associated with a subscriber.
- Rename a subscriber. This changes the subscriber identification string for all hosts and facilitates a move from a default subscriber identification string to a string that complies with a particular naming scheme.
- Configure an intermediate destination identifier, such as a DSLAM name, for a static host and obtain the identifier from the Option 82 information in a DHCP packet. This function enables the listing of the hosts for a specific DSLAM and facilitates interworking with other TISPAN components.
- Configure a default subscriber identification string for the hosts on a SAP. This enables residential subscriber management functionality without the use of a subscriber identification policy or RADIUS authentication. This default string is associated with hosts when a string is not available from another source.
- Interpret an SLA profile string or subscriber profile string from a host as the profile name when a profile with a matching profile string is not found. This eliminates the need to map a profile string to a profile.

74.2 Residential subscriber components

74.2.1 Overview

Configuring residential subscriber management on the NFM-P involves the configuration of a variety of components, depending on the service delivery model.

For detailed NE-specific information about residential subscriber management, see the NE documentation. For information about configuring residential subscriber policies and profiles, see Chapter 64, “Residential subscriber policies”.

A host requires subscriber-profile and SLA-profile associations to gain access to the network. Profiles define service attributes for hosts such as scheduling, accounting, security, and traffic prioritization by application type. A profile uses existing NFM-P policy definitions and allows the customization of policy parameters using override values.

An NFM-P user who is assigned the administrator, subscriber management, service management, or policy management scope of command role can perform all residential subscriber management functions, such as managing profile or DHCP-lease information.

The NFM-P allows the configuration of multiple components in one operation, but limits configuration to those parameters that are not specific to a component. For example, when a policy is applied to multiple profiles, the NFM-P removes pre-existing policy override values in the profiles.

The figure below shows a conceptual model of the main component relationships in residential subscriber management. The model does not represent a specific type of service or service delivery mechanism and does not show all possible component configurations.

Figure 74-2 Residential subscriber management component relationships

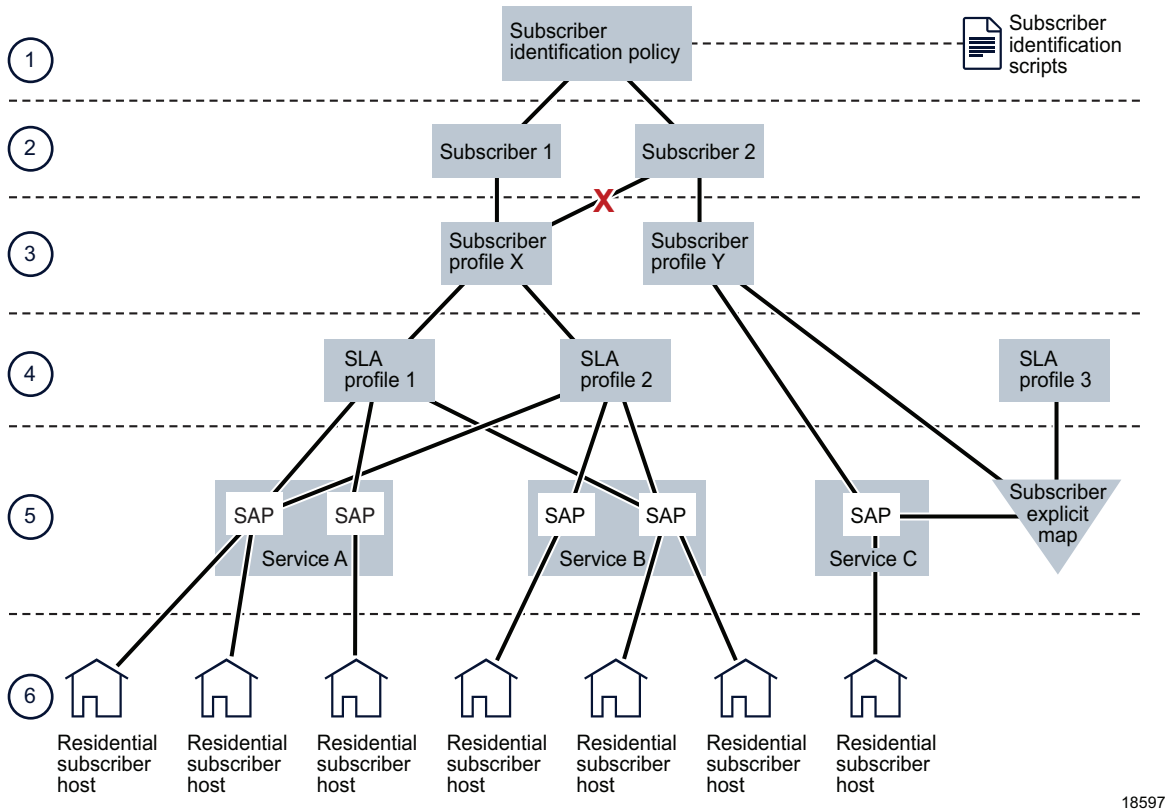

As shown in Figure 74-2, “Residential subscriber management component relationships” (p. 1993), there is a one-to-one relationship between a subscriber and a subscriber profile and a many-to-many relationship between subscriber profiles and SLA profiles. A subscriber profile or an SLA profile can apply to multiple SAPs or NEs, and can be specified as the default profile for the SAP or the NE.

The table below lists where to find configuration information for residential subscriber components. The numbered levels segregate components for clarity only.

Table 74-2 Residential subscriber management components

Leve	Component	Description	For more information, see	Procedure
1	Subscriber identification policy	Associates a dynamic host with a subscriber	64.1.2 “Subscriber identification policies” (p. 1827)	64.3 “To configure a subscriber identification policy” (p. 1838)
	Subscriber identification script	Parses DHCP Option 82 field to extract subscriber identification for a host, and optional profile, ANCP, and intermediate destination identification strings.	64.1.2 “Subscriber identification policies” (p. 1827)	74.19 “To modify the primary subscriber identification script or URL” (p. 2040)
2	Subscriber	Unique identifier that associates a group of subscriber hosts with policies	74.1 “Residential subscriber management” (p. 1989)	—
2	Subscriber instance	An instantiation of a subscriber on an NE	74.1 “Residential subscriber management” (p. 1989)	—
3	Subscriber profile	Names existing ingress and egress scheduler policies and an accounting policy for use by all hosts of a subscriber	64.1.3 “Subscriber profiles” (p. 1828)	64.4 “To configure a subscriber profile” (p. 1840)
4	SLA profile	Names existing QoS policies to define the traffic shaping for different classes of service within an application such as HSI, VoIP, VoD, or BTV Names existing ACL policies to filter the ingress and egress traffic	64.1.4 “SLA profiles” (p. 1829)	64.5 “To configure an SLA profile” (p. 1845)
5	Subscriber explicit map	Lists subscriber identifiers and the associated SLA and subscriber profiles	64.1.5 “Subscriber explicit maps” (p. 1829)	64.6 “To configure a subscriber explicit map entry” (p. 1848)

Table 74-2 Residential subscriber management components (continued)

Leve	Component	Description	For more information, see	Procedure
6	Residential subscriber host (static)	A device such as a computer or set-top box that receives service traffic using a fixed IP address	74.2.7 “Static subscriber hosts” (p. 1998)	74.24 “To create a static host for residential subscriber management on a SAP” (p. 2045) , 78.55 “To configure anti-spoofing filters for an IES L3 access interface” (p. 2501), and79.111 “To configure anti-spoofing filters for a VPRN L3 access interface” (p. 2686) access interface anti-spoofing configuration steps
	Residential subscriber host (dynamic)	A device such as a computer or set-top box that receives the service traffic using a temporarily assigned IP address	64.1.2 “Subscriber identification policies” (p. 1827), 64.1.3 “Subscriber profiles” (p. 1828), 64.1.4 “SLA profiles” (p. 1829), 64.1.5 “Subscriber explicit maps” (p. 1829), 74.4.3 “Managed SAP (MSAP)” (p. 2018), and 64.2.3 “PPPoE policies” (p. 1831)	74.34 “To view and configure residential subscriber hosts on a SAP” (p. 2057)
7	Local DHCP server and PPPoE group interface	A local user database is created and associated to a local DHCP server and PPPoE group interface configuration	74.2.4 “Local DHCP server” (p. 1996)	64.9 “To configure a PPP policy” (p. 1851) 74.9 “To configure a local user database for subscriber host authentication” (p. 2025) 79.5 “To create a VPRN service” (p. 2534) 78.19 “To configure a group interface on an IES” (p. 2449)

The NFM-P treats residential subscriber management components like policies. Depending on the distribution mode configuration of a local instance, when you modify a global component using the NFM-P, all local instances of the residential subscriber management component on the associated NEs can be updated. The Local Definitions tab of a component management form lists the local instances of the component.

Note: You cannot delete an active component of a residential subscriber management configuration.

74.2.2 Residential subscriber component and policy deployment

Global residential subscriber management components are created in draft mode. This allows operators to verify the configuration before they distribute it to the network elements; see Chapter 49, “Policies overview” for information about global and local policy instances and policy distribution modes.

The NFM-P deploys the components associated with a subscriber, such as policies and profiles, to an NE when one of the following occurs:

- The deployment of a subscriber identification policy with profile specifications
- The provisioning of the first static host for the subscriber on a SAP
- The provisioning of a capture SAP with an MSAP Policy to automatically create an MSAP; see 74.4.3 “Managed SAP (MSAP)” (p. 2018)

After the NFM-P deploys the subscriber and SLA profiles, the NE creates the queues, ACL filters and HQoS schedulers specified in the profiles. The NE ignores pre-existing queue or ACL filter policies on a SAP when subscriber management is enabled on the SAP.

74.2.3 Subscriber identification strings

The NFM-P identifies the subscriber for a dynamic host by obtaining a subscriber identification string from a source such as the following:

- information that the host passes in the Option 82 field of a DHCP packet
- a RADIUS server
- the local user database

A static host requires explicit provisioning on the NFM-P that includes the association of a subscriber identification string.

You can optionally specify a default subscriber identification string for a SAP. You can enter the string or configure the NFM-P to use the SAP ID as the string.

Obtaining subscriber identification strings from DHCP packets

When a dynamic residential subscriber host requests an IP address using DHCP, the host can include a subscriber ID string and optional profile ID strings in the Option 82 field of the DHCP packet. If the request is approved by a DHCP server, the NE obtains the subscriber ID string for the host from the Option 82 information in the returning ACK message. The subscriber ID string is a match criterion for the assignment of SLA and subscriber profiles to the host unless profile ID strings are explicitly encoded in the DHCP option information by the host or configured as SAP default profiles. You can also configure a default subscriber ID string on a SAP.

Note: If adding Option 82 information to a DHCP relay packet causes the packet to exceed the DHCP relay maximum of 1500 bytes, the NE forwards the DHCP request without the Option 82 information. For this reason, short strings can be used as aliases for profile names.

74.2.4 Local DHCP server

A local DHCP server leases IP addresses to clients in the network. Options are configured to define the IP address properties, such as, the length of time an IP address is active and which DNS server

must be used. A local user database is used to authenticate and authorize clients requesting IP addresses from the local DHCP server. If the local DHCP server does not use the local user database, the server can use the GI address to assign free IP addresses, however it is not possible to configure match or authentication parameters.

Three applications are targeted for the local DHCP server.

- Subscriber aggregation using a single NE or TPSDA.
- Business services running VPRN and locally attached to the host can request and obtain IP addresses directly from the server.
- The DHCP server identifies an IP request from a PPPoE client and provides an IP address and options.

DHCP servers can be integrated with Enhanced Subscriber Management for DHCP and PPPoE clients. A local DHCP server can be created in the routing instance window or VPRN service site window. A local DHCP server created in the VPRN service site can be associated with the L3 access interface on a VPRN service only. A local DHCP server created in the routing instance window can be associated with a network interface or L3 access interface on an IES.

Failover support

You can configure failover support on local DHCPv4 and DHCPv6 servers, as well as on individual IP address pools on the servers. The same IPv6 prefix or IPv4 address range can be shared between the two peering DHCP servers in a redundant configuration. Either DHCP server can allocate a new lease from the same address range or prefix. To avoid IP address duplication, a single active path is available from the clients to one of the DHCP servers in a redundant configuration. The single active path is ensured through a protection mechanism in a dual homed environment in the access side of the network. The supported protection mechanisms are SRRP or MC-LAG.

Sticky leases

You can configure a sticky lease object on a local DHCPv4 or DHCPv6 server to reserve an IP address, preventing the address from being allocated to another device. The reserved IP address can be defined as a public address for a dynamic host; see 34.22 “To configure an IES or VPRN IPsec gateway” (p. 1256).

74.2.5 Local user database

A local user database is configured and associated with the local DHCP server to provide local authentication. The local DHCP server must have a pool of IP addresses configured, otherwise it is not able to lease IP addresses.

A create local user database configuration form is available from the Manage Residential Subscribers form; see 74.9 “To configure a local user database for subscriber host authentication” (p. 2025). Once a local user database is configured, it can be associated with a local DHCP server and PPPoE configuration on a group interface.

When a local user database is not configured, you can use GI addresses to access free IP addresses, however the clients requesting the IP address will not be authenticated.

74.2.6 ARP host

You can use the NFM-P to configure and retrieve ARP hosts on a 7750 SR. ARP hosts are a subtype of enhanced subscriber host objects. The creation of the object is triggered by the reception of ARP messages from end-user devices. ARP hosts can be configured as an alternative to DHCP subscriber hosts or PPPoE sessions.

ARP host configuration can be performed on VPLS and MVPLS L2 access interfaces, IES and VPRN group interfaces, and VPRN retailer subscriber interfaces. The configuration includes enabling the functionality, setting the maximum limit of hosts per SAP or interface, and setting the default authentication interval.The configuration must be performed in conjunction with other enhanced subscriber management configuration on these interfaces. ARP hosts can be retrieved from the VPLS, MVPLS, and IES or VPRN SAPs. ARP hosts can also reside on MSAPs.

74.2.7 Static subscriber hosts

Static subscriber hosts require explicit provisioning rather than an association with a subscriber identification policy. Static hosts for VPLS are configurable on the anti-spoofing tab of the L2 access interface management form. For IES and VPRN services, static hosts are configurable using the management form for a SAP on a group interface.

Static host configuration involves the following elements:

- IP address
- MAC address (optional)
- subscriber ID string, or the use of the SAP ID as the subscriber ID
- valid subscriber profile
- valid SLA profile
- ANCP string
- intermediate destination, such as a DSLAM
- application profile

Note: When residential subscriber management is enabled on a VPLS SAP that is part of an operational MC ring group, the following must be configured on each static host:

- an intermediate destination that is a ring node in the MC ring group
- subscriber identification, such as a subscriber ID string or the use of the SAP ID as the subscriber ID
- a subscriber profile
- an SLA profile

74.2.8 Subscriber host connectivity verification

Aside from relatively infrequent IP-address lease renewal, DHCP has no session-monitoring or connection-monitoring capability. Residential subscriber management provides this functionality for DHCP hosts and static hosts using subscriber host connectivity verification, or SHCV.

When SHCV is enabled on a SAP, an NE issues a periodic ARP request to each host on the SAP. If the NE receives no reply to an ARP request within the specified interval, the NE raises an event. When SHCV is configured to drop a lost host, the NE immediately removes the host from its active subscriber host table.

SHCV records the state information that it collects for hosts on a SAP and maintains a history of connectivity-related events for troubleshooting purposes. The size of the history log is restricted by a size-constraint policy.

SHCV operates in conjunction with DHCP snooping and is configurable on VPLS, VPRN, and IES SAPs and on IES group interfaces. Because it uses ARP, SHCV automatically populates the FIBs of bridging devices in the access and provider networks.

The configurable items for SHCV on a SAP include:

- frequency of connectivity checks
- source of the ARP request
- action to take when a host loses connectivity

Note: On an L2 access interface, SHCV uses the NE system IP address as the source of the ARP request. On an L3 access interface, SHCV uses the IP- and MAC-address combination of the interface or the VRRP state for the interface as the ARP source. On IES group interfaces, SHCV uses the address of the subscriber interface as the source.

The behavior that an NE exhibits toward SHCV events is configurable using the residential subscriber management form.

The configurable NE SHCV items include:

- maximum host connectivity loss rate
- action to take when the connectivity loss rate exceeds the maximum
- action to take when the NE drops an event trap, as such an event may indicate a high connectivity loss rate

SHCV event handling on an NE is disabled by default. When SHCV is enabled on a SAP but is disabled on the NE, the NFM-P records the blocking of SHCV events in the SHCV log. This ensures that an operator is aware of SHCV activity when viewing the SHCV log, even if the NE is configured to suppress SHCV events.

The following conditions represent an SHCV host connectivity loss:

- absence of a host response
- inconsistency between the reply data and the DHCP lease state

An NE makes more than one SHCV attempt before it raises an event against a host to ensure that the absence of a host response indicates a host connectivity problem and is not simply the result of occasional packet loss.

When a SAP becomes operationally down, the NE generates one trap for the SAP rather than one trap for each host on the SAP. When the SAP returns to service, the NE forwards a trap for each host as it verifies the restoration of connectivity.

SHCV generates the following event traps:

- SAP up
- SAP down
- host up
- host down
- trap dropped

SHCV policies

The SHCV policy provides consistent BNG behavior when handling RG replacement and mobility across different (IP and/or MAC) address types. The SHCV policy is configured on VPRN or IES group interfaces, and VPRN or IES L3 access interfaces.

An SHCV policy specifies uniform Action, Interval, Retry Count, and Timeout parameters for periodic SHCV. The policy also specifies uniform Retry Count and Timeout parameters for each of the following SHCV event triggers:

- Host limit exceeded: the number of hosts specified on the SLA profile is exceeded.
- IP conflict: Can occur when an IP address/prefix is assigned to a SAP.
- Mobility: A MAC or IP address conflict resulting from a UE moving from SAP to SAP.
- MAC learning: Triggered at static host no-shutdown. SHCV attempts to continue until the MAC address is learned for the static host.
- Inactivity: Host inactivity.

To create an SHCV policy, see 64.31 “To configure an SHCV policy” (p. 1872).

74.2.9 Routed CO

A broadband access network typically requires the aggregation of the traffic from access equipment before routing of the traffic is possible. Routed CO functionality allows a network operator to directly connect a DSLAM or similar multiplexer to a router such as the 7750 SR.

Residential subscriber management, combined with the use of DSCP or dot1p values, supports access-integration models such as the following:

- one SAP for all subscribers and service offerings
- one SAP per service offering
- one SAP per subscriber

Routed CO allows the configuration of subscriber interfaces and group interfaces on IES and VPRN. A subscriber interface defines the subnets that are available to a subscriber. A group interface is a child object of a subscriber interface that allows the configuration of multiple SAPs as part of a single interface. Routed CO functionality depends on residential subscriber management to maintain the subscriber host information.

There is no direct association between the subnets and the group interfaces. Subnets can be shared between group interfaces. If the host IP address is not in one of the subnets, packets from the network to that host IP are not received by the subscriber or group interfaces.

A subscriber interface defines a maximum of 16 subscriber subnets and acts as the DHCP relay agent for a subscriber. For the forwarding of DHCP packets to a subscriber server, a subscriber interface requires the specification of a gateway address that is in one of the subscriber subnets. Group interfaces below the subscriber interface inherit this gateway address but can specify an override value if required.

All SAPs in a group interface use the same port. The first SAP on a group interface determines which port the group interface uses. Additional SAPs for the group interface must be associated with the group-interface port during SAP creation.

A group interface is similar to a regular IES interface except for the following.

- Multiple SAPs are configurable.
- No IP addresses are specified.
- No broadcast traffic is permitted.
- No loopback mode exists.
- No static ARP can be specified.
- VRRP is not supported.

When a subscriber interface or a group interface is operationally disabled, no packets are sent to or from the subscriber hosts on the SAPs of the interface, but the state information for a static or dynamic host persists until the static host is removed from the NE by an operator or the DHCP lease of the dynamic host expires.

Because the configuration of multiple SAPs on a group interface makes normal forwarding of DHCP responses to hosts impossible, DHCP relay for routed CO maintains a cache of DHCP requests. If a DHCP server response does not arrive within the specified interval time, the NE discards the cache entry. The NFM-P operator can choose to have the NE use the Option 82 circuit ID field in a DHCP request as the match criterion for the returning DHCP ACK message.

You can configure NAT for dynamic subscriber hosts in a routed CO deployment. NAT for routed CO requires a NAT policy that is associated with a subscriber profile. In an IES routed CO deployment, NAT is configured on the base routing instance of an NE. In a VPRN routed CO deployment, NAT is configured on the VPRN routing instance; see Chapter 30, “NAT” for information about configuring and deploying NAT, and about configuring a NAT policy; see 64.4 “To configure a subscriber profile” (p. 1840) for information about associating a NAT policy with a subscriber profile.

Lease populate is enabled on a group interface by default, and the number of allowable lease entries is set to one. All SAPs on the group interface inherit this configuration during creation.

See Chapter 78, “IES management” and Chapter 79, “VPRN service management” for more information on configuring subscriber and group interfaces for IES and VPRN services.

Wholesale and retail configurations for VPRN services

A VPRN routed CO allows a service provider to resell wholesale carrier services while providing direct DSLAM connectivity. You can create a VPRN service for the retailer and also define

subscriber access and configuration information for the retailer network. The implementation of configuration changes occurs as if the VPRN is a standalone router using the routed CO model. The benefit of this model is flexibility for the retailer and decreased involvement for the wholesaler.

You must use a subscriber SAP to support shared access for multiple retailers. Another dependency of shared access is the requirement for the wholesaler to maintain separate access nodes for each retailer with network scaling issues.

In the wholesale and retail model, the wholesaler instance connections that are common to the access nodes are distributed to many retail instances. Upstream subscriber traffic ingresses into the wholesaler instance and, after identification, is then forwarded into the retail instance. The reverse principle occurs for traffic in the opposite direction. The wholesale and retail traffic flow is controlled with minimal communication to the RADIUS server. A RADIUS policy is defined in the wholesaler instance. The RADIUS response used during the subscriber instantiation provides the service context of the retailer VPRN. If the wholesaler has a retail business, the operator can configure a separate VPRN for their retail services.

The retailer subscriber interface primarily controls the DHCP configuration. The single exception to this model is the lease-populate value. The lease-populate value in the wholesale context controls the individual SAP limits. The lease-populate value in the retail subscriber interface controls the limits for that retailer interface. The wholesale and retail limits must be met before the instantiation of a new subscriber.

Figure 74-3 Routed CO for traditional and wholesale/retail VPRNs

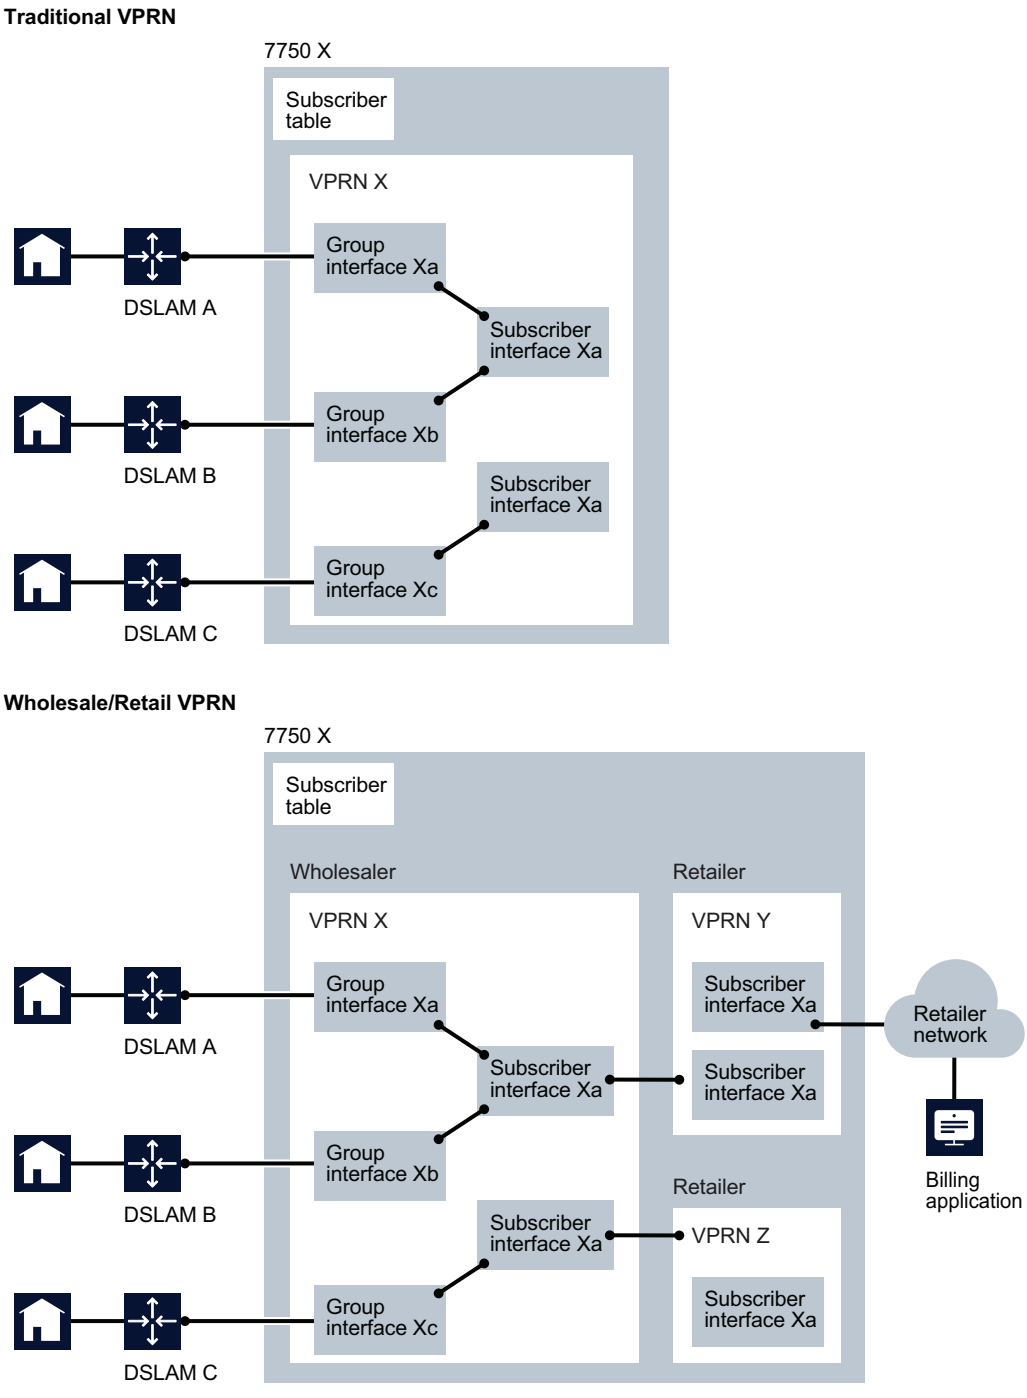

See Chapter 79, “VPRN service management” for information on configuring the forwarding service component for wholesale and retail VPRN services.

74.2.10 Subscriber host polling and monitoring

Subscriber hosts can be periodically polled and monitored for certain DHCP event changes. A subscriber host monitoring configuration form is available from the Manage Residential Subscribers form.

The maximum number of hosts that can be selected for monitoring is configurable through XML; the default is five. Host monitoring has a performance impact, so only a small number are typically monitored at one time.

The Monitored Subscriber Host form displays a variety of polled information about the host that is obtained from the NE. You can also open a DHCP event log for the host from the Host Properties form. The logged events include information on DHCP lease renewals, lease expiration, and profile changes. An entry is added to the table whenever a new event occurs. You can purge event entries using the Host Properties form.

Host monitoring can be started, stopped, or removed from the monitoring configuration form, and the monitoring period and polling interval can be configured.

The Host monitored objects and associated events are stored in the NFM-P database. The monitoring continues until the monitoring period elapses or the monitoring is stopped by an NFM-P operator.

74.2.11 Subscriber services

The NFM-P supports the retrieval of subscriber services and related attributes from PPP or IPoE subscriber hosts. A subscriber service specifies a mode of operation in which a subscriber-specific service is activated and deactivated on a PPPoE session via a RADIUS Access-Accept message or a CoA-Request message. This function is sometimes referred to as Radius Based Policy control, or RaBaPol.

Single or multiple subscriber service activation and/or deactivation calls can be included in Access-Accept messages for new PPPoE sessions, or in CoA-Request messages for existing PPPoE sessions. PPPoE sessions can be IPv4, IPv6 or dual stack. Subscriber services can also operate on the PPPoE session itself.

Examples of subscriber services are:

- apply or change the dynamic bandwidth of a PPPoE session
- apply or change a QoS profile
- activate or deactivate HTTP redirection
- apply or change a rate-limiting service
- a Disconnect Request for a PPPoE session

Figure 74-4 Subscriber services model

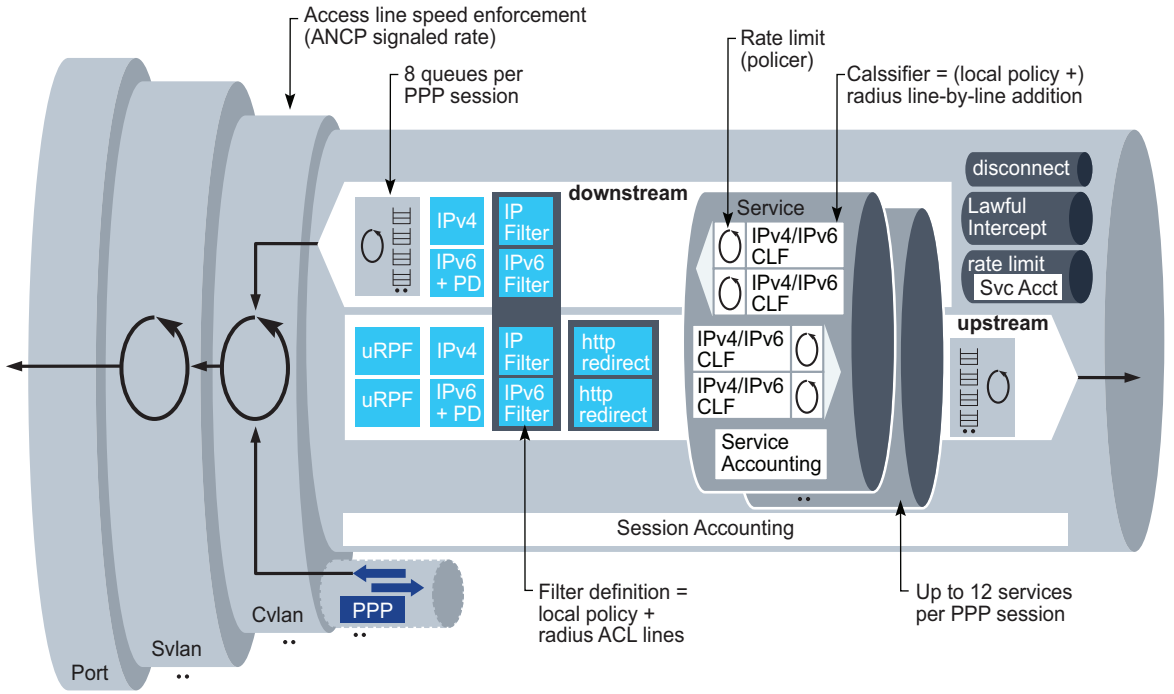

74.2.12 BNG integrated virtual CPE

In a redundant configuration, the vRGW runs in an active/standby redundancy model. When a switch-over occurs, subscriber hosts are re-created via data trigger. In order to avoid assigning the same address to multiple hosts, allocated IP addresses are tracked in the address pool assigned to a subscriber. To support this functionality, the Standby IP Lifetime is configured on the DHCP pool on a BRG profile. Additionally, the BRG instance displays the standby IP address, remaining hold time, and a list of BRG address pools.

74.2.13 Managed routes

The NFM-P can retrieve managed route (also known as framed route) information from the NE for hosts instantiated on IES and VPRN services. A routed subscriber host is a subscriber with subnets behind the routed gateway in order to support subscriber hosts with regular routers. The subscriber host associated subnets are learned through RADIUS VSAs (up to 16 routes).

Framed routes (IPv4 subnets) are displayed on the following ESM session types:

- DHCP lease
- ARP host

- PPP session

Framed IPv6 routes (IPv6 prefixes) are displayed on the following ESM session types:

- DHCPv6 lease
- PPP DHCPv6 session
- PPP SLAAC session
- IPoE SLAAC host (displayed within the subscriber host object)

74.2.14 IPoE sessions

An IPoE session is a logical grouping of multiple subscriber hosts that represent different IP stacks on the same end user device. The grouping is based on a configurable data key, available in a trigger packet. An IPoE session can be single-stack IPv4, single-stack IPv6,or dual-stack, depending which host types or stacks are associated with the session.

The creation of a new IPoE session triggers authentication in the same way as for a subscriber host. A unique accounting session ID is generated for the IPoE session. The authentication data is stored in the IPoE session state for future reference.

Subscriber hosts with a matching key in the trigger packet are associated with an existing session without re-authentication: the data stored in the session state is used to instantiate the subscriber host. Deleted subscriber hosts are no longer associated with the IPoE session.

During the lifetime of an IPoE session, the state is maintained by control plane events for the associated IP stacks (such as DHCP lease state renewals or SHCV checks). Mid-session changes can occur.

An IPoE session is terminated when the last associated IP stack is deleted, or by policy control. The corresponding IPoE session state is then deleted from the system.

The NFM-P specifies IPoE session trigger packet information and session termination control by means of an IPoE session policy. The IPoE session policy is bound to a VPLS capture SAP, and to an IES or VPRN service group interface.

74.2.15 HSQ

HSQ (High Scale QoS) is an evolution of egress Q chip traffic management functions on an enhanced version 4 IOM (iom4-e-hs). The functionality of HSQ is similar to the HSMDAv2, with expanded queue and shaping scale, and an added level of aggregate shaping. Because HSQ functions at the egress forwarding plane level, all egress ports on supported MDA types utilize HSQ QoS functions. Ingress ports on attached MDAs are provisioned using the standard HQoS (Hierarchical Quality of Service) feature set. HSQ supports ports, LAGs, and PW ports.

The following services support HSQ:

- VPRN
- IES
- VPLS, B-VPLS, I-VPLS, and MVPLS
- VLL Epipe and Ipipe

HSQ supports two SLA modes for ESM: Expanded and Single. The SLA mode is configurable on the subscriber profile. In Single mode, the subscriber aggregate rate and the SLA profile queuing are both performed using a single CVLAN context for the subscriber.

A unique SLA profile instance is required when one or both of the following conditions occur:

- A subscriber host references a unique SLA profile name not currently associated with an existing subscriber host.
- A subscriber host is associated with a different egress SAP than an existing subscriber host.

In Single mode, the first condition causes the new SLA profile context to alter the existing SLA profile instance associated with the subscriber. The second condition is not supported for a subscriber configured in Single mode. In Expanded mode, multiple SLA profile instances are supported by moving the subscriber aggregate shaping function to an SVLAN specifically created for the subscriber instance. A unique CVLAN is created for each unique SLA profile instance for the subscriber.

For configuration information, see 74.3.5 “Workflow to configure HSQ” (p. 2013).

74.2.16 ESM over GTP hybrid access

Hybrid access technology provides higher bandwidth to areas where there is limited network infrastructure (typically copper wire-based), and provides the option for the service provider to make a smooth transition to LTE technology in the future. GPRS Tunnelling Protocol (GTP) is used to terminate wireless connections on S11 interfaces, allowing the service provider to use ESM over GTP on hybrid access (WLAN GW and BNG). GTP can be used with and without bonding. The Access Point Name (APN) policy defines subscribers with an access point name and authentication parameters. GTP is configured on VPRN and base routing instances, including APN policy configuration on S11 interfaces.

For configuration information, see 74.3.6 “Workflow to configure ESM over GTP hybrid access” (p. 2014).

74.2.17 Home LAN extension

In a home LAN extension configuration, a bridge domain is created for each residential subscriber on the BNG, the bridge domain is extended by a VXLAN tunnel to the VM in the domain controller. All L3 routing and services on a residential CPE are moved out of the CPE to run in the service provider’s network. These L3 services can be integrated on the BNG or run in a VM per CPE in the data-center. The home CPE runs in bridged mode. It provides local-bridging for traffic between devices on the home LAN, and any traffic destined to an address that is not on the local subnet is bridged towards the BNG. From BNG’s perspective this looks like a bridged residential subscriber with ESM and L2-aware NAT for routed traffic. Each device on the home LAN gets a unique address from the DHCP Server or DHCP proxy on the BNG. In this manner, the VM appears to be on the home LAN.

For more information, see 74.3.8 “Workflow to configure home LAN extension functionality in a network” (p. 2016).

74.2.18 ISA service chaining

Value Added Service (VAS) is a set of service functions in a data center (DC) providing optional services to ESM hosts on a vRGW or BNG. These service functions in the DC can live on virtual machines (VMs) behind a network virtualization edge (NVE) or a third party NVE. ESM hosts corresponding to residential subscribers on the vRGW or BNG (and corresponding NAT outside IP addresses) can be in a different subnet than the VMs in the DC. The vRGW or BNG can access the VAS in DC over an EVPN or VxLAN overlay.

Based on ingress/egress VAS filters associated with L2-aware subscribers, traffic is steered towards service functions (VMs or appliances behind an NVE in an EVPN-enabled data center). The vRGW/BNG shares a local subnet and an overlay routing context (VPRN) with the NVE (behind which the service functions sit), and acts as the next-hop router for the NVEs to reach the NAT outside IP addresses on the vRGW/BNG. NAT outside service can be same as the overlay routing context shared with the NVEs in the DC, or can be a different context. In case it is a different context, NAT routes are leaked into the overlay VPRN, and announced via EVPN prefix routes.

Traffic after NAT on the vRGW/BNG is directed to the MAC address of the service function, resolved by tracking EVPN updates, and is VxLAN tunneled towards VTEP on an NVE in the DC. The NVE bridges this traffic to the service function, or traffic after NAT on vRGW/BNG is directed to the MAC address of the NVE (resolved by tracking appropriate EVPN routes), and is VxLAN tunneled towards VTEP on the NVE. The NVE routes the traffic in appropriate routing context (based on the VNI) to the VM/VA sitting behind it.

Workflow to configure ISA service chaining

Complete the following procedures to configure ISA service chaining:

1. Configure an NE with a MAC prefix for ISA service chaining; see 12.27 “To configure ISA service chaining on an NE” (p. 363).
2. Configure the base routing instance for ISA service chaining, with ISA groups and VXLAN VTEP ranges; see 27.2 “To configure a routing instance or a VRF instance” (p. 826).
3. Configure a service chaining EVPN policy, see 64.33 “To configure a service chaining EVPN policy” (p. 1873).
4. Configure a service chaining VAS Filter policy, see 64.34 “To configure a service chaining VAS filter policy” (p. 1874).

74.3 Residential subscriber management configuration workflows

74.3.1 Workflow to manage residential subscribers

1

Create a service that effectively delivers the proposed residential service offering; see Chapter 70, “Service management and QoS”.

2

Depending on your service delivery model, configure the following policies as required:

a. subscriber identification policy; see 64.3 “To configure a subscriber identification policy” (p. 1838).

b. ANCP policy; see 64.7 “To configure an ANCP policy” (p. 1849).

c. PPP policy; see 64.9 “To configure a PPP policy” (p. 1851).

d. MSAP policy; see 64.10 “To configure an MSAP policy” (p. 1852).

e. host tracking policy; see 64.12 “To configure a host tracking policy” (p. 1855).

f. category map policy; see 64.13 “To configure a category map policy” (p. 1855).

g. credit control policy; see 64.14 “To configure a credit control policy” (p. 1857).

h. IGMP policy; see 64.15 “To configure an IGMP policy” (p. 1858).

i. BGP Peering policy; see 64.17 “To configure a BGP Peering policy” (p. 1859).

j. diameter policy; see 64.18 “To configure a diameter policy” (p. 1860).

k. diameter application policy; see 64.26 “To configure a diameter application policy” (p. 1867).

l. subscriber multicast CAC policy; see 64.19 “To configure a subscriber multicast CAC policy” (p. 1861).

m. mobile gateway/peer profile; see 64.20 “To configure a mobile gateway/peer profile” (p. 1862).

n. host lockout policy; see 64.21 “To configure a host lockout policy” (p. 1863).

o. RADIUS script policy; see 64.23 “To configure a RADIUS script policy” (p. 1865).

See Chapter 64, “Residential subscriber policies” for more information about policies related to residential subscribers.

3

Create ingress and egress scheduler policies for the subscriber hosts in accordance with the customer SLA; see 50.55 “To configure a scheduler policy” (p. 1596).

4

Create an accounting policy for the customer; see the *NSP NFM-P Statistics Management Guide*.

5

Create access ingress and access egress QoS policies for the different applications, for example, HSI and VoIP, and levels of service, for example, gold, silver, and bronze, that subscriber hosts are to receive; see 50.28 “To configure a SAP access ingress policy” (p. 1538) and 50.30 “To configure a SAP access egress policy” (p. 1550). Associate the queues in the QoS policies with the previously created ingress and egress scheduler policies.

6

Specify a unique subscriber identification string for the subscriber, according to your customer specifications; see 74.2.1 “Overview” (p. 1992) for the different options for specifying a subscriber identification string.

7

Create a RADIUS-based accounting policy for the subscriber, as required; see 57.3 “To configure a RADIUS-based accounting policy” (p. 1793).

8

Create a policer control policy for the subscriber, as required; see 50.62 “To configure a policer control policy” (p. 1605).

9

Create SLA profiles that name the previously created QoS policies; see 64.5 “To configure an SLA profile” (p. 1845).

1. Create a different SLA profile for each class of service offering.
2. Use override values to customize the policy values, as required.
3. Use override values to customize the policer values, as required.

10

Create a subscriber profile; see 64.4 “To configure a subscriber profile” (p. 1840).

1. Choose the previously created ingress and egress scheduler policies.
2. Choose the previously created accounting policy.
3. Enable accounting.
4. Assign the previously created RADIUS-based accounting policy, as required.
5. Assign the previously created ingress and egress policer control policies, as required.
6. Associate one or more of the previously created SLA profiles with the subscriber profile.

11

Create, configure, and manage SAPs, as required.

a. Enable automatic generation of subscriber IDs; see 74.5 “To enable automatic generation of subscriber IDs” (p. 2021).

b. Assign the subscriber profile and SLA profiles to the service SAP; see 74.22 “To configure residential subscriber management components on a SAP” (p. 2043).

c. Enable or disable residential subscriber management on a SAP; see 74.23 “To enable or disable residential subscriber management on a SAP” (p. 2044).

d. Specify a subscriber identifier, subscriber profile, and SLA profile for each static host, as required, if default values are not configured on the SAP; see 74.24 “To create a static host for residential subscriber management on a SAP” (p. 2045).

e. Configure a MEP on a SAP; see 74.25 “To configure a MEP on a SAP” (p. 2047).

12

Create, configure, and manage MSAPs, as required.

a. Create a capture SAP that will enable the creation of an MSAP; see 74.26 “To configure a capture SAP” (p. 2049).

b. List MSAPs and view MSAP properties; see 74.27 “To list MSAPs and view MSAP properties” (p. 2052).

c. View an MSAP event log, modify the global MSAP log policy, and purge MSAP log records; see 74.28 “To view an MSAP event log, modify the global MSAP log policy, and purge MSAP log records” (p. 2052).

d. Modify and re-evaluate an MSAP policy on an MSAP; see 74.29 “To modify and re-evaluate an MSAP policy on an MSAP” (p. 2053).

e. Modify an MSAP policy and re-evaluate the MSAPs; see 74.30 “To modify an MSAP policy and re-evaluate the MSAPs” (p. 2054).

13

Configure or monitor DHCP or NE SHCV events or databases, as required.

a. Configure DHCP event monitoring for a SAP; see 74.32 “To configure DHCP event monitoring for a SAP” (p. 2056).

b. Monitor DHCP events for a SAP; see 74.33 “To monitor DHCP events for a SAP” (p. 2056).

c. Configure DHCP lease management for a subscriber host; see 74.6 “To renew or terminate a DHCP lease on a subscriber host” (p. 2022).

d. Configure DHCP event monitoring for a subscriber host; see 74.7 “To configure DHCP event monitoring for a subscriber host” (p. 2023).

e. Monitor DHCP events for a subscriber host; see 74.8 “To manage DHCP event monitoring for a subscriber host” (p. 2024).

f. Configure a local user database to authenticate DHCP clients; see 74.9 “To configure a local user database for subscriber host authentication” (p. 2025).

g. Configure NE SHCV event handling; see 74.13 “To configure NE SHCV event handling” (p. 2034).

14

Turn up the service.

15

Provide the customer with the necessary IP information for provisioning dynamic and static hosts in the customer network.

16

Perform one or more of the following, as required.

a. Create a subscriber explicit map, if required; see 64.6 “To configure a subscriber explicit map entry” (p. 1848).

b. Configure a MEP on an SDP Binding; see 74.14 “To configure a MEP on an SDP Binding” (p. 2035).

c. Configure L2Aware static port forwarding on a subscriber instance; see 74.15 “To configure L2Aware static port forwarding on a subscriber instance” (p. 2036).

d. Resynchronize static port forwarding entries; see 74.16 “To resynchronize static port forwarding entries” (p. 2037).

74.3.2 Workflow to configure IPoE sessions

1

Configure an IPoE session policy; see 64.16 “To configure an IPoE session policy” (p. 1859) .

2

Configure a VPLS capture SAP with IPoE session information. Associate the capture SAP with the IPoE session policy configured in Stage 1; see 74.26 “To configure a capture SAP” (p. 2049).

3

Configure IPoE session information, including the IPoE session policy configured in Stage 1, on the following service objects, as required:

- IES group interface (78.19 “To configure a group interface on an IES” (p. 2449))
- VPRN group interface (79.37 “To configure a group interface on a VPRN” (p. 2586))

74.3.3 Workflow to configure data-triggered host creation on IPoE sessions

1

Configure a local user database IPoE host with a Match List type of IP, and a host identification option of IP Address Prefix; see 74.10 “To configure IPoE hosts on a local user database” (p. 2026).

2

Configure a VPLS capture SAP with a Trigger Packet type of Data; see 74.26 “To configure a capture SAP” (p. 2049).

3

Configure the Data-triggered Administrative State parameter on the following service objects, as required:

- IES group interface (78.19 “To configure a group interface on an IES” (p. 2449))
- VPRN group interface (79.37 “To configure a group interface on a VPRN” (p. 2586))

74.3.4 Workflow to configure diameter NASREQ

1

Configure one or more diameter peer policies, specifying the NASREQ application type; see 57.13 “To configure a diameter peer policy” (p. 1804) .

2

Configure a diameter application policy, specifying the NASREQ application type. Associate the diameter application policy with the diameter peer policies configured in Stage 1; see 64.26 “To configure a diameter application policy” (p. 1867).

3

Configure the diameter application policy (as a Diameter Authentication Policy) on the following objects, as required:

- Local user database IPoE host (74.10 “To configure IPoE hosts on a local user database” (p. 2026))
- Local user database PPP host (74.11 “To configure PPP hosts on a local user database” (p. 2030))
- VPLS capture SAP (74.26 “To configure a capture SAP” (p. 2049))
- IES group interface (78.19 “To configure a group interface on an IES” (p. 2449))
- VPRN group interface (79.37 “To configure a group interface on a VPRN” (p. 2586))

74.3.5 Workflow to configure HSQ

1

Configure the HS Egress Aggregate Rate Limit and HS SLA Profile Handling Mode parameters on a subscriber profile; see 64.4 “To configure a subscriber profile” (p. 1840).

2

Configure an SLA profile with the following (see 64.5 “To configure an SLA profile” (p. 1845)):

- HS Aggregate Rate Limit and HS Queue Stats Mode parameters
- Access egress queue override with Override WRR Weight and Override Class Weight queue override parameters, and an HS override slope policy
- HS WRR group override with Override PIR and Override Class Weight parameters

3

Configure an L2 or L3 access interface with the following (see 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654)):

- HS WRR group override with Override PIR and Override Class Weight parameters

- Access egress queue override with HS WRR Weight, HS Class Weight, and Burst Limit queue override parameters, and an HS WRED queue slope policy

4

Configure an HS secondary shaper on an access interface on any of the following services:

- Service tunnel; see 33.10 “To configure a service tunnel” (p. 1198).
- VLL; see 76.42 “To assign ingress and egress QoS policies to a VLL L2 access interface” (p. 2181).
- VPLS; see 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340).
- VPLS or MVPLS B-L2 access interface; see 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366).
- VPLS I-L2 access interface; see 77.88 “To create a VPLS I-L2 access interface” (p. 2372).
- IES group interface; see 78.21 “To configure a SAP on an IES group interface” (p. 2457).
- VPRN group interface; see 79.39 “To configure a SAP on a VPRN group interface” (p. 2594).

74.3.6 Workflow to configure ESM over GTP hybrid access

1

Configure an APN policy, specifying an authentication type; see 64.38 “To configure an APN policy” (p. 1878).

2

As required, configure any of the following policy types with the appropriate attributes:

- Subscriber authentication policy RADIUS attributes: APN, IMEI, MSISDN, RAT Type, ULI, GPRS Negotiated QoS Profile; see 57.11 “To configure a subscriber authentication policy” (p. 1802)
- RADIUS accounting policy RADIUS attributes: MSISDN, IMEI, APN; see 34.15 “To configure a RADIUS accounting policy” (p. 1244)
- Diameter application policy Gx diameter attributes: PDN Connection ID, APN AMBR, User Location Info, SGSN MCC MNC, RAI; see 64.26 “To configure a diameter application policy” (p. 1867)
- Diameter application policy NASREQ diameter attributes: IMEI, User Location Info, RAT Type; see64.26 “To configure a diameter application policy” (p. 1867)
- Diameter application policy Gx 3GPP QoS downlink mapping types: arbiter, policer, queue, scheduler, aggregateRate, and hsSlaAggregateRate; see64.26 “To configure a diameter application policy” (p. 1867)
- Diameter application policy Gx 3GPP QoS uplink mapping types: arbiter, policer, queue, and scheduler; see 64.26 “To configure a diameter application policy” (p. 1867)

3

Configure a mobile gateway/peer profile with an Interface Type of S11, and MME node parameters (QoS, Dowlink, Uplink, and AMBR); see 64.20 “To configure a mobile gateway/peer profile” (p. 1862).

4

Configure subscriber management extensions on an FPE; see 12.41 “To create an FPE” (p. 374).

5

Configure GTP on a VPRN site; see 79.16 “To configure GTP on a VPRN site” (p. 2549).

6

Configure GTP on a base routing instance; see 27.14 “To configure GTP on a routing instance” (p. 853).

7

Configure a VPRN group interface as Interface Type GTP, and configure GTP settings; see 79.37 “To configure a group interface on a VPRN” (p. 2586).

8

Configure an IES group interface as Interface Type GTP, and configure GTP settings; see 78.19 “To configure a group interface on an IES” (p. 2449).

74.3.7 Workflow to configure ESM connection bonding

1

Configure subscriber management extensions on an FPE; see 12.41 “To create an FPE” (p. 374).

2

Configure a VPRN bonding group interface or an IES bonding group interface; see 79.38 “To create a bonding group interface on a VPRN” (p. 2593) or 78.20 “To create a bonding group interface on an IES” (p. 2456).

3

Configure RADIUS accounting policy with new attributes: Bonding ID, Active Bonding connections; see 34.15 “To configure a RADIUS accounting policy” (p. 1244).

4

Configure a routing policy statement with the From Criteria > Origin parameter set to Bonding; see 54.5 “To configure a routing policy statement” (p. 1745).

5

Configure an IPv4 ACL filter policy entry with the Primary Action parameter set to Forward (Bonding Connection) and Bonding Connection ID parameter configured; see 51.5 “To configure an ACL IP filter policy” (p. 1671).

6

Configure an IPv6 ACL filter policy entry with the Primary Action parameter set to Forward (Bonding Connection) and Bonding Connection ID parameter configured; see 51.6 “To configure an ACL IPv6 filter policy” (p. 1677).

7

Configure an SLA Profile with Bonding Selection parameters; see 64.5 “To configure an SLA profile” (p. 1845).

74.3.8 Workflow to configure home LAN extension functionality in a network

1

Configure a RADIUS accounting policy with the following attributes: LAN Extension Route Distinguisher, LAN Extension Route Target, LAN Extension Bridge ID, LAN Extension Route Device Type, LAN Extension Route VNI; see 57.3 “To configure a RADIUS-based accounting policy” (p. 1793).

2

Configure an ISA RADIUS policy with the following attributes: Cross-connect Tunnel Local IPv6 Address, Cross-connect Tunnel Remoute IPv6 Address, Cross-connect Tunnel Service, Cross-connect Tunnel Type: see 57.12 “To configure an ISA RADIUS policy” (p. 1803).

3

Configure an NE with a router target AS number under Subscriber Management; see 12.26 “To configure home LAN extension functionality on an NE” (p. 362).

4

Enable home LAN extension functionality on a base routing instance. Configure home LAN extension parameters on the routing instance and cross-connect parameters on the WLAN GW; see 27.2 “To configure a routing instance or a VRF instance” (p. 826).

5

Configure Cross-Connect parameters on a VPRN site; see 79.65 “To configure WLAN GW functionality on a VPRN site” (p. 2630).

6

Configure home LAN extension and cross-connect parameters on a WLAN GW on a VPRN group interface; see 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631).

7

Configure home LAN extension and cross-connect parameters on a WLAN GW on a IES group interface; see 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462).

74.4 SAP and MSAP management overview

74.4.1 General information

You can use residential subscriber profiles and policies to manage SAPs. Enabling residential subscriber management on a SAP does not affect an existing dynamic host on a SAP until the host DHCP lease expires, at which point a subscriber identification policy manages lease renewal. Static hosts require the provisioning of a subscriber identification string for continued network access.

When residential subscriber management on a SAP is enabled, each subscriber host that connects to the SAP requires a subscriber identification string to associate the host with a subscriber instance on the NE, or network access is denied. A subscriber identification string uniquely identifies a subscriber in the NFM-P.

The maximum number of subscribers that can use a SAP is configurable. When a SAP is limited to one subscriber, you can specify the treatment of traffic that does not match the subscriber profile, such as BTV traffic, which uses an IP address in the multicast range as the destination address. During single-subscriber SAP configuration, the NFM-P operator specifies whether to allow non-profile traffic and specifies the SLA and subscriber profiles to use for non-profile traffic, if it is allowed.

74.4.2 SAP DHCP monitoring

You can monitor the DHCP events on one or more SAPs. A SAP monitoring form is available from the Manage Residential Subscribers form.

Table 74-3 SAP DHCP monitoring events

DHCP event	Description
DHCP Lease Entries Exceeded	The number of DHCP lease state entries on a SAP has reached the configured maximum.
DHCP Lease State Overridden	An existing DHCP lease state was overridden by a new lease state with the same IP address and a new MAC address.
Suspicious DHCP Packet	A DHCP packet containing suspicious content was received.
DHCP ACK Dropped	A DHCP ACK message was discarded because the related lease state could not be updated.
Host Connectivity Lost	The connection to the specified host was lost.
Host Connectivity Restored	The connection to the specified host was regained.

The maximum number of SAPs that can be selected for monitoring is configurable through XML; the default is five. SAP monitoring has a performance impact, so only a small number are typically monitored at one time.

The Monitored Access Interface form displays a variety of DHCP-related events or notifications about subscriber hosts on a SAP that are compiled using traps received from NEs. You can also open an event log for the SAP from the SAP Properties form. The logged events include information on DHCP lease entries, lease states, and host connectivity. An entry is added to the table whenever a new event occurs. You can purge event entries using the Monitoring Configuration form.

SAP monitoring can be started, stopped, or removed from the monitoring configuration form, and the monitoring period can be configured.

The SAP monitored objects and the associated events are stored in the database. The monitoring continues until the monitoring period elapses or the monitoring is stopped by an NFM-P operator.

74.4.3 Managed SAP (MSAP)

You can configure the NFM-P to automatically create a SAP for a subscriber; this SAP is called a managed SAP, or MSAP. MSAPs are used in network models where each subscriber has a separate VLAN, and multiple subscriber hosts share a SAP. MSAP creation uses the authentication mechanisms that an NE supports to provide an MSAP for the one-subscriber-per-SAP model. An MSAP behaves the same as a regular SAP, but the MSAP configuration is not editable. MSAPs count towards the maximum number of SAPs allowed on an NE.

An MSAP is typically used by only one subscriber, although the NFM-P supports MSAPs shared by multiple subscribers. However, if more than one subscriber is allowed and an MSAP has been defined by a host, when a new host installation attempts to change the MSAP policy the installation fails and raises an event.

The creation of an MSAP is triggered by a capture SAP. When a capture SAP is configured, triggering packets initiate RADIUS authentication, and the RADIUS reply provides a service context. The authentication and the service context are used by the NE to create an MSAP; see 74.4.4 “Capture SAP” (p. 2018).

Although MSAPs are not configurable, a user with a Policy Management or Subscriber Management role can create, list, delete, or modify an MSAP policy to control how the parameters apply during MSAP creation.

An MSAP remains active as long as there is at least one subscriber host on the MSAP; see 74.4.5 “MSAP management” (p. 2019).

MSAPs are supported in HA and dual-homing environments. In dual homing, the MSAP is synchronized with the other NE during MSAP creation. For both HA and dual-homing environments, each participating NE must use the same MSAP policy.

74.4.4 Capture SAP

A capture SAP must be defined to trigger the process that automatically creates an MSAP. A capture SAP does not forward traffic; see 74.26 “To configure a capture SAP” (p. 2049).

You can create a capture SAP only in a VPLS. A capture SAP can create MSAPs in the routed CO configuration of an IES or VPRN, and in a VPLS TPSDA configuration. IES and VPRN MSAPs are

restricted to group interfaces. You can configure the capture SAP to allow Dot1Q MSAP creation when the terminating port for the capture SAP is configured for QinQ encapsulation.

The following triggers are supported to initiate MSAP creation:

- DHCPv4 and DHCPv6 trigger packets—DHCP Discover (or Requests if configured) for DHCP clients; the MSAP lifetime is defined by the lease time.
- PPPoE trigger packets—PPPoE PADI for PPPoE client; the MSAP lifetime is defined by the session time. The MSAP is installed after an IP address is provided.
- IPoE trigger packets—IPoE session key, as defined in IPoE session policy bound to the capture SAP. The MSAP lifetime is also defined on the IPoE session policy.
- ARP trigger packets—ARP for static-ip hosts; the MSAP lifetime is defined by ARP entry time. Current ARP entry refresh behavior is maintained.
- Data trigger packets—receipt of data packets triggers RADIUS authentication and subscriber host creation.

After you configure the capture SAP, every DHCP packet, PPPoE packet, or ARP packet received on the SAP is sent to the CPM, which triggers RADIUS authentication to provide a service context. The MSAP is created in the specified service. Non-triggering packets captured by the capture SAP are dropped.

If RADIUS does not provide all of the required information to install the host (for example, RADIUS lacks the IP address), the MSAP is created with a short timer while waiting for the host to be installed. Default SAP polices are available on the NE and will be used if the MSAP policy is not configured.

The authentication policy used in the capture SAP must be the same as the policy used for the MSAP. For L3, the authentication policy is defined under the group-interface.

Note: When PPPoE is used with MSAPs, the authentication policy must not use the username for MSAP creation.

The MSAP is not created (and an event is generated) if the group-interface name returned from RADIUS points to a different authentication policy than the policy defined by the capture SAP.

For MSAPs, the authentication policy is defined in the MSAP policy. Based on the configuration, the system reauthenticates the sessions when they are renewed. If the authentication policy is not used or when only PPPoE is used, the MSAPs stay active when the session is renewed.

An MSAP is created with dual-pass (shared) queuing. The SLA profile of the host may change the queuing later in the process. An MSAP is always created with the default QoS and scheduling.

74.4.5 MSAP management

An MSAP cannot be modified directly, so MSAP management is performed using MSAP policies. To make changes to an MSAP, alter the MSAP policy for the MSAP and then re-evaluate the MSAP. During re-evaluation, the NE updates the MSAP based on the changes to the MSAP policy; see 74.29 “To modify and re-evaluate an MSAP policy on an MSAP” (p. 2053).

The NFM-P treats the setup and teardown of MSAPs as state transitions to offset the load imposed by constant SAP creation and deletion. There are two states that an MSAP can have, active or inactive. These states cannot be configured. The state is automatically set to active when the

MSAP is created on the NE. When the NE deletes the MSAP, the state changes to inactive, and the MSAP continues to be available in the NFM-P but is not used until it is reactivated by the NE. When the MSAP is reactivated, it will have the same identification so that the NFM-P can identify it with the same FDN.

The state information is not automatically updated in the NFM-P GUI. You must perform a resynchronization to retrieve the current state information. When the state is inactive, performance statistics for the MSAP are not retrievable; however, historical statistical records are available. You can schedule MSAP statistics or collect the statistics on demand; see the *NSP NFM-P Statistics Management Guide* for more information about scheduling and collecting statistics.

Inactive services that previously contained MSAPs consume database space. MSAPs that have been in an inactive state for a long period of time or that are no longer used must be manually deleted.

Note: If there is a disruption on the NE and the NE recovers by restarting, all MSAPs are re-evaluated and any policy changes are applied.

Note: The creation of a SAP that uses the same port and encapsulation values as an existing inactive MSAP fails under the following conditions:

- If you try to use the NFM-P to create a SAP, the configuration fails and the NFM-P displays an error message.
- If you use a CLI to create a SAP in a service other than the service that contains the MSAP, the configuration succeeds but the MSAP remains inactive until the regular SAP is deleted. Although the NFM-P displays the SAP and MSAP, the MSAP remains inactive and consumes resources.
- If you use a CLI to create a SAP in the service that contains the MSAP, the SAP creation fails.

Nokia recommends that you delete an inactive MSAP from the NFM-P if you need to create a regular SAP on the same port using the same encapsulation values.

74.4.6 MSAP Re-evaluate Lease States function

The NFM-P Re-evaluate Lease States function allows you to push SLA profile or subscriber profile changes out to residential subscribers without reinstantiating the hosts; see 74.31 “To re-evaluate lease states for an MSAP” (p. 2055).

74.4.7 MSAP alarm suppression

An IES or VPRN MSAP that is created on a LAG in an SRRP deployment uses an NFM-P mechanism to suppress unnecessary AccessInterfaceDown alarms. The mechanism uses the SRRP states of the redundant BNGs, and determines whether to raise the alarm based on the following rules:

- When the active BNG is in the master SRRP state and the standby BNG is in the initialize or backupShunt SRRP state, no AccessInterfaceDown alarm is raised against the standby BNG.
- If the standby BNG reboots, or if there is a link failure between the standby BNG and the access network, the MSAP AccessInterfaceDown alarms are suppressed for the standby BNG. The

active BNG continues to remain in the master SRRP state. In this scenario, there may be other alarms on the standby BNG for card or link failures.

- After a fault that results in an SRRP failover, for example, active BNG reboot or loss of access network connectivity on the active BNG, an AccessInterfaceDown alarm is raised against each BNG. The alarms clear after SRRP state convergence.
- After a fault that affects both BNGs concurrently and disables SRRP redundancy, an AccessInterfaceDown alarm is raised against each BNG. The alarms clear when a BNG regains connectivity and assumes the role of SRRP master.

Note: The alarm-suppression mechanism can function correctly only when SRRP is correctly configured. If an SRRP peer is missing or disabled, the alarms may not be suppressed for the BNG that the NFM-P perceives to be the standby BNG.

74.4.8 MSAP event logs

An MSAP event log records the date, time, and active state for each state change that occurs after the MSAP is created. An event log is not recorded when the MSAP is created, although the information is recorded in a time stamp; see 74.28 “To view an MSAP event log, modify the global MSAP log policy, and purge MSAP log records” (p. 2052).

74.4.9 MSAP time stamps

You can view the following MSAP state time stamps:

- date and time that the MSAP was created
- date and time of the last active state change

See 74.27 “To list MSAPs and view MSAP properties” (p. 2052).

74.4.10 Workflow to manage MSAPs

1

Create an MSAP policy; see 64.10 “To configure an MSAP policy” (p. 1852).

2

Create a capture SAP that references the MSAP policy you created in Stage 1; see 74.26 “To configure a capture SAP” (p. 2049).

3

Change an existing MSAP by modifying and re-evaluating an MSAP policy, as required. To modify a single MSAP, see 74.29 “To modify and re-evaluate an MSAP policy on an MSAP” (p. 2053); to modify multiple MSAPs that use the same policy, see 74.30 “To modify an MSAP policy and re-evaluate the MSAPs” (p. 2054).

4

Perform the following MSAP maintenance tasks, as required:

a. List MSAPs; see 74.27 “To list MSAPs and view MSAP properties” (p. 2052).

b. View or purge MSAP logs, or modify the MSAP log policy; see 74.28 “To view an MSAP event log, modify the global MSAP log policy, and purge MSAP log records” (p. 2052).

74.5 To enable automatic generation of subscriber IDs

74.5.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Auto Subscriber ID Key (Residential Subscribers) and click Search.

3

Select the site ID of the NE you want to query and click Properties. The Auto Subscriber ID Key (Edit) form opens.

4

As required, disable the Default parameter and configure the following parameters in the PPP Subscriber ID Key panel:

- PPP Key 1
- PPP Key 2
- PPP Key 3
- PPP Key 4
- PPP Key 5

5

As required, disable the Default parameter and configure the following parameters in the IPoE Subscriber ID Key panel:

- IPoE Key 1
- IPoE Key 2
- IPoE Key 3
- IPoE Key 4

6

Save your changes and close the form.

END OF STEPS

74.6 To renew or terminate a DHCP lease on a subscriber host

74.6.1 Purpose

DHCP is used to assign IP addresses to hosts or workstations on the network. This function is usually performed by a DHCP server, which leases out addresses for specific times to the various hosts. If a host does not use a specific address for a set period of time, that IP address can be assigned to another machine by the DHCP server.

In order to allow the DHCP client to lease out an address for specific times, the NFM-P allows an operator to force the DHCP client to change its state by manually sending a Renew command. Upon receipt of the Renew command from the NFM-P, the DHCP client changes to the renewal state and negotiates with the DHCP server for lease times.

Conversely, an operator can also manually terminate (clear) a lease for a particular host from within the Residential Subscriber Host form of the NFM-P.

Note: The DHCP lease renewal and termination are supported only on dynamic hosts; not on static hosts or subscriber hosts.

74.6.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Select Residential Subscriber Host (Residential Subscriber) and search for the SAP on which the subscriber hosts you want to manage are located.

3

Select one or more host entries.

You can renew a lease for multiple entries, but you can only terminate leases one host at a time.

4

Perform one of the following steps:

a. To renew the lease, click Force Renew.

b. To terminate the lease, click Clear Lease State.

5

Confirm the action.

6

Repeat Step 3 and Step 4 to perform additional lease renewals or terminations.

7

Close the form.

END OF STEPS

74.7 To configure DHCP event monitoring for a subscriber host

74.7.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Select Monitored Subscriber Host (Monitor) and search for the SAP on which the subscriber hosts you want to manage are located.

3

Select a host entry and click Create Monitored Host. The Monitored Subscriber Host form opens.

A maximum of five hosts or five SAPs can be monitored simultaneously. If more than the maximum number of hosts is selected and configured for monitoring an error message will be displayed.

4

Configure the required parameters.

5

Save your changes and close the form.

Note: You must complete 74.8 “To manage DHCP event monitoring for a subscriber host” (p. 2024) to enable and schedule subscriber host monitoring.

END OF STEPS

74.8 To manage DHCP event monitoring for a subscriber host

74.8.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Select Monitored Subscriber Host (Monitor) and search for the SAP on which the monitored subscriber hosts you want to manage are located.

3

Select a host entry and click Properties. The Monitored Subscriber Host form opens.

4

Perform one or more of the following:

a. To start monitoring the chosen subscriber host, click Start Monitor.

b. To stop monitoring the chosen subscriber host, click Stop Monitor.

c. To remove monitoring of the chosen subscriber host, click Delete.

d. To change the monitoring period of the chosen subscriber host, click Stop Monitor, type the desired value in the field, and select an appropriate unit from the adjacent menu.

e. To change the Polling Interval of the chosen subscriber host, click Stop Monitor and select the desired value from the menu.

5

Save your changes and close the form.

END OF STEPS

74.9 To configure a local user database for subscriber host authentication

74.9.1 Purpose

Perform this procedure to configure a base local user database object on a local DHCP server. Additionally, you may need to configure IpoE hosts (74.10 “To configure IPoE hosts on a local user database” (p. 2026)) or PPP hosts (74.11 “To configure PPP hosts on a local user database” (p. 2030)) on the local user database.

74.9.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Local User Database (Local User DB) and click Search.

3

Click Create Local User Database or select an existing local user database entry and click Properties. The Local User Database (Create|Edit) form opens.

4

Select a site for the local user database.

5

Configure the required parameters.

6

Click Apply.

7

To configure IPoE hosts for the local user database, see 74.10 “To configure IPoE hosts on a local user database” (p. 2025).

8

To configure PPP hosts for the local user database, see 74.11 “To configure PPP hosts on a local user database” (p. 2030).

9

Save your changes and close the forms.

END OF STEPS

74.10 To configure IPoE hosts on a local user database

74.10.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Local User Database (Local User DB) and click Search.

3

Select a local user database entry and click Properties. The Local User Database (Edit) form opens.

4

Click on the IPoE tab and configure the Match Type IPoE parameters.

You cannot set two parameters to the same value.

5

Click on the Circuit ID Mask tab and configure the required parameters.

6

Click on the Masks tab.

1. Click Create or select an existing IPoE mask entry and click Properties. The Local User Database IPoE Mask form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

7

Click on the Hosts tab.

1. Click Create or select an existing host entry and click Properties. The Local User Database IPoE Host (Create|Edit) form opens.
2. Configure the required parameters.
3. Select a diameter application policy.
4. Select a subscriber authentication policy.

   If you are using LLID pre-authentication, choose a subscriber authentication policy that is associated with the pre-authentication server.
5. Select a diameter authentication policy.
6. Select a RIP policy.
7. In the MSAP Defaults panel, select the following objects:
   - MSAP policy
   - MSAP service ID
   - MSAP group interface

   You can also specify an MSAP group interface by manually typing its name. If the port ID is not included in the typed name, set one of the following parameters to Port ID, depending on whether the port ID is to be the prefix or suffix of the group interface name:
   - MSAP Group Interface Prefix
   - MSAP Group Interface Suffix
8. Select a RADIUS accounting policy and a duplicate RADIUS accounting policy.
9. Configure IPv6 lease times for the IPoE host, if required.

   The Valid Life and Preferred Life parameters can be set to infinite by placing check marks in the Infinite check boxes.

8

Click on the Address tab.

1. Configure the required parameters.
2. If you configure IPv6 parameters, select an IPv6 SLAAC prefix pool.

   You can also manually type an IPv6 SLAAC prefix pool name.

9

Click on the Host Identification tab.

1. Configure up to four of the following parameter groups.
   - Circuit ID Format, Circuit ID
   - MAC Address
   - Option 60 Format, Option 60
   - Remote ID Format, Remote ID
   - SAP ID
   - Service ID
   - IPoE String
   - System ID
   - Derived ID
   - IP Address Prefix, Prefix Length
2. If one of the match criteria specified in Step 4 is Encapsulation Tag Range, configure the Type parameter.

   Configure the following parameters to define the outer encapsulation range:
   - Outer Start Value (Tag1)
   - Outer End Value (Tag1)
3. If you set the Type parameter to Q in Q, configure the following parameters to define the inner encapsulation range:
   - Inner Start Value (Tag2)
   - Inner End Value (Tag2)

10

Click on the Identification Strings tab.

1. Configure the Option Number parameter. The form displays additional parameters.
2. Configure the required identification strings parameters.

11

Click on the WPP tab to configure WPP-based host authentication.

1. Configure the Routing Instance parameter.

2. If you specify the VPRN option in 1, you must configure the VPRN Service parameter. You can manually type a VPRN service ID, or select a VPRN service.
3. Configure the Portal Name parameter. You can manually type a portal name, or select a WPP portal.
4. Configure the Initial Subscriber Profile parameter. You can manually type a subscriber profile name, or select a subscriber profile.
5. Configure the Initial SLA Profile parameter. You can manually type an SLA profile name, or select an SLA profile.
6. Configure the Initial Application Profile parameter. You can manually type an application profile name, or select an application profile.
7. Configure the Restore Default Profile On Disconnection parameter.

12

Click on the Options tab.

13

To configure IPv4 options, perform the following steps:

1. Click Create or select an existing IPv4 option entry and click Properties. The Local User Database DHCP IPv4 Option (Create|Edit) form opens.
2. Configure the Option parameter.
3. Configure any additional parameters that appear. The configurable parameters vary, depending on the Option parameter value.
4. Save your changes and close the form.

14

Click on the IPv6 tab to configure IPv6 options.

1. Click Create or select an existing IPv4 option entry and click Properties. The Local User Database DHCP IPv6 Option (Create|Edit) form opens.
2. Configure the required IPv6 address parameters.
3. Save your changes and close the form.

15

Click on the RADIUS Proxy Cache tab to configure a RADIUS proxy cache for the IPoE host.

1. Configure the Routing Instance parameter.
2. If you set the Routing Instance parameter to Base, select a RADIUS proxy server.
3. If you set the Routing Instance parameter to VPRN, select a VPRN service.
4. Configure the Match Type parameter.
5. Configure any additional parameters that appear. The configurable parameters vary, depending on the Match Type parameter value.

6. Configure the Match Failure Action parameter.

16

Click on the To Client Options tab to configure relay to client options for the IPoE host.

1. Click Create or select an existing options instance and click Properties. The To Client Options (Create|Edit) form opens.
2. Configure the required parameters.
3. Configure any additional parameters that appear in the Option Value panel. The configurable parameters vary, depending on the Type parameter value.
4. Save your changes and close the form.

17

Save your changes and close the forms.

END OF STEPS

74.11 To configure PPP hosts on a local user database

74.11.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Local User Database (Local User DB) and click Search.

3

Select a local user database entry and click Properties. The Local User Database (Edit) form opens.

4

Click on the PPP tab and configure the Match Type PPP parameters.

You cannot set two parameters to the same value.

5

Click on the Masks tab.

1. Click Create or select an existing PPPmask entry and click Properties. The Local User Database PPP Mask form opens.
2. Configure the required parameters.

You must configure the Mask Type parameter before you configure the remaining parameters.

3. Save your changes and close the form.

6

Click on the Hosts tab.

1. Click Create or select an existing host entry and click Properties. The Local User Database PPP Host (Create|Edit) form opens.
2. Configure the required parameters.
3. Select a diameter application policy.
4. Select a subscriber pre-authentication policy.
5. Select a subscriber authentication policy.
6. Select a diameter authentication policy.
7. Select a RIP policy.
8. Select a user database.
9. In the MSAP Defaults panel, select the following objects:
   - MSAP policy
   - MSAP service ID
   - MSAP group interface

   You can also specify an MSAP group interface by manually typing its name. If the port ID is not included in the typed name, set one of the following parameters to Port ID, depending on whether the port ID is to be the prefix or suffix of the group interface name:
   - MSAP Group Interface Prefix
   - MSAP Group Interface Suffix
10. Select a RADIUS accounting policy and a duplicate RADIUS accounting policy.
11. Configure access loop parameters for the PPP host, if required.

    The Circuit ID parameter must be configured if the Circuit ID Type parameter is set to ASCII. The Remote ID parameter must be configured if the Remote ID Type parameter is set to ASCII.
12. Configure IPv6 lease times for the PPP host, if required.

    The Valid Life and Preferred Life parameters can be set to infinite by placing check marks in the Infinite check boxes.

7

Click on the Address tab.

1. Configure the required parameters.
2. If you configure IPv6 parameters, select an IPv6 SLAAC prefix pool.

   You can also manually type an IPv6 SLAAC prefix pool name.

8

Click on the Host Identification tab.

1. Configure up to three of the following parameter groups:
   - Circuit ID
   - MAC Address
   - Remote ID
   - Service Name
   - SAP ID
   - Derived ID
2. Configure the parameters in the User panel.
3. If one of the match criteria specified in Step 4 is Encapsulation Tag Range, configure the Type parameter.

   Configure the following parameters to define the outer encapsulation range:
   - Outer Start Value (Tag1)
   - Outer End Value (Tag1)
4. If you set the Type parameter to Q in Q, configure the following parameters to define the inner encapsulation range:
   - Inner Start Value (Tag2)
   - Inner End Value (Tag2)

9

Click on the Identification Strings tab.

1. Configure the Option Number parameter. The form displays additional parameters.
2. Configure the required identification strings parameters.

10

Click on the L2TP tab to select a tunnel group.

L2TP configuration is required on an NE that has the LAC role if RADIUS authentication is not used for PPP clients.

11

Click on the Options tab.

12

To configure IPv4 options, perform the following steps:

1. Click Create or select an existing IPv4 option entry and click Properties. The Local User Database DHCP IPv4 Option (Create|Edit) form opens.
2. Configure the Option parameter.
3. Configure any additional parameters that appear. The configurable parameters vary, depending on the Option parameter value.

4. Save your changes and close the form.

13

Click on the IPv6 tab to configure IPv6 options.

1. Click Create or select an existing IPv4 option entry and click Properties. The Local User Database DHCP IPv6 Option (Create|Edit) form opens.
2. Configure the required IPv6 address parameters.
3. Save your changes and close the form.

14

Click on the Access Loop tab.

1. Click Create or select an existing access loop entry and click Properties. The Access Loop (Create|Edit) form opens.
2. Configure the required parameters

   If the Encap Offset Mode parameter is set to Auto, you must configure the Encap Offset parameter.
3. Save your changes and close the form.

15

Click on the To Client Options tab to configure relay to client options for the PPP host.

1. Click Create or select an existing options instance and click Properties. The To Client Options (Create|Edit) form opens.
2. Configure the required parameters.
3. Configure any additional parameters that appear in the Option Value panel. The configurable parameters vary, depending on the Type parameter value.
4. Save your changes and close the form.

16

Save your changes and close the forms.

END OF STEPS

74.12 To enable or disable subnet draining on a local DHCPv4 server

74.12.1 Purpose

Perform this procedure to configure IP address draining for a subnet on a local DHCPv4 server. When subnet draining is enabled, the server stops leasing IP addresses from the subnet to new subscriber hosts; existing leases are not affected.

74.12.2 Steps

1

Perform one of the following.

a. Configure subnet draining on a base routing instance:

1. Choose Routing from the navigation tree view selector.
2. Expand *Network→NE→Routing Instance*.
3. Right-click on the routing instance and choose Properties. The Routing Instance (Edit) form opens.

b. Configure subnet draining on a VPRN routing instance:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.
3. In the navigation tree, expand the Sites object.
4. Select a routing instance object. The Site (Edit) form opens.

2

Click on the RADIUS/DHCP/Diameter tab. The Local DHCP Servers tab is displayed.

3

Select the required DHCPv4 or DHCPv6 server and click Properties. The Local DHCP Server (Edit) form opens.

4

Click on the IP Address Pools tab.

5

Select a pool and click Properties. The IP Address Pool (Edit) form opens.

6

Click on the Subnets tab.

7

Select a subnet and click Properties. The Subnet (Edit) form opens.

8

Enable or disable the Subnet Drain parameter and click Apply.

If the Subnet Drain parameter is enabled, the DHCP server stops leasing IP addresses from the subnet.

If the Subnet Drain parameter is disabled, the DHCP server leases IP addresses from the subnet.

9

Close the forms.

END OF STEPS

74.13 To configure NE SHCV event handling

74.13.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Click View SHCV. The Subscriber Host Connectivity Verification form opens.

3

Choose Network Element SHCV and click Search.

4

Select an NE and click Properties. The Network Element SHCV (Edit) form opens.

5

Configure the required parameters.

6

Save your changes and close the forms.

END OF STEPS

74.14 To configure a MEP on an SDP Binding

74.14.1 Purpose

You can configure a MEP on an SDP binding for VPLS, VPRN, and VLL Epipe services.

74.14.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Properties. The service configuration form opens.

3

On the navigation tree, expand Sites→*Routing Instance*→Spoke SDP Bindings→*SDP*.

4

Select one or more Spoke SDP Bindings, right-click, and choose Properties. The Spoke SDP Binding configuration form opens.

5

Click on the OAM tab.

6

On the ETH-CFM tab, MEPs panel, click Create or select an existing MEP entry and click Properties. The MEP (Create|Edit) form opens.

7

Select a maintenance entity group.

8

Configure the required parameters.

The CCM Padding Packet Size parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

9

If the MD for the MEP has a Name Type of none and the associated MEG has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable. Click on the Y.1731 Tests tab and configure the required parameters.

The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

10

Click on the AIS tab and configure the required parameters.

The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.

11

Save your changes and close the forms.

END OF STEPS

74.15 To configure L2Aware static port forwarding on a subscriber instance

74.15.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search.

3

Select a subscriber instance and click Properties. The Residential Subscriber Instance (Edit) form opens.

4

Click on the NAT Static Port Forwarding tab.

1. Click Create or select an existing static port forwarding entry and click Properties. The NAT Static Port Forwarding Display (Create|Edit) form opens.
2. Configure the required parameters.

   You cannot specify the same set of Inside Port and Protocol values in more than one static port mapping to an Inside IP Address.

   You can specify the same Outside Port value in multiple mappings to an Inside IP Address.
3. Save your changes and close the form.

5

Perform the following steps to synchronize NAT static port forwards.

1. Select one or more static port forward entries and click Sync. The NAT Static Port Forwarding Sync form opens. You can select up to 100 static port forward entries.
2. Select a targeted router instance.
3. Configure the Lifetime parameter.

6

Save your changes and close the forms.

END OF STEPS

74.16 To resynchronize static port forwarding entries

74.16.1 Purpose

Because of the potentially large number of static port forwarding entries, the entries are not resynchronized during a full NE resynchronization. Perform this procedure to manually resynchronize the entries.

Note: The NFM-P static port forwarding entries are current and do not require resynchronization if the following conditions are true.

- The NFM-P is used to create, maintain, and delete each entry.
- Each entry has an infinite lifetime, so is truly static.

74.16.2 Steps

1

Right-click on an NE in the Equipment navigation tree and choose Resync→Customized Resync. The Resync Sites wizard opens.

2

Select Choose MIB Entries.

3

Click Next. The Choose MIB Entries step is displayed.

4

Select tmnxNatFwd2Entry.

5

Click Next. The Force Resync step is displayed.

6

Select the Ignore Timestamps parameter.

7

Click Finish. The entries are refreshed.

You can view the static port forwarding entries on the NAT Static Port Forwarding tab of the NAT Configuration form for the following:

- base or service routing instances
- subscriber instances

END OF STEPS

74.17 To change the identification of a subscriber

74.17.1 Purpose

Perform this procedure to change the identification of a subscriber, for example, when you no longer want hosts on the SAP to use the SAP ID as the subscriber ID.

CAUTION

Service Disruption

Renaming a subscriber can be service-affecting because it changes the subscriber identification string of all associated subscriber hosts. Before you proceed, ensure that no subscriber hosts require the subscriber identification string associated with the subscriber.

Renaming a subscriber changes the SAP default subscriber identification strings that are associated with it.

74.17.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search

3

Select a residential subscriber and click Modify. The Residential Subscriber (Edit) form opens.

4

Configure the New Subscriber Identification parameter.

5

Save your changes and close the form.

END OF STEPS

74.18 To reset a subscriber’s subscription credit limit

74.18.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form appears.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search

3

Perform one of the following:

a. Select a host from the list and click Reset Credit. Go to Step 7.

b. Select a host from the list and click Properties. The Subscriber Host form opens.

4

Click on the Credit Control Operational values tab.

5

Select an entry and perform one of the following:

a. Click Reset Credit. Go to Step 7.

b. Click Properties. A form opens.

6

Review the credit control data and close the form.

7

Save your changes and close the forms.

END OF STEPS

74.19 To modify the primary subscriber identification script or URL

74.19.1 Purpose

This procedure allows an NFM-P operator to modify the primary subscriber identification script without service disruption. Perform this procedure only when the primary script and URL are operational.

CAUTION

Service Disruption

Modifying the primary subscriber identification script is potentially service-affecting if no functional backup script is administratively enabled. Modifying a backup (secondary or tertiary) subscriber identification script is unlikely to be service-affecting if the other backup script (secondary or tertiary) functions properly and is administratively enabled. Ensure that at least one administratively enabled backup script is accessible to the NFM-P and the NEs to which it applies before you proceed.

74.19.2 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Manage Subscriber Policies form opens.

2

Select a subscriber identification policy and click Properties. The Subscriber Identification Policy (Edit) form opens.

3

Copy the operational primary script URL and paste it to the secondary position. This action ensures that an operational backup script is in place in the event that there is a problem with the new primary script or URL.

1. Set the Secondary Script Administrative State parameter to Disabled.
2. Click Apply and confirm the operation. The NFM-P administratively disables the secondary script.
3. Configure the Secondary Script URL parameter with the value of the Primary Script URL parameter.
4. Set the Secondary Script Administrative State parameter to Enabled.
5. Click Apply and confirm the operation. The NFM-P administratively enables the secondary script.

   You must administratively disable and enable a script URL to cause the NEs to which the subscriber identification policy applies to load the script using the URL.

4

Modify a renamed copy of the former primary script or create a replacement script, as required. Record the new or modified script URL.

5

Configure the new URL as the primary script URL.

1. Set the Primary Script Administrative State parameter to Disabled.
2. Click Apply and confirm the operation. The NFM-P administratively disables the primary script. The secondary (former primary) script is the active script.
3. Configure the Primary Script URL parameter with the new URL value.
4. Set the Primary Script Administrative State parameter to Enabled.
5. Click Apply and confirm the operation. The NFM-P administratively enables the primary script. The new primary script is the active script.

   You must administratively disable and enable a script URL to cause the NEs to which the subscriber identification policy applies to load the script using the URL.

6

Save your changes and close the forms.

END OF STEPS

74.20 To delete an inactive subscriber instance

74.20.1 Purpose

Use this procedure to remove the record of an inactive subscriber instance from the NFM-P database. Subscriber instances become inactive in the NFM-P when the subscriber is deleted from the NE.

Note: AA statistics data can be lost when you delete a subscriber instance before the NFM-P is finished collecting and processing AA statistics. Nokia recommends waiting at least two statistics collection intervals before deleting an inactive subscriber instance.

74.20.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search.

The Active property indicates the status of each subscriber instance in the list. An inactive subscriber instance has no check mark under the Active heading.

3

Select the inactive subscriber instance you want to delete and click Delete. The inactive subscriber instance is removed from the list.

4

Close the form.

END OF STEPS

74.21 To view WLAN GW UEs on an NE

74.21.1 Purpose

Use this procedure to query different types of WLAN GW UEs configured on an NE ISA.

74.21.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose WLAN GW Query-based User Equipment and click Search.

3

Select the site ID of the NE you want to query.

4

Specify a WLAN GW ISA group, WLAN GW ISA group member, or UE MAC address to narrow your search, and then click Search.

5

Select a UE item in the list and click Properties. The WLAN GW DSU UE properties form opens.

6

View the read-only information as required and then close the form.

END OF STEPS

74.22 To configure residential subscriber management components on a SAP

74.22.1 Purpose

Perform this procedure to manage the subscriber configuration on one or more SAPs. You can configure the subscriber-related SAP parameters and assign one or more of the following residential subscriber management components:

- subscriber identification policy
- default subscriber identification string
- default subscriber profile
- default SLA profile
- non-subscriber traffic subscriber profile
- non-subscriber traffic SLA profile

74.22.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Properties. The service configuration form opens.

3

On the navigation tree, expand *site*→*routing instance*.

4

Click on the Access Interfaces icon to display the site SAPs.

5

Select one or more SAPs, right-click, and choose Properties from the menu. The SAP configuration form opens.

Note: You can configure subscriber management on multiple SAPs at once from any list of SAP search results that the NFM-P generates. For simplicity, the procedure uses the term SAP to refer to one or more SAPs, and focuses only on SAPs that belong to the same service site.
The NFM-P commits changes to a SAP configuration only when subscriber management is enabled on all SAPs that you are configuring.

6

Click on the Subscriber Management tab. The Host Connectivity tab is displayed.

7

Click on the Profiles tab.

8

Perform one of the following steps for a component.

a. Associate a component with the SAP by selecting a component in the Policies panel.

b. Remove a component, by clicking Clear beside the component in the Policies panel.

Note: Removing or replacing a residential subscriber management component on a SAP that is in use by subscriber hosts can be service-affecting to hosts that attempt to join the network. Ensure that removing the component does not affect the subscriber hosts before you proceed.

9

Save your changes and close the form.

END OF STEPS

74.23 To enable or disable residential subscriber management on a SAP

74.23.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Properties. The service configuration form opens.

3

On the navigation tree, expand *site*→*routing instance*.

4

Click on the Access Interfaces icon to display the site SAPs.

5

Select one or more SAPs, right-click, and choose Properties from the menu. The SAP configuration form opens.

6

Click on the Subscriber Management tab.

7

Click on the Profiles tab.

8

Configure the Admin Status parameter by performing one of the following.

a. Set the parameter to Enabled to enable residential subscriber management on the SAP.

b. Set the parameter to Disabled to disable residential subscriber management on the SAP.

9

Save your changes and close the form.

END OF STEPS

74.24 To create a static host for residential subscriber management on a SAP

Note: If you are configuring an IPv6 static host, ensure that the IPv6 Allowed parameter is enabled on the parent group interface and subscriber interface for the SAP.

74.24.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click on the Properties button. The service configuration form opens.

3

On the navigation tree, expand the site on which you want to create a static host for residential subscriber management on a SAP.

4

Navigate to the SAPs that support subscriber management. Perform one of the following, depending on the type of service that you are configuring:

a. For an IES, expand *IES*→*site*→Subscriber Interfaces→*subscriber_interface*→*group_interface*→Service Access Points.

b. For a VPLS, expand *VPLS_service*→*site*→Access Interfaces→*SAP*.

c. For a VPRN service, expand *VPRN_service*→*site*→Subscriber Interfaces→*subscriber_interface*→*group_interface*→Service Access Points→*SAP*.

5

Select one or more SAPs, right-click, and choose Properties from the menu. The SAP configuration form opens.

6

If you are configuring an IPv6 static host, click on the Subscriber Management tab > Profiles tab, and ensure that the Admin Status parameter is set to Enabled.

7

Click on the Anti-Spoofing tab and configure the Anti-Spoofing parameter.

You must set the Anti-Spoofing parameter to IP-address matching or to IP- and MAC-address matching before you can enable subscriber management for the static hosts on a SAP.

8

To configure an IPv4 static host, click on the Static Hosts tab.

1. Click Create or select an existing static host and click Properties. The Access Interface Anti-Spoofing Static Host (Create|Edit) form opens.
2. Configure the required parameters.

To enable residential subscriber management for a static host, you must specify values for the IP Address and Subscriber Identification parameters.

3. Select a subscriber profile for the static host.
4. Select an SLA profile for the static host.
5. Select an application profile for the static host, if required.
6. Configure the Scope parameter, if applicable.
7. Select a RIP policy for the static host, if required.
8. Set the Administrative State parameter to Up and confirm the operation.
9. Save your changes and close the form.

9

To configure an IPv6 static host, click on the IPv6 Static Hosts tab.

1. Click Create or select an existing IPv6 static host and click Properties. The Access Interface Anti-Spoofing IPv6 Static Host (Create|Edit) form opens.
2. Configure the required parameters.

   To link the IPv6 static host to an existing IPv4 static host, configure the MAC Linking parameter with the IP address of the IPv4 static host.

   To enable residential subscriber management for a static host, you must specify values for the IP Address and Subscriber Identification parameters.
3. Select a subscriber profile for the IPv6 static host.
4. Select an SLA profile for the IPv6 static host.
5. Select an application profile for the IPv6 static host, if required.
6. Configure the Scope parameter, if applicable.
7. Set the Administrative State parameter to Up and confirm the operation.
8. Save your changes and close the form.

10

Click on the Managed Routes tab.

1. Click on the Create button. The Access Interface Anti-Spoofing Static Host Managed Route Display (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the forms.

11

Save your changes and close the forms.

END OF STEPS

74.25 To configure a MEP on a SAP

74.25.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Properties. The service configuration form opens.

3

On the navigation tree, expand the Sites icon.

4

Perform one of the following steps.

a. To create a MEP on an Access Interface SAP, expand *site_n*→Access Interfaces. The site SAPs are displayed.

   MEPs can be configured on an Access Interface SAP for VPLS, VPRN, IES, and VLL Epipe services.

b. To create a Down MEP on a Subscriber Group Interface SAP, expand *site_n*→Group Interfaces. The site group interface SAPs are displayed.

   Down MEPs can be configured on a Subscriber Group Interface SAP for VPRN and IES.

5

Select one or more SAPs and click Properties. The SAP configuration form opens.

6

Click on the MEP tab.

1. Click Create or select an existing MEP entry and click Properties. The MEP (Create|Edit) form opens.
2. Configure the required parameters.

   The Control MEP parameter is only displayed if the Interface Type parameter is set to Ethernet Tunnel Path Endpoint.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

   When configuring a Down MEP on a Subscriber Group Interface SAP, the Direction parameter cannot be configured.
3. If the MD for the MEP has a Name Type of none and the associated MEG has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable. Configure the required parameters on the tabs.

The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.

7

Save your changes and close the forms.

END OF STEPS

74.26 To configure a capture SAP

74.26.1 Purpose

Perform this procedure to create a capture SAP to trigger MSAP creation.

Created MSAPs are on the same port as the associated capture SAP. A capture SAP must be the default SAP on a port.

74.26.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPLS in the list and click Properties. The VPLS Service (Edit) form opens.

3

In the navigation tree, expand Sites→*site_n*→Access Interfaces.

4

Right-click on the Access Interfaces icon and choose Create VPLS L2 Access Interface, or right-click an existing access interface object and choose Properties. The VPLS L2 Access Interface (Create|Edit) form opens.

1. Set the SAP Sub Type parameter to Capture. The form displays a different array of tabs and panels.
2. Configure the Description and Administrative State parameters.
3. Select a host lockout policy, if required.
4. Select a subscriber authentication policy, if required.

5

Click on the Port tab.

1. Select a terminating port for the capture SAP.

2. Configure the encapsulation values for the SAP, if required. The encapsulation type of the terminating port determines the encapsulation values that are available. The table below describes the encapsulation values for each encapsulation type.

   The capture SAP is used if a more specific match for the dot1q or QinQ tags is not found.

   Some providers use QinQ encapsulation to represent a service.sub or sub.service SAP. When configured, the full QinQ represents a subscriber.

Table 74-4 Encapsulation types

Encapsulation type	Encapsulation values
dot1q	4095 only, which appears as an asterisk (*)
QinQ	1 to 4095 for the outer encapsulation value, and 4095 only for the inner encapsulation value
LAG	1 to 4095 for both encapsulation values

6

Click on the ACL tab to select an ACL filter.

7

Click on the QoS tab to specify an ingress or egress queue group policy.

1. Select an NE DoS protection policy.
2. Configure the MAC Monitoring parameter OR the Ethernet CFM Monitor Flags parameter.
3. Select a DDoS protection policy.

8

If the capture SAP is an ATM SAP, click on the ATM tab.

1. Configure the AAL5 Encapsulation parameter.
2. Click on the VC Ranges tab.
3. Click Create or select an existing ATM configuration and click Properties. The ATM Configuration (Create|Edit) form opens.
4. Configure the required parameters.
5. Save your changes and close the form.

9

Click on the Capture Access Interface tab.

1. Enable the Allow Dot1Q MSAPs parameter, if required. This parameter can only be configured when the terminating port for the capture SAP is configured for QinQ encapsulation.
2. Select one or more of the following parameters on the Trigger Packet panel to specify the types of packets that will trigger RADIUS authentication:
   - DHCPv4
   - PPP
   - PPPoE
   - ARP
   - DHCPv6
   - Data

   The PPPoE parameter is selectable only on an ATM SAP. The ARP, DHCPv4, and DHCPv6 parameters are selectable only on a non-ATM SAP.
3. Select a MSAP policy in the MSAP Defaults panel.
4. Select or manually type an MSAP Service ID.
5. Select or manually type an MSAP Group Interface Name.
6. Select a PPP policy in the PPP panel, if required.

   A PPP policy is configurable only on an ATM SAP.
7. Select a PPP policy in the PPPoE panel, if required.
8. Select a local user database for the following protocols, as required:
   - PPP
   - PPPoE
   - DHCPv4
   - DHCPv6

   A PPP local user database is selectable on only an ATM SAP.

   A DHCPv4 or DHCPv6 local user database is selectable on only a non-ATM SAP.
9. Select a python policy for the following protocols, as required:
   - PPPoE
   - DHCPv4
   - DHCPv6
10. Select an SRRP instance in the Redundancy panel.
11. Select a diameter authentication policy.
12. If the capture SAP is to be used for data-triggered authentication in a dynamic services configuration, select a dynamic service policy and configure its administrative state.

10

If the capture SAP is to be used to instantiate IPoE sessions, click on the IPoE Session tab.

1. Configure the Administrative State and Description parameters.

The Administrative State parameter must be set to Disabled before IPoE session policy and local user database configurations can be changed.

2. Select an IPoE session policy.
3. Select a local user database.

11

Save your changes and close the forms.

END OF STEPS

74.27 To list MSAPs and view MSAP properties

74.27.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Perform one of the following to display the MSAPs associated with each service:

a. For VPLS MSAPs:
   1. Click on the object type drop-down menu and expand Access Interface (Service Management)→L2 Access Interface (Service Management)→Abstract L2 Access Interface (VPLS)→L2 Access Interface (VPLS)→VPLS L2 Access Interface.
   2. Select VPLS MSAP and click Search.

b. For IES MSAPs:
   1. Click on the object type drop-down menu and expand Access Interface (Service Management)→Service Access Point (Service Management)→IES Service Access Point (IES).
   2. Select IES MSAP and click Search.

c. For VPRN MSAPs:
   1. Click on the object type drop-down menu and expand Access Interface (Service Management)→Service Access Point icon (Service Management)→VPRN Service Access Point (VPRN).
   2. Select VPRN MSAP and click Search.

3

Select an MSAP and click on the MSAP Properties tab.

END OF STEPS

74.28 To view an MSAP event log, modify the global MSAP log policy, and purge MSAP log records

74.28.1 Steps

1

Perform 74.27 “To list MSAPs and view MSAP properties” (p. 2052) to list the MSAPs and choose the appropriate MSAP.

2

Click on the Events tab to display the MSAP event log.

3

Perform any of the following:

a. View an event log record:
   1. Select a record in the event log and click Properties. The Statistics Record - MSAP Event Log form opens, displaying the MSAP record properties.
   2. Close the form.

b. Modify the global MSAP event log policy:
   1. Click Log Policy. The Log Policy - Ressubscr. MSap Event Log form opens.

      The global MSAP log policy affects all MSAP event logs.
   2. Configure the Retention Time and Administrative State parameters.
   3. Save your changes and close the form.

c. Purge the MSAP event log records:
   1. Click Purge Log Records.
   2. Confirm the purge and close the form.

4

Close the form.

END OF STEPS

74.29 To modify and re-evaluate an MSAP policy on an MSAP

74.29.1 Purpose

Perform this procedure to re-apply an MSAP policy on an existing MSAP. MSAPs cannot be edited; however, MSAP policies on existing MSAPs can be changed and re-applied.

74.29.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form appears.

2

Configure the filter criteria; see 74.27 “To list MSAPs and view MSAP properties” (p. 2052) for information about configuring the filter to list MSAPs. A list of MSAPs for the service appears.

3

Select an MSAP and click Properties. The service configuration form opens.

4

Click on the MSAP Properties tab.

5

Click Properties in the Creation MSAP Policy panel and enter the required changes.

6

Click Apply and confirm the changes.

7

Click on the MSAP Properties tab.

8

Specify the Do Action option for the Creation MSAP Policy Re-evaluation parameter to apply the new policy parameter values.

If you specify Not Applicable, the new policy parameter values are not applied, but remain on the policy until you specify the Do Action option.

9

Save your changes and close the form.

END OF STEPS

74.30 To modify an MSAP policy and re-evaluate the MSAPs

74.30.1 Purpose

Perform this procedure to re-apply an MSAP policy on existing MSAPs. MSAPs cannot be edited; however, MSAP policies on existing MSAPs can be changed and re-applied.

74.30.2 Steps

1

Choose Policies→Residential Subscriber from the NFM-P main menu. The Residential Subscriber Policies form appears.

2

Select MSAP Policy from the object type drop-down menu and click Search. A list of MSAP policies appears.

3

Select an MSAP policy and click Properties. The MSAP Policy (Edit) form appears.

4

Enter the required changes.

5

Click Apply and confirm the changes.

6

Click Reevaluate MSAPs. The MSAPs are re-evaluated with the changed MSAP policy.

7

Save your changes and close the form.

END OF STEPS

74.31 To re-evaluate lease states for an MSAP

74.31.1 Purpose

Perform this procedure to push SLA profile or subscriber profile changes out to an MSAP. MSAPs cannot be edited; however, SLA profiles or subscriber profiles on existing default SAPs can be changed and re-applied. The Re-evaluate Lease States command updates all residential subscriber hosts associated with the MSAP without reinstantiating the hosts.

This procedure describes how to perform a Re-evaluate Lease States function from a residential subscriber GUI context. The same function can be performed from an IES or VPRN group interface GUI context.

74.31.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose IES Service Access Point or VPRN Service Access Point from the object menu and click Search.

3

Select the site ID of the IES or VPRN site on which you want to perform a Re-evaluate Lease States function and click Search.

A list of applicable SAPs appears.

4

Select a SAP and click Re-evaluate Lease States.

An information window appears, indicating the number of hosts that were sent a Re-evaluate Lease Sates request, and the number of hosts that were skipped.

5

Close the forms.

END OF STEPS

74.32 To configure DHCP event monitoring for a SAP

74.32.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Select one of the available SAP types and click Search.

3

Select a SAP entry and click Create Monitored SAP. The Monitored Access Interface (Create) form opens.

You can configure a maximum of five hosts or five SAPs for simultaneous DHCP event monitoring.

4

Configure the Monitoring Period and Units parameters.

5

Save your changes and close the forms.

Note: You must complete 74.33 “To monitor DHCP events for a SAP” (p. 2056) to enable and schedule SAP monitoring.

END OF STEPS

74.33 To monitor DHCP events for a SAP

74.33.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Monitored Access Interface (Monitor) and click Search.

3

To initiate or view DHCP monitoring for a particular SAP, select the entry and click Properties. The Monitored Access Interface (Edit) form opens.

4

You can perform one or more of the following on the form:

a. To start monitoring the chosen SAP, click Start Monitor.

b. To stop monitoring the chosen SAP, click Stop Monitor.

c. To remove monitoring of the chosen SAP, click Delete.

d. To change the monitoring period, click Stop Monitor, type the desired value in the Monitoring Period field, and select an appropriate unit from the drop-down menu.

5

To view DHCP events for the chosen SAP, click on the Events tab. The DHCP Events list is displayed. The NFM-P supports monitoring of the following DHCP events:

- DHCP Lease Entries Exceeded
- DHCP Lease State Overriden
- Suspicious DHCP Packet
- DHCP ACK Dropped
- Host Connectivity Lost
- Host Connectivity Restored

6

Select an event and click Properties to view detailed information about the event.

7

Save your changes and close the forms.

END OF STEPS

74.34 To view and configure residential subscriber hosts on a SAP

74.34.1 Purpose

Use this procedure to retrieve the residential subscriber hosts configured on a single SAP. You can view the configuration of individual subscriber hosts for troubleshooting purposes, and make configuration changes where required.

Note: The NFM-P retrieves active host data directly from NEs and displays results only after all of the hosts have been retrieved. Because this process can take minutes, Nokia recommends that you specify a single SAP address when searching subscriber hosts. Although it is possible to search subscriber hosts for an entire site or service ID, the time required could cause the NFM-P GUI to timeout before the search process completes, with no search results displayed. Searching without specifying a SAP should be reserved for LNS or WiFi Offload subscriber hosts.

74.34.2 Steps

1

Select Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Select Residential Subscriber Host in the drop-down list.

3

Select the site ID, service ID, and SAP from which you want to retrieve subscriber hosts.

4

Click Search to retrieve the subscriber hosts.

5

To view the information for a subscriber host, select the host entry and click Properties. The Subscriber Host form opens with the subscriber host information displayed.

If you need to make configuration changes to the subscriber host, proceed to the remaining steps in this procedure. Otherwise, close the forms.

6

CAUTION

Service Disruption

Changing the subscriber profile string directly on a subscriber host is potentially service-affecting to all hosts on all SAPs that belong to the subscriber. Because a subscriber is associated with only one subscriber profile, when you change the subscriber profile string for one host, all dynamic or PPPoE hosts associated with the subscriber on all SAPs are automatically configured with the new subscriber profile.

Configure the Subscriber Profile String parameter for the dynamic host or PPPoE host, if required. Select a subscriber profile string, or simply type a subscriber profile string in the field.

7

Configure the SLA Profile String parameter for the dynamic host or PPPoE host, if required. Select an SLA profile, or simply type an SLA profile string in the field.

8

Configure the ANCP String and Intermediate Destination ID parameters, if required.

9

Configure the parameters on the Subscriber Profile/SLA Profile Aliases panel, as required.

10

Configure the Application Profile parameter for the dynamic host PPPoE host, if required. Select an application profile string, or simply type an application profile string in the field.

11

Save your changes and close the forms.

END OF STEPS

74.35 To clear IPoE sessions from an IES or VPRN SAP or MSAP

74.35.1 Purpose

Perform this procedure to clear IPoE sessions from one or more selected SAPs or MSAPs on a specified IES or VPRN site.

You can also clear individual IPoE sessions from a VPRN or IES service site configuration form (*[Subscriber_interface]*→*[Group_interface]*→Subscriber Management tab→*[SAP]*).

74.35.2 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose one of the following from the object menu.

- VPRN Service Access Point
- VPRN Service Access Point→VPRN MSAP
- IES Service Access Point
- IES Service Access Point→IES MSAP
- IPoE Session

3

Specify the following objects to narrow the search results.

- IES or VPRN site ID
- Service ID (for individual IPoE session search only)
- Service Access Point (for individual IPoE session search only)

4

Click Search.

5

Select a SAP, MSAP, or IPoE session and click Clear IPoE Session(s).

6

Close the form.

END OF STEPS

75 VLAN service management

75.1 Overview

75.1.1 Purpose

VLANs are created to provide the segmentation services traditionally provided by routers in LAN configurations. VLANs address issues such as scalability, security, and network management. Routers in VLAN topologies provide broadcast filtering, security, address summarization, and traffic flow management. The following table lists the types of VLAN services that are supported on the NFM-P by device type:

Table 75-1 Device VLAN support

Device	Supported VLAN types
7450 ESS	• Standard VLAN • L2 VPN (TLS/VLAN-Stacking) VLAN • Broadcast TV (MVR/IPMV) VLAN
All OmniSwitches	• Standard VLAN • L2 VPN (TLS/VLAN-Stacking) VLAN
All OmniSwitches except for the OS 6900 and OS 10K	• Broadcast TV (MVR/IPMV) VLAN
OS 6900 and OS 10K	• IPC VLAN • VIP VLAN
Wavence SM [1]	• Wavence (dot1ad) VLAN • Wavence P2P (dot1q) VLAN • Wavence P2MP (dot1q) VLAN

Notes:

1. See “Wavence microwave backhaul service management” in the *NSP Wavence Device Support Guide* for more information.

Several 7450 ESS, Wavence and OmniSwitch VLAN services can be interconnected through a backbone VPLS.

The NFM-P supports end-to-end VLAN configuration using tabbed configuration forms with an embedded navigation tree.

75.1.2 Contents

VLAN service management overview

75.2 VLAN service management overview

75.2.1 VLAN ring groups

VLAN ring groups are used to send traffic across an Ethernet ring using copper or fiber optic connections from the source traffic device, for example, from a 7450 ESS, to all devices in the ring. STP configuration on OmniSwitch devices ensures that there is a constant stream of traffic in either direction by rerouting traffic around breaks in the physical links between the devices. The following figure shows an example of a ring group that is part of a larger metropolitan area network.

Figure 75-1 Ring in a metropolitan and broadcast network

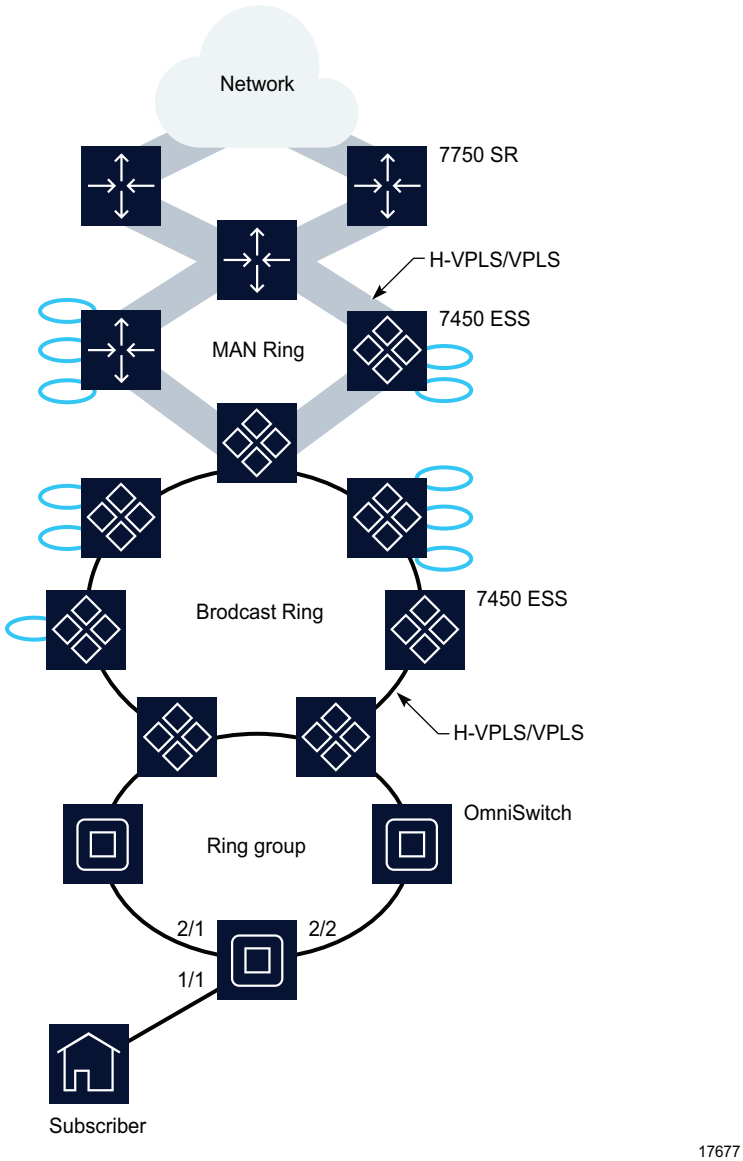

The NFM-P provides OAM tools for service validation and for troubleshooting service and network transport issues.You can run the OAM Validation test suite for the VLAN service by clicking on the Validate button or by clicking on the More Actions button and choosing Validate. Alternatively, you can also perform a One Time Validation. If a check mark appears beside the OAM Validation Failed state cause indicator, the test has failed. In addition, the Tested Entity Result tab on the Tests tab

displays more detailed information about the OAM test result. See Chapter 89, “Service Test Manager” for more information about how to configure OAM validation test suites.

The Aggregated Service Site Operational State parameter has four possible values. The value is derived from the operational states of the sites that are part of the service, as follows:

- Up—all sites are operationally up
- Partially Down—at least one site is operationally down
- Down—all sites are operationally down
- Unknown—the service has no provisioned sites

When the Aggregated Service Site Operational State parameter is Partially Down or Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P operator. You can view alarms on the Faults page.

When you use the NFM-P to create or discover a service, the NFM-P assigns a default tier value to the service. The Service Tier parameter value is relevant only in the context of composite service topology map views. See Chapter 85, “Composite service management” for more information about the hierarchical organization of composite services.

75.2.2 VLAN groups

The NFM-P supports the configuration and provisioning of VLAN groups which are used by the following NEs.

- OmniSwitch devices to configure VLAN services
- Wavence SM ANSI/ETSI (2.x and 3.x streams) to configure VLAN and VLL services

You can use VLAN groups to:

- logically group OmniSwitch and Wavence SM devices to represent a typical network topology. An OmniSwitch and a Wavence SM cannot belong to the same VLAN group.
- manage the VLAN IDs that are assigned to Wavence SM VLAN group members.

See 75.13 “To create a VLAN group” (p. 2083) for more information about creating OmniSwitch VLAN groups. See “Wavence microwave backhaul service management” in the *NSP Wavence Device Support Guide* for information about creating Wavence VLAN groups.

75.2.3 OmniSwitch VLAN service policies

Policies can be assigned to ports on an OmniSwitch. Policies are defined at a global level and then applied to components of the service, such as a port.

The policy on the component is then a local version of the global policy. The following policies are common to VLAN services:

- QoS policies define ingress classification, policing, shaping, and marking on the device.
- UNI policies define how control frames that are received on a port are processed. UNI policies are applied to a port that is used as a SAP in a L2 VPN TLS (stacked) VLAN.

- SAP policies define traffic engineering parameters for bandwidth sharing, rate limiting, CVLAN translation (or double-tagging), and priority bit mapping. SAP policies are applied to the service access multi-point.
- Filter policies control network traffic into or out of an interface or device based on IP or MAC matching criteria.

See Chapter 49, “Policies overview” for more information about policies.

75.2.4 OmniSwitch network management

The NFM-P manages OmniSwitch devices using SNMP and CLI messages. OmniSwitch devices support termination of all VLAN types on Ethernet ports.

75.2.5 Spanning tree protocols

The STP configuration on an OmniSwitch device in a VLAN group detects loops in the topology and ensures that there is a constant stream of traffic by rerouting traffic around breaks in the physical links between the devices. OmniSwitch devices support STP, MSTP and RSTP.

Sample VLAN configurations

75.3 Sample L2 VPN VLAN configuration

75.3.1 Overview

The following figure shows a sample L2 VPN VLAN service configuration. Transparent LAN services such as a L2 VPN are used to transport large numbers of customer VLANs while keeping the traffic in each VLAN segregated. The configuration depends on the specific network requirements.

Figure 75-2 Sample L2 VPN VLAN configuration

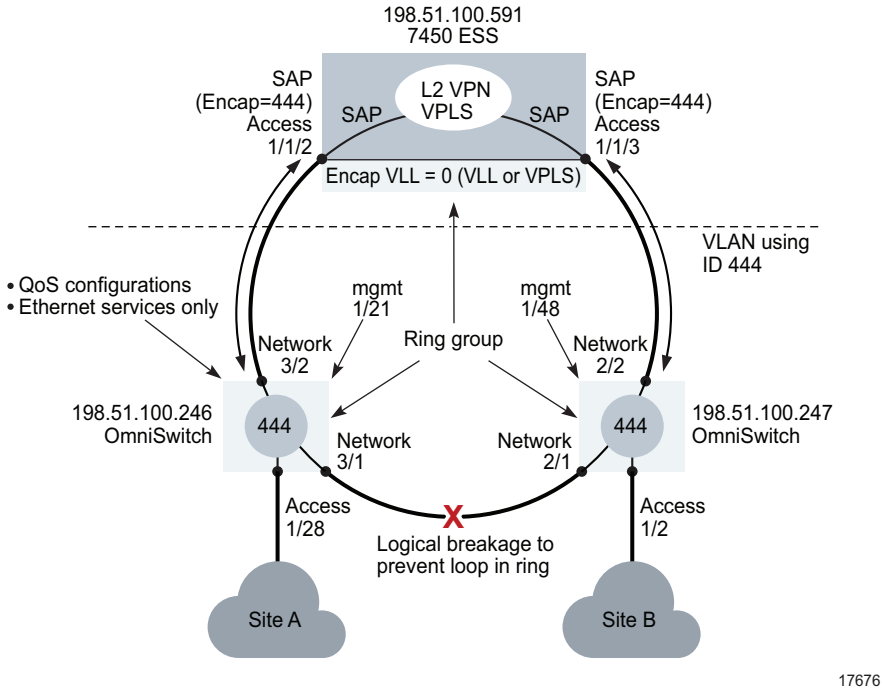

75.3.2 Pre-configuration tasks

Verify that the following preconfigurations are complete.

- Ensure that the appropriate preconfigurations have been performed on the OmniSwitch devices.
  - pre-discovery CLI modifications
  - discovery including mediation configuration with the CLI user names and passwords
  - configure protocol to manage topology loops, such as MSTP or RSTP

- Ensure that the VPLS that feeds the ring VLAN service is configured on the 7450 ESSs. The encapsulation of the SAPs that belong to the VPLS on the 7450 ESSs must match the VLAN ID of the ring VLAN.
- Configure the OmniSwitch ports as access (ports that are part of the VLAN) and network (ports that are used for uplinks), as required, from the navigation tree.
- Devices that belong to the ring, and the 7450 ESS that the ring connects to, must be added to the ring group.
- Ensure that the appropriate preconfigurations have been performed on the 7450 ESS.
  - A VLL is created with 0 encapsulation between the SAPs (1/1/2 and 1/1/3) on the 7450 ESS
  - The SAPs on the 7450 ESS that are part of the VPLS use the same encapsulation as the VLAN ID for the ring group
  - A L2 VPN VPLS is created on the 7450 ESS.

Note: Only Ethernet services are supported on an OmniSwitch L2 VPN VLAN service.

75.3.3 Configuration tasks

The following high-level tasks are required to configure this sample VLAN L2 VPN service:

1

Ensure that the required configurations are completed to equipment, including configuring access and network ports, enabling CLI configuration on managed OmniSwitch devices.

2

Configure the required policies, including QoS, scheduling and IGMP filter policies.

3

Distribute configured policies to the devices. The policies are used during the creation of VLAN services.

4

Create VLAN services using a series of configuration forms.

Ensure that the VLAN created for the ring group uses ID 444.

5

Associate access interfaces (similar to SAPs) with VLAN services, that are the physical ports to which end users connect. Use the L2 Interfaces tab on the VLAN properties form to associate VLAN services with the ports used by end users.

75.4 Sample BTV VLAN configuration

75.4.1 Overview

The following figure shows a sample broadcast TV VLAN configuration. BTV VLANs are shared VLANs, where the multicast broadcast channels or pay per view channels are available across the ring, and based on subscriptions and privileges determined using IGMP snooping, the user gains or is denied access. The configuration depends on the specific network requirements.

Figure 75-3 Sample BTV VLAN configuration

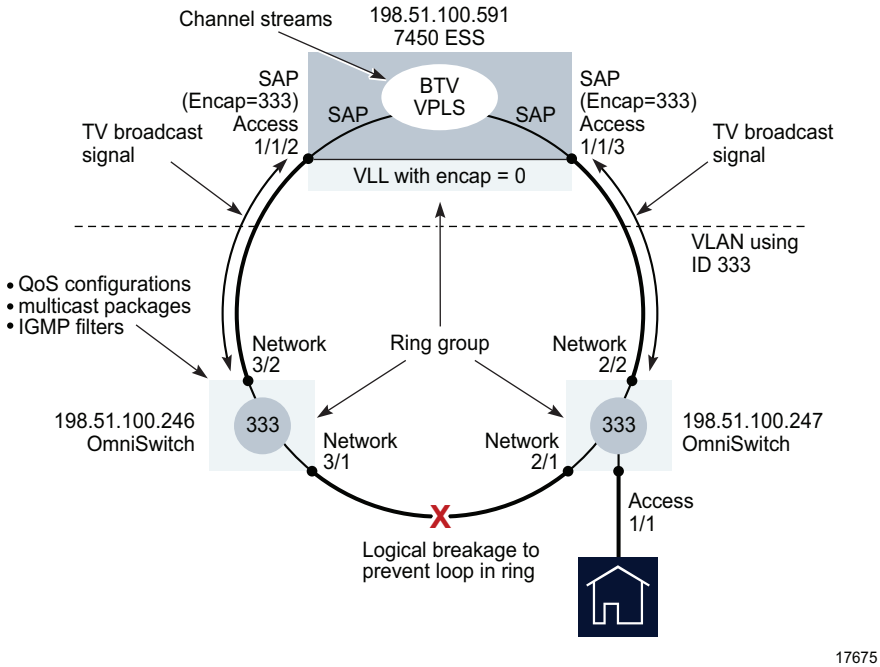

75.4.2 Pre-configuration tasks

Verify that the following preconfigurations are complete.

- Ensure that the appropriate preconfigurations have been performed on the OmniSwitch devices:
  - pre-discovery CLI modifications
  - SNMP trap forwarding to the NFM-P
  - discovery, including mediation configuration with the CLI user names and passwords
  - configure protocols to manage topology loops, such as MSTP or RSTP
- Ensure that the BTV VPLS that feeds the ring VLAN service is configured on the 7450 ESSs. The encapsulation of the SAPs that belong to the VPLS on the 7450 ESSs must match the VLAN ID of the ring VLAN.

- Configure OmniSwitch ports as access (ports that are part of the VLAN) and network (ports that are used for uplinks), as required, from the navigation tree.
- OmniSwitch devices that belong to the ring, and the 7450 ESS that the ring connects to, must be added to the ring group.
- Ensure that the appropriate preconfigurations have been performed on the 7450 ESS.
  - A VLL is created with 0 encapsulation between the SAPs (1/1/2 and 1/1/3) on the 7450 ESS
  - The SAPs on the 7450 ESS that are part of the VPLS use the same encapsulation as the VLAN ID for the ring group
  - A BTV VPLS is created on the 7450 ESS.
- Create and manage the necessary broadcast TV policies, including multicast package and IGMP filtering.

75.4.3 Configuration tasks

The following high-level tasks are required to configure this sample broadcast TV VLAN service.

1

Ensure that the required configurations are completed to equipment, including configuring access and network ports, enabling CLI configuration on managed OmniSwitch devices, and the creation of ring groups.

2

Configure the required policies, including QoS, scheduling and IGMP filter policies.

3

Distribute configured policies to the devices. Policies are used during the creation of VLAN services.

4

Create VLAN services using a series of configuration forms:

Ensure that the VLAN created for the ring group uses ID 333.

5

Associate access interfaces (similar to SAPs) with VLAN services, that are the physical ports to which end users connect. You use the L2 Interfaces tab on the VLAN properties form to associate VLAN services with the ports used by end users.

75.5 Sample interconnection VLAN configuration

75.5.1 Overview

The following figure shows a sample configuration of an interconnection of VLANs across a VPLS backbone.

Figure 75-4 Sample interconnection VLAN

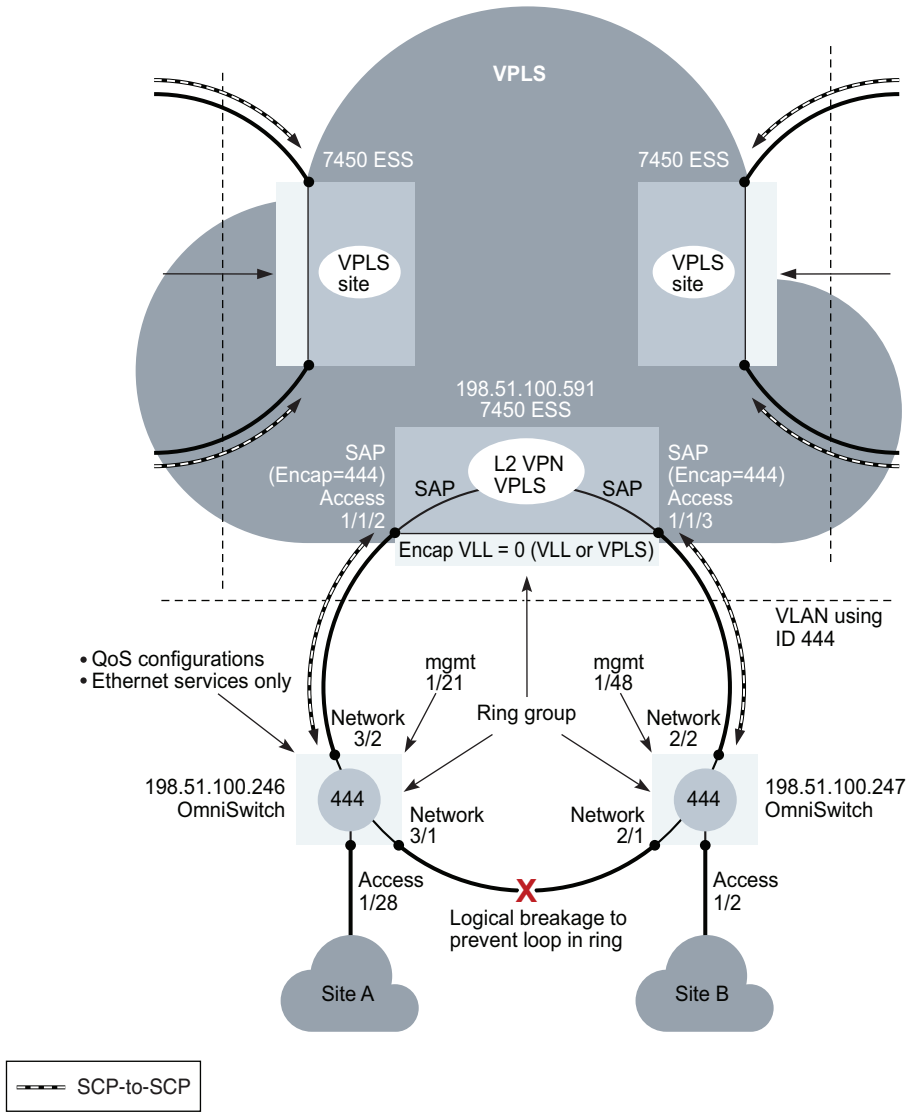

75.5.2 Pre-configuration tasks

Verify that the following preconfigurations are complete:

- Ensure that the appropriate preconfigurations have been performed on the OmniSwitch devices.
  - pre-discovery CLI modifications
  - SNMP trap forwarding to the NFM-P
  - discovery including mediation configuration with the CLI user names and passwords

  - ports are configured as access (ports that are part of the VLAN) and network (ports that are used for uplinks)
  - configure protocol to manage topology loops, such as MSTP or RSTP, using CLI
- Configure OmniSwitch ports as access and network, as required, from the navigation tree.
- OmniSwitch devices that belong to the ring, and the 7450 ESS that the ring connects to, must be added to the ring group.

75.5.3 Configuration tasks

The following high-level tasks are required to configure this sample interconnection of VLAN services.

1

Ensure that the required configurations are completed to equipment as described in 75.3.3 “Configuration tasks” (p. 2068) , including access and network ports, enabling CLI configuration on managed OmniSwitch devices, and configuring and distributing policies.

2

Create VLAN services using a series of configuration forms.

Ensure that the VLAN created for the ring group uses ID 444.

3

Configure the VPLS:

- Ensure that the encapsulation value of the ports that contain the SAPs on the 7450 ESS matches the VLAN ID of each VLAN service in the ring group.
- Ensure that the encapsulation value of the ports that contain the SAPs on the 7450 ESS at each VPLS site is Dot1 Q.

See 77.5 “To create a VPLS” (p. 2249) for more information about creating a VPLS.

4

Create a composite service to interconnect the VLAN services with the VPLS.

- Create the composite service.
- Define general properties for the composite service.
- Specify the services that are participating in the composite service. You can specify multiple services in one operation.
- Create SCP-to-SCP connectors to link the dot1q-encapsulated VPLS SAP and the adjacent L2 switch uplinks of the VLAN ring group.

See Chapter 85, “Composite service management” for more information about creating composite services.

VLAN service management procedures

75.6 Workflow to create VLAN services (OmniSwitch)

75.6.1 Overview

The following workflow lists the high-level steps required to create a VLAN service on OmniSwitch devices. As a prerequisite for creating a VLAN service, this workflow assumes that the following appropriate preconfigurations have been performed:

- pre-discovery CLI modifications
- discovery including mediation configuration with CLI usernames and passwords
- addition of the OmniSwitches that participate in the VLAN to a VLAN group
- configure OmniSwitch network ports, as required, from the navigation tree. The Automatic VLAN Binding parameter associated with a network port or LAG must be enabled before the NFM-P can identify the network port or LAG as a network interface. This parameter is only supported on Stacked VLANs.
- configure QoS policies, as required

75.6.2 Stages

1

Provision a Standard VLAN service. See 75.7 “To create a standard VLAN service on OmniSwitch devices” (p. 2075) .

1. Create the VLAN (setting the application parameter to Standard VLAN) and associate a customer with the VLAN.
2. Associate access interfaces with the VLAN service. See 75.11 “To associate an access interface with a VLAN service” (p. 2081) .

2

Provision a L2 VPN VLAN service. See 75.8 “To create an OmniSwitch L2 VPN TLS VLAN service” (p. 2076) .

1. Create the VLAN (setting the application parameter to L2-VPN (TLS/VLAN-Stacking)) and associate a customer with the VLAN.
2. Add Ethernet services to the VLAN site. See 75.12 “To add Ethernet services to a VLAN Site” (p. 2082) .

3

Provision a BTV VLAN service. See 75.9 “To create an OmniSwitch BTV VLAN service” (p. 2078) .

1. Create the VLAN (setting the application parameter to Broadcast TV (MVR/IPMV)) and associate a customer with the VLAN.

2. Associate access interfaces with the VLAN service. See 75.11 “To associate an access interface with a VLAN service” (p. 2081) .
3. Add Ethernet services to the VLAN site. See 75.12 “To add Ethernet services to a VLAN Site” (p. 2082) .

4

Provision a VIP VLAN service. See 75.10 “To create an OmniSwitch VIP VLAN service” (p. 2080) .

1. Create the VLAN (setting the application parameter to VIP VLAN) and associate a customer with the VLAN.
2. Associate access interfaces with the VLAN service. See 75.11 “To associate an access interface with a VLAN service” (p. 2081) .

5

Provision a Backbone VLAN service. See 80.3 “To create an OmniSwitch Backbone VLAN service” (p. 2695) .

6

As required:

a. Create VLAN groups and add members to each group as required. See 75.13 “To create a VLAN group” (p. 2083) .

b. Delete a VLAN group or group member. See 75.14 “To delete a VLAN group or group member” (p. 2084) .

c. Add a MEP to an OmniSwitch VLAN service access interface. See 75.15 “To manually add MEPs to an OmniSwitch VLAN service access interface” (p. 2085) .

d. Configure IGMP on an OmniSwitch VLAN site. See 75.16 “To configure IGMP on an OmniSwitch VLAN site” (p. 2086) .

7

Turn up the service.

Network interface VLAN bindings are created between the newly-created VLAN and all of the network ports on the node.

8

As required, run an OAM validation test on a VLAN service. See 75.19 “To run an OAM validation test on a VLAN service” (p. 2090) .

9

View the VLAN service operational status. See 75.20 “To view the VLAN service operational status” (p. 2091) .

75.7 To create a standard VLAN service on OmniSwitch devices

Note: Nokia recommends that you specify the OmniSwitch ports that will be network interfaces before you configure a standard VLAN service.

75.7.1 Steps

1

Choose Create→Service→VLAN from the NFM-P main menu. The VLAN (Create) Service form opens.

2

Select a customer to associate with the VLAN.

3

Configure the required general parameters.

The Service ID and SVC Mgr Service ID parameters are configurable when the Auto-Assign ID parameter is disabled.

4

Click on the VLAN tab and set the Application parameter to Standard VLAN.

5

Select a group to associate with the VLAN in the Group panel.

6

On the service navigation tree, right-click on VLAN and choose Create VLAN Site to choose a site for the VLAN. The Site (Create) form opens.

7

Configure the required general parameters.

8

Click on the VLAN tab and configure the required parameters.

9

Perform 75.11 “To associate an access interface with a VLAN service” (p. 2081) to create one or more access interfaces on the site.

10

To enable DHCP/DHCPv6 snooping on the VLAN:

1. Click on the DHCP/DHCPv6 Snooping tab.

2. Click Create. The VLAN Level DHCP/DHCPv6 Snooping (Create) form opens.
3. Configure the required parameters.

   These parameters are automatically enabled when DHCP/DHCPv6 snooping is enabled on a VLAN.
4. Save the changes and close the form.

11

To add one or more static entries to the DHCP binding table:

1. Click on the Binding Database tab.
2. Click Create. The DHCP Snooping Binding Database (Create) form opens.
3. Configure the MAC Address parameter.
4. Select a port.
5. Configure the required parameters.
6. Save the changes and close the form.

12

Save the changes and close the forms.

END OF STEPS

75.8 To create an OmniSwitch L2 VPN TLS VLAN service

Note: The TLS Mode must be set to Ethernet Service before a sL2 VPN TLS (stacked) VLAN service can be created on an OmniSwitch. See 28.128 “To configure bridging on an OmniSwitch” (p. 1043) to configure the TLS Mode.

Nokia recommends that you configure OmniSwitch network ports before you configure a L2 VPN TLS VLAN service.

75.8.1 Steps

1

Choose Create→Service→VLAN from the NFM-P main menu. The VLAN Service (Create) form opens.

2

Select a customer to associate with the L2 VPN TLS VLAN.

3

Configure the required general parameters.

The Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Click on the VLAN tab and set the Application parameter to L2-VPN (TLS/VLAN-Stacking).

5

Select a group to associate with the L2 VPN in the Group panel.

6

On the navigation tree, right-click on VLAN and choose Create VLAN Site to select a site for the VLAN. The Site (Create) form opens.

7

Configure the required general parameters.

8

Perform 75.12 “To add Ethernet services to a VLAN Site” (p. 2082) to add one or more Ethernet services to the VLAN site.

9

Click on the STP tab and configure the parameters.

10

To enable DHCP snooping on the VLAN:

1. Click on the DHCP Snooping tab.
2. Click Create. The VLAN Level DHCP Snooping (Create) form opens.
3. Configure the required parameters.

   These parameters are automatically enabled when DHCP snooping is enabled on a VLAN.
4. Save the changes and close the form.

11

To add one or more static entries to the DHCP binding table:

1. Click on the Binding Database tab.
2. Click Create. The DHCP Snooping Binding Database (Create) form opens.
3. Configure the MAC Address parameter.
4. Select a port.
5. Configure the required parameters.
6. Save the changes and close the form.

12

Save the changes and close the forms.

END OF STEPS

75.9 To create an OmniSwitch BTV VLAN service

75.9.1 Steps

1

Choose Create→Service→VLAN from the NFM-P main menu. The VLAN (Create) Service form opens.

2

Select a customer to associate with the VLAN.

3

Configure the required general parameters:

The Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Click on the VLAN tab and set the Application parameter to Broadcast TV (MVR/IPMV).

5

Select a group to associate with the VLAN in the Group panel.

6

Configure the Type parameter in the IPMVLAN panel.

7

On the navigation tree, right-click on VLAN and choose Create VLAN Site to select a site for the VLAN. The Site (Create) form opens.

8

Configure the required general parameters.

9

Perform 75.11 “To associate an access interface with a VLAN service” (p. 2081) to create one or more access interfaces on the site.

10

Perform 75.12 “To add Ethernet services to a VLAN Site” (p. 2082) to add one or more Ethernet services to the VLAN site.

11

Click on the STP tab and configure the parameters.

12

To enable DHCP snooping on the VLAN:

1. Click on the DHCP Snooping tab.
2. Click Create. The VLAN Level DHCP Snooping (Create) form opens.
3. Configure the required parameters.

   These parameters are automatically enabled when DHCP snooping is enabled on a VLAN.
4. Save the changes and close the form.

13

To add one or more static entries to the DHCP binding table:

1. Click on the Binding Database tab.
2. Click Create. The DHCP Snooping Binding Database (Create) form opens.
3. Configure the MAC Address parameter.
4. Select a port.
5. Configure the required parameters.
6. Save the changes and close the form.

14

To add one or more multicast group addresses to the site:

1. Click on the Multicast Groups tab.
2. Click Create. The Multicast Group (Create) form opens.
3. Configure the Multicast Address parameter.
4. Save the changes and close the form.

15

To add one or more customer VLAN tags to the site:

1. Click on the Customer VLAN Tags tab.
2. Click Create. The Customer VLAN Tag (Create) form opens.
3. Configure the Customer VLAN Tag parameter.
4. Save the changes and close the form.

16

Save the changes and close the forms.

END OF STEPS

75.10 To create an OmniSwitch VIP VLAN service

75.10.1 Steps

1

Choose Create→Service→VLAN from the NFM-P main menu. The VLAN (Create) Service form opens.

2

Select a customer to associate with the VLAN.

3

Configure the required general parameters:

The Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Click on the VLAN tab and set the Application parameter to VIP VLAN.

5

Select a group to associate with the VLAN in the Group panel.

6

On the navigation tree, right-click on VLAN and choose Create VLAN Site to select a site for the VLAN. The Site (Create) form opens.

7

Configure the required general parameters.

8

Click on the CFM tab and configure the required parameters.

9

Save the changes and close the forms.

10

Perform 75.11 “To associate an access interface with a VLAN service” (p. 2080) to create one or more access interfaces on the site.

END OF STEPS

75.11 To associate an access interface with a VLAN service

75.11.1 Purpose

Associate access interfaces with VLAN services. These access interfaces are the physical ports to which end users connect. Use the L2 Interfaces tab on the VLAN properties form to associate VLAN services with the ports used by end users.

75.11.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLAN service.

3

On the service navigation tree, expand the site on which you want to associate the new access interface, right-click on L2 Access Interfaces and choose Create VLAN Access Interface. The VLAN Access Interface (Create) form opens.

4

Configure the required general parameters.

5

Click on the Port tab and select a port for the VLAN access interface.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click on the Search button.
To configure an access interface, a L2 VPN TLS VLAN user port must be created. See 75.8 “To create an OmniSwitch L2 VPN TLS VLAN service” (p. 2076) .

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

75.12 To add Ethernet services to a VLAN Site

75.12.1 Purpose

During Ethernet service creation, you can create one or more service access multipoints. A service access multipoint is assigned a user-selected or auto-generated ID when the multipoint is created.

VLAN service access points, a SAP policy, and customer VLANs are associated with service access multipoints. Customer VLANs associate customer traffic with a VLAN-stacking SAP and identify the type of customer traffic that is received on the SAP UNI ports.

75.12.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLAN service.

3

On the service navigation tree, expand the site on which you want to add the Ethernet service, right-click on Ethernet Services and choose Create Ethernet Service. The Ethernet Service (Create) form opens.

4

Configure the Ethernet Service Name parameter.

5

On the service navigation tree, right-click on Ethernet Service and choose Create Service Access Multipoint. The Service Access MultiPoint (Create) form opens

6

Configure the required general parameters.

7

If you need to apply a SAP policy other than the default one, click on the Clear button to clear the default SAP policy from the SAP Profile panel. Otherwise, go to Step 9 .

8

Select a SAP profile to associate with the service access multipoint.

9

On the service navigation tree, right-click on Service Access Points and choose Create VLAN Service Access Point. The VLAN Service Access Point (Create) form opens.

10

Configure the required general parameters.

11

Click on the Port tab.

12

Select a port for the VLAN SAP.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click on the Search button.

13

Save the changes and close the form.

14

To add one or more customer VLANs to the service access multipoint:

1. On the Service Access MultiPoint (Create) form, click on the Customer VLANs tab.
2. Click Create. The Customer VLAN (Create) form opens.
3. Configure the required parameters.

15

Save the changes and close the forms.

END OF STEPS

75.13 To create a VLAN group

75.13.1 Steps

1

Choose Manage→VLAN Groups from the NFM-P main menu. The Manage VLAN Groups form opens.

2

Click Create. The VLAN Group (Create) form opens.

3

Configure the parameters.

You can configure the Head Ends parameter when the Node Type parameter is set to OMNI.

4

Click Apply.

5

Click on the Group Members tab.

6

Click Create and choose one or more NEs. Only NE types specified by the Node Type parameter are displayed in the list.

7

To add head end NEs to the group:

1. Ensure that the Head Ends parameter is enabled in Step 3 .
2. Click Add Headend Node to select one or more NEs. Only 7750 SR, and 7450 ESS NE types can be selected.

8

To apply a span of control to a group, other than the default:

1. Click on the Spans tab.
2. Click Create to select one or more spans of control to apply to the VLAN group.

9

Save the changes and close the forms.

END OF STEPS

75.14 To delete a VLAN group or group member

Note: You must delete all members from a group before you can delete the group.

75.14.1 Steps

1

Choose Manage→VLAN Groups from the NFM-P main menu. The Manage VLAN Groups form opens.

2

Choose a VLAN group and click Properties. The VLAN Group (Edit) form opens.

3

Click on the Group Members tab.

4

Choose one or more members of the VLAN group and click Delete button.

If the group does not contain any members, go to Step 5 .

5

Close the form.

6

On the Manage VLAN Groups form choose the VLAN group that you need to delete.

7

Click Delete.

8

Close the Manage VLAN Groups form.

END OF STEPS

75.15 To manually add MEPs to an OmniSwitch VLAN service access interface

75.15.1 Overview

See Chapter 91, “Ethernet CFM” for more information about MEPs.

75.15.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLAN Service and click Properties. The VLAN - *Service Name* (Edit) form opens.

3

On the service navigation tree, expand the site to which you want to add the MEP, right-click on the access interface and click Properties. The VLAN Access Interface (Edit) form opens.

4

Click on the MEPs tab and click Create. The Select Maintenance Entity Group form opens.

5

Select an entry and click OK.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

75.16 To configure IGMP on an OmniSwitch VLAN site

75.16.1 Purpose

OmniSwitch IGMP parameters can be configured globally and per VLAN site. You can configure VLAN site IGMP parameters only after you create the VLAN. VLAN IGMP configuration settings override global IGMP settings. See 28.105 “To configure IGMP on an OmniSwitch” (p. 1019) for information about configuring global OmniSwitch IGMP parameters.

75.16.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an OmniSwitch VLAN service and click Properties. The VLAN *Service Name* (Edit) form opens.

3

On the service navigation tree, right-click on the site on which you want to configure IGMP parameters. The Site (Edit) form opens.

4

Click on the IGMP tab.

5

Choose an entry from the list and click Properties. The IGMP (Edit) form opens.

6

Configure the required parameters.

When the Administrative State parameter is Up, IP multicast switching is enabled on the OmniSwitch for standard and L2 VPN (TLS/VLAN-Stacking) VLAN services.

7

To create one or more static IGMP groups:

1. Click on the Multicast Group tab.
2. Click Create. The Group (Create) form opens.
3. Select a terminating port.
4. Configure the Multicast Group IP Address parameter.
5. Save the changes and close the form.

8

To create one or more static IGMP neighbors:

1. Click on the Multicast Neighbor tab.
2. Click Create. The Neighbor (Create) form opens.
3. Select a terminating port.
4. Save the changes and close the form.

9

To create one or more static IGMP queriers:

1. Click on the Multicast Querier tab.
2. Click Create. The Querier (Create) form opens.
3. Select a terminating port.
4. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

75.17 To configure MLD on an OmniSwitch VLAN site

75.17.1 Purpose

OmniSwitch MLD parameters can be configured globally and per VLAN site. You can configure VLAN site MLD parameters only after you create the VLAN.

75.17.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an OmniSwitch VLAN service and click Properties. The VLAN *Service Name* (Edit) form opens.

3

On the service navigation tree, right-click on the site on which you want to configure MLD parameters. The Site (Edit) form opens.

4

Click on the MLD tab.

5

Choose an entry from the list and click Properties. The MLD (Edit) form opens.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

75.18 To configure RA filtering on an OmniSwitch VLAN site

75.18.1 Purpose

Perform this procedure to configure RA (Router Advertisement) filtering that can be used to prevent the spread of unwanted RAs from unauthorized systems. Once enabled on an interface, any received RAs will be dropped without being forwarded to other connected IPv6 clients. One or more trusted ports or link aggregation group members can be specified for an interface. RAs received on these trusted ports or link aggregation group members will be allowed to continue on to all other IPv6 clients reached via the interface.

Note: You can configure RA filtering only after the VLAN site is created. RA filtering cannot be enabled on a VLAN site that is used as an OpenFlow VLAN, Ethernet service VLAN, mirrored port, or IPM VLAN.

75.18.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an OmniSwitch VLAN service and click Properties. The VLAN *Service Name* (Edit) form opens.

3

On the service navigation tree, right-click on the site on which you need to configure RA filtering parameters. The Site (Edit) form opens.

4

Click on the RA Filter tab. By default, the RA Filter Status parameter is set to Disabled.

If you are configuring the RA Filter for 6900/6860/6865/6465, an additional RA Filter Interface parameter is displayed.

5

Choose Enable from the drop-down menu.

If you are configuring the RA Filter for 6900/6860/6865/6465, enter the RA Filter Interface name. The RA Filter status can be enabled only if the RA Filter Interface name is entered. Providing the RA Filter Interface name along with the RA Filter in enabled status will create a default IPv6 interface on the node.

6

Click on the Trusted Ports/LAGs tab.

7

Click on the Create tab and configure RA filtering on the trusted ports/LAGs. By default, all the ports/LAGs are untrusted.

8

Click Apply to save the changes and close the form.

END OF STEPS

75.19 To run an OAM validation test on a VLAN service

75.19.1 Prerequisites

A validator test suite must be created for the tested entity. See Chapter 90, “OAM diagnostic tests” for more information about how to create a validator test suite.

Note: As an alternative, you can also run an OAM validation test on the service by performing a One Time Validation. This is a mostly automated procedure and is described in 90.55 “To run a one-time validation test on a service” (p. 3084) .

75.19.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VLAN (Edit) form opens.

3

Click Validate. If the Validate button is not visible, click on the More Actions button and choose Validate. If a validator test suite is not associated to the service, a dialog box appears.

Perform the following steps:

1. Click OK to associate the service with an existing validator test suite. The Choose Validator Test Suite form appears.
2. Select a validator test suite and click OK button. The Choose Validator Test Suite form closes.

4

View the State Cause indicators. When the validation test fails, a check mark beside the OAM Validation Failure indicator.

5

Click on the Tests tab.

6

Click on the Tested Entity Result tab.

7

Select an entry and click Properties. The Tested Entity Result form opens and displays information about the validation test.

8

Close the forms.

END OF STEPS

75.20 To view the VLAN service operational status

75.20.1 Purpose

The Aggregated Operational State and State Cause indicators on the General tab of the service management form display information about service faults.

75.20.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VLAN *Service Name* (Edit) form opens.

3

View the Aggregated Operational State and State Cause indicators. When the Aggregated Operational State is Down or Partially Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The Object Alarms tab appears.

6

Click on the Aggregated Alarms tab to view the aggregated alarms for the object. The Aggregated Alarms tab appears.

7

Close the forms.

END OF STEPS

75.21 To delete a VLAN service

CAUTION

Service Disruption

Network service disruption

Deleting a service may result in a service disruption for customers. Consider the implication of deleting the service before proceeding.

CAUTION

Service Disruption

NFM-P service disruption

Do not delete any VLANs that are used for network management, otherwise connectivity between the NFM-P and the OmniSwitch devices is lost.

75.21.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service or a range of services from the list.

3

Click Delete. A warning form appears. This message is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 4 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 4 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 4 .

4

For all services regardless of how their priority is configured, acknowledge the check box that prompts you confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

5

Click Yes to confirm the action. The service is deleted and removed from the list.

6

Close the forms.

END OF STEPS

76 VLL service management

76.1 Overview

76.1.1 Purpose

The NFM-P supports the provisioning of VLL services on edge devices. A VLL service is an L2 point-to-point service that connects access interfaces. A VLL service is completely transparent to customer or subscriber data and to control protocols. Because of this, the device performs no MAC learning in a VLL service.

The NFM-P supports multiple variations of VLL services. See 76.2 “VLL service management overview” (p. 2100) in this section for more information.

A VLL service that connects access interfaces on one device is called a local VLL service. As there is no need for signaling between devices, a local VLL service uses no SDPs.

A VLL service that connects access interfaces on two devices is called a distributed VLL service. Subscriber or customer data enters a distributed VLL service through access interfaces on different edge devices. The VLL service encapsulates the data and transports it across a service provider IP/MPLS network through GRE, MPLS, or MPLS-TP service tunnels.

Packets that arrive at an edge device are associated with a VLL service based on the access interface on which they arrive. An access interface is uniquely identified using these parameters.

- physical port or POS port and channel
- encapsulation type
- encapsulation identifier (if required, depending on encapsulation type)

A VLL service uses T-LDP signaling, and uses MPLS or GRE as the service tunnel transport.

A new or existing VLL can be configured as the spoke of an HVPLS. See 77.2.1 “HVPLS” (p. 2210) in “VPLS management overview” (p. 2210) and Chapter 85, “Composite service management” for more information.

The NFM-P supports end-to-end VLL configuration using tabbed configuration forms with an embedded navigation tree.

Common to all device services, such as VLL, are policies that are assigned to the service. Policies are defined at a global level and can then be applied to components of the service, such as interfaces or circuits, when the service is configured or modified. The following policies are common to all device services:

- QoS policies define ingress classification, policing, shaping, and marking on the ingress side of the interface. QoS policies are configured using the SAP Access Ingress Policy form, the SAP Access Egress Policy form, and the ATM QoS Policy form. Because a VLL service is a point-to-point service, ingress QoS policies create only the unicast queues that are defined in the policy and not the multicast queues.

- (Epipe service only) Policer control policies to control access ingress policers and access egress policers under a common hierarchy. Policer control policies are configured using the Policer Control Policy Manager.
- Scheduling policies define hierarchical rate limiting and scheduling to govern the scheduling of queues. Scheduler policies are configured using the Scheduler Policy and HSMDA Scheduler Policy forms.
- Port scheduler policies define hierarchical bandwidth allocation and scheduling at the egress port level. Port scheduler policies are configured using the Port Scheduler Policy form.
- Filter policies control network traffic into or out of an interface or circuit based on IP or MAC matching criteria. Filter policies are configured using the ACL IP Filter form and the ACL MAC Filter form.
- Accounting policies measure the traffic on a service to ensure proper billing and enforcement of SLAs. Accounting policies are configured using the Accounting Policy form.
- ANCP policies provide status and control information based on port-up and port-down messages and changes to the current access line rate between the edge device and the access node. ANCP policies are configured using the Manage Subscriber Policies form.
- Time of day suites specify time and day restriction policies that are assigned to QoS policies and schedulers, ACL filters, and aggregation schedulers. Time of day suites and time range policies are configured using the Time of Day Suite form and Time Range form, respectively.

See Chapter 49, “Policies overview” for more information about policies.

OAM diagnostics can be performed on a per-service basis. See Chapter 90, “OAM diagnostic tests” for more information.

The General tab of the VLL service management form displays information about the operational state of the service and its sites through the Operational State and State Cause indicators.

The Operational State indicator identifies the states of the sites that are part of the service, as follows:

- Up—one operational path in both directions (end-to-end)
- Down—path is not operationally complete

When the Operational State is Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P operator.

For Epipe services, you can configure SAPs so that if their operational state goes down, the service operational state remains up, to allow performance monitoring and ETH CFM on the service. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) and the NE documentation for more information.

You can run the OAM Validation test suite for the service by clicking Validate. Alternatively, you can also perform a One Time Validation. If a check mark appears beside the OAM Validation Failed state cause indicator, the test has failed. In addition, the Tested Entity Result tab on the Tests tab displays detailed information about the OAM test result. See Chapter 89, “Service Test Manager” for more information about how to configure OAM validation test suites.

The NFM-P also monitors the status of a peer SAP after a VLL has been created and put into service. Status information includes faults detected on the service tunnel, and access and network

SAP transmissions and receptions. The States tab of the Spoke SDP Binding form displays indicators of failure in the VLL in the State Cause panel.

When you use the NFM-P to create or discover a service, the NFM-P assigns a default Service Tier value to the service. The Service Tier parameter value is relevant only in the context of composite service topology views. See Chapter 85, “Composite service management” for more information about the hierarchical organization of composite services.

76.1.2 Contents

VLL service management overview

76.2 VLL service management overview

76.2.1 VLL types

The NFM-P supports the creation of the following VLL service types:

- Epipe, or Ethernet VLL service
- Apipe, or ATM VLL service
- Fpipe, or frame relay VLL service
- Hpipe, or HDLC service
- Ipipe, or IP interworking VLL service
- Cpipe, or circuit emulation VLL service

See the *NSP Wavence Device Support Guide* for Wavence device VLAN service management configuration information.

Epipe (Ethernet VLL)

An Epipe, or Ethernet VLL service, provides a point-to-point Ethernet service. One endpoint of an Epipe uses Ethernet encapsulation, and the other endpoint uses Ethernet, ATM, frame relay, or CEM encapsulation. An Epipe effectively provides ATM and frame relay bridged encapsulation termination for interworking. The NFM-P supports local cross-connecting when the Epipe endpoints are on the same device.

The device supports these Epipe connectivity scenarios:

- a frame relay or ATM user in an ATM network communicating with an Ethernet user on an IP/MPLS network
- a frame relay or ATM user who connects to a PE device in an IP/MPLS network and communicates with an Ethernet user who connects to another PE device in the same network

ATM users connect through a UNI using AAL-5 or bridged Ethernet PDUs, and use the VCI/VPI as the ATM SAP identifier. Frame relay users connect through a UNI that uses Multiprotocol Interconnect over frame relay or bridged Ethernet PDUs, or over an Ethernet UNI interface. The DLCI is the frame relay SAP identifier. The VCI/VPI and DLCI identifier tags are transparent to the service and remain unaffected during transport.

Devices which support MEF 8 can support TDM services encapsulated on Epipe services using circuit emulation mode. A SAP with a TDM port can then be assigned to an Epipe service using CEM encapsulation. The TDM SAP defines a local and remote ECID as well as a remote MAC address. The MAC address of the TDM port is used as the source MAC address for the emulated circuit.

MEF 8 can be configured on an Epipe service in the following two scenarios:

- one TDM SAP and one Ethernet SAP
- one TDM SAP and one spoke SDP

Using MEF 8 functionality to emulate TDM services, Epipe services can interoperate with devices that support MEF 8 Epipe services, but not MPLS Cpipe services, such as the Wavence SM. MEF 8 is supported on the 7450 ESS, 7705 SAR, and 7750 SR.

The NFM-P supports the following Ethernet SAP encapsulations for an Epipe service:

- null
- dot1q
- QinQ
- SONET/SDH BCP-null
- SONET/SDH BCP-dot1q

Epipe services allow you to designate a dot1q encapsulated SAP as the default SAP for a specific port. For more information about how you can use default SAPs on an Ethernet port, see 77.2.26 “Default SAPs” (p. 2234) in “VPLS management overview” (p. 2210) .

A default SAP can co-exist with other SAPs on a port, but it cannot be implemented on a null encapsulated port. You can use the NFM-P GUI to create a default SAP by specifying an outer encapsulation value of 4095 or * for a SAP. If the XML API is used, the outer encapsulation value is always 4095.

When you enable the Enable Q in Q Untagged Sap parameter on an NE, the NFM-P allows the creation of the following two default SAP types:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- The SAP type *.* functions as a default SAP for double-tagged frames on a Q in Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

SAP types X.0 and X.* as well as SAP types *.null and *.* can be configured on the same Q in Q port on an Epipe service.

QTag Manipulation

The NFM-P supports QTag Manipulation on VLL access interfaces, on supporting NEs. QTag Manipulation allows you to define actions for inner and outer VLAN IDs, and assign tag values where required. For example, you can configure QinQ Tunneling, VLAN Translation, or other manipulations for L2 traffic that ingresses and egresses the interface. QTag Manipulation is configured on the Port tab of the L2 Access Interface properties form. See the NE documentation for more information about QTag Manipulation actions.

3-plus-tag Epipe service

You can configure a 3-plus-tag Epipe service on supporting 7210 SAS NEs. A 3-plus-tag Epipe service enables processing of packets with more than two tags when received on a QinQ SAP. A 3-plus-tag Epipe service requires a QinQ SAP as one endpoint and one of the following as the other endpoint:

- a spoke SDP binding
- a dot1q L2 access interface
- a QinQ L2 access interface

The spoke SDP binding must be of VC type VLAN, and the encapsulation values, or tags, must be configured to match. See 76.28 “To configure a VLL site” (p. 2154) , 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) , 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) , and the NE documentation for more information.

Support for 3-plus-tag Epipe services on 7210 SAS NEs varies, depending on the chassis type and release. See the NE documentation for support information.

Ethernet CFM is supported for 3-plus-tag Epipe services.

Apipe (ATM VLL)

An Apipe, or ATM VLL service, provides a point-to-point ATM service between users who connect to 7750 SR, 7705 SAR, 7950 XRS, or 7450 ESS NEs in an IP/MPLS network directly or through an ATM access network. One endpoint of an Apipe uses ATM encapsulation and the other endpoint uses ATM or frame relay encapsulation. An ATM PVC—for example, a VC or a VP—is configured on the managed device. As a result, the ATM switches at the service endpoints appear to be directly connected over an ATM link. The NFM-P supports VPI/VCI translation in an Apipe and supports local cross-connecting when the Apipe endpoints are on the same managed device.

An Apipe encapsulates standard UNI/NNI cells that ingress the ATM SAP into a pseudowire packet using N:1 cell mode encapsulation or AAL-5 SDU mode encapsulation. When using N:1 cell mode encapsulation, an Apipe supports cell concatenation into a pseudowire packet and the setup of both VC- and VP-level connections.

For ATM and frame relay interworking, an Apipe provides a point-to-point service between a user who connects to an existing ATM network and another user who connects to a PE in an IP/MPLS network. An ATM AAL-5 SDU pseudowire or a frame relay 1-to-1 mode pseudowire connects the nodes. The PE performs an FRF.5 interworking function to join the ingress and egress data paths.

An ATM VT SAP on a PE is identified by the physical port and VPI range. Cells that arrive on a specified port and are within the specified VPI range go into a single pseudowire for transport through the IP/MPLS network. A user can configure the whole ATM port as a VT and does not need to specify a VPI range. There is no ingress or egress VPI/VCI translation or loss of cell order.

The 7705 SAR uses N>1 cell mode encapsulation to support multiple VCs on an Apipe. You can multiplex VCs on a 7705 SAR Apipe service using SAP aggregation groups. See 76.19 “To create a SAP aggregation group on a 7705 SAR Apipe” (p. 2140) for more information about how to configure SAP aggregation groups.

The NFM-P supports the ATM N:1-N>1 VC type to allow interoperability between 7705 SAR and 7750 SR Apipe services, despite each site using different cell mode encapsulation. An interoperating Apipe service uses the ATM N:1-N>1 VC type, while the 7705 SAR site uses ATM VCC and the 7750 SR uses ATM Cell. An interoperating Apipe service combines SAP aggregation groups from the 7705 SAR Apipe with connection profiles from the 7750 SR Apipe onto a single Apipe service. See Chapter 63, “Connection profile policies” for more information about connection profiles.

Fpipe (frame relay VLL)

An Fpipe, or frame relay VLL service, provides a point-to-point frame relay service between users who connect to PE 7750 SR or 7450 ESS NEs in an IP/MPLS network. Both endpoints of an Fpipe use frame relay encapsulation. An Fpipe connects users through frame relay PVCs. An Fpipe

receives standard Q.922 core frames on the frame relay SAP and encapsulates them in a pseudowire packet according to the 1-to-1 frame relay encapsulation mode. This is the VC type used on the SDP by default. The NFM-P does not support the many-to-one, or port encapsulation, mode. Fpipe creation supports local cross-connecting when the endpoints are on the same managed device.

Hpipe (HDLC VLL)

An Hpipe, or HDLC VLL service, is used to carry HDLC PDUs over an MPLS network. HDLC PWs enable service providers to offer emulated HDLC services over existing MPLS networks. HDLC mode provides port-to-port transport of HDLC-encapsulated traffic.The HDLC PDU is transported in its entirety, including the HDLC address and control fields, but the HDLC flags and the FCS are excluded. If the optional control word is used, the flag bits in the control word are not used and must be set to 0 for transmitting and must be ignored upon receipt.

Before HDLC SAPs can be configured, the mda-mode command must be set to cem-fr-hdlc-ppp at the card level. Only ports that are configured with HDLC encapsulation can be mapped to an Hpipe SAP.HDLC encapsulating ports do not terminate the HDLC. The ports pass the HDLC frames through the Hpipe. HDLC encapsulated ports can pass through any HDLC-framed traffic, such as Cisco-HDLC, FR, PPP, etc.

Ipipe (IP interworking VLL)

An Ipipe, or IP interworking VLL service, uses the IP/MPLS network to provide Layer 3 connectivity between different Layer 2 technologies. An Ipipe service provides point-to-point IP connectivity between a user on a frame relay, ATM, cHDLC, or PPP access circuit with routed PDU IPv4 encapsulation and a user on an Ethernet interface. The Ethernet SAP interface can terminate on a 7705 SAR, 7750 SR, 7450 ESS, or 7950 XRS.

The following table summarizes the supported SAP types.

Table 76-1 Supported SAP types

SAP Types	Frame relay	ATM	PPP/IPCP	cHDLC	Ethernet
Frame relay					✓
ATM					✓
PPP/IPCP			✓		✓
cHDLC				✓	✓
Ethernet	✓	✓	✓	✓	✓

In an Ipipe service, both CE devices appear to be on the same IP interface. The PE devices must therefore resolve Layer 2 addresses when different resolution protocols are used on either SAP. Each PE device is manually configured with the IP addresses of both CE devices, or alternatively, can be set to automatically discover the IP addresses of the CE routers. The PE device maintains an ARP cache context for each IP interworking VLL, and responds to ARP request messages received on the Ethernet SAP. The PE device responds with the Ethernet SAP configured MAC address as a proxy for an ARP request for the frame relay, ATM, or PPP user access circuit IP address, and silently discards any ARP request message received on the Ethernet SAP for any

other address. The PE device maintains a record of the association of IP addresses with MAC addresses for ARP requests that it receives over the Ethernet SAP.

An Ipipe SAP can be bound to a physical or logical port with PPP, cHDLC, ATM, or FR encapsulation. ATM users connect through a UNI using AAL-5 MUX IP or AAL-5 SNAP routed PDU encapsulation. Frame relay users connect using routed PDU IPv4 encapsulation. PPP interfaces use PPP/IPCP encapsulation of an IPv4 packet. Users of cHDLC connect using routed IPv4 encapsulation.

The NFM-P supports the following Ethernet SAP encapsulations for an Ipipe service:

- Null
- dot1q
- QinQ

Note: IPCP SAPs on the 7705 SAR can be configured to assign primary and secondary DNS addresses to the remote peer.

The following identifiers are used for packet forwarding:

- VCI/VPI as the ATM SAP identifier
- DLCI as the frame relay SAP identifier

Cpipe (circuit emulation VLL)

A Cpipe, or circuit emulation VLL service, provides a point-to-point CEM service between supporting devices. The endpoints of a Cpipe use CEM encapsulation.

The Cpipe L2 access interface can be bound to a unstructured DS1 or E1 channel, a channelized DS0 channel group, or a DS0 group with CAS signaling. Consider the following when creating a Cpipe:

- The Time Slots parameter of the DS0 channel must be configured with at least one time slot.
- Time slots are automatically configured for unstructured E1 and T1 endpoints.
- The Clock Source parameter of the DS1 channel must be set to Node-Timed.

76.2.2 VLL spoke switching

VLL spoke switching allows you to create a VLL service by cross-connecting two spoke SDPs. Spoke switching allows you to scale L2 services, such as VLLs and H-VPLS, over a multi-area network without the requirement for a full mesh T-LDP. The NFM-P supports spoke switching on all VLL types, however, all service instances must be the same type.

The NFM-P supports VLL services with spoke switching on the 7210 SAS, 7250 IXR, 7450 ESS, 7750 SR, and 7950 XRS. Support for 7210 SAS NEs varies depending on the chassis type and release; see the NE documentation for information.

The following table describes the VLL site types that you can use in a spoke switching configuration.

Table 76-2 VLL site types

VLL site type	Description
Terminating	VLL instance has one or two VLL SAPs.
Switching	VLL instance cross-connects two spoke bindings.

The following figure shows a switching VLL that connects two VPLS. The configuration uses two T-PEs, three S-PEs, and two pairs of spoke bindings to connect the two VPLS.

Figure 76-1 VLL spoke switching

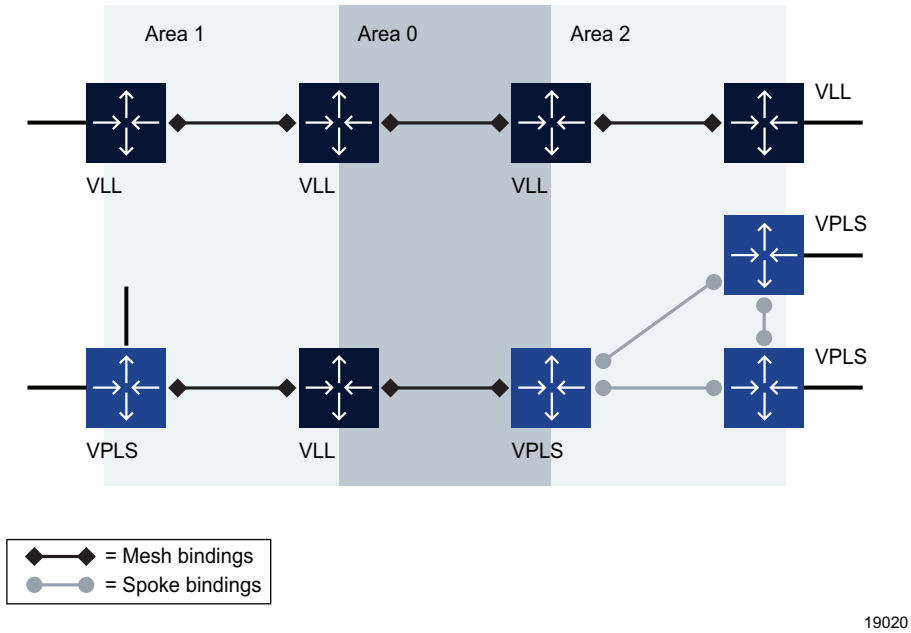

Configuration requirements

You must consider the following configuration rules when provisioning VLL services with spoke switching:

- Service type must be a VLL at the switching node.
- VC IDs can be different for the segments in the service.
- Segments in the service can run on different types of tunnels; for example, LDP, GRE, and RSVP SDPs.
- VLL instances in a service must use the same service ID to avoid the discovery of multiple and composite services.
- VLL with one or more switching sites must have two or more terminating sites.
- Autobinding creation service is not supported for switching VLLs.

76.2.3 VLL redundancy

VLL redundancy requires that you associate the SAP or SDP bindings to an endpoint. You can configure the endpoint association as active or standby so that you can create a redundant configuration. The associated nodes use signaling to determine the active SAP or SDP binding. The NFM-P supports VLL redundancy on SR OS NEs.

A VLL service site can have up to two local endpoints. A local endpoint combines a SAP with a binding (access) or a group of bindings (network). A SAP or an SDP binding can also exist without an endpoint association.

The 7450 ESS, 7750 SR, and 7950 XRS support HSDPA offload fallback for VLL Apipe and Epipe services by allowing fallback from an active PW on a primary spoke SDP to a secondary SAP.

The following table describes the components in redundant VLL configurations.

Table 76-3 VLL components for redundant configurations

VLL component	Description
Primary or Redundant binding	Primary or redundant binding is the same as a regular spoke binding. Up to four spoke bindings can form a VLL instance network endpoint. Only one binding can be configured as the primary; up to 3 others can only be configured as redundant spokes. Each redundant spoke has a precedence value to decide which spoke is the immediate backup. Only the terminating VLL instance can have multiple bindings on the network side endpoint. A switching VLL instance has up to 2 bindings, one on each side.
Inter-Chassis Backup	You can use an ICB in conjunction with a redundant SAP to provide protection for the SAP. The SAP must also be associated with a MC-LAG or MC-APS port. The ICB transports network traffic to the SAP on the second PE when the local SAP is unavailable. You must define a switching state for the redundant SAP. You must also configure the return ICB on the opposite endpoint of the protected site.

Configuration options

The following table describes the redundancy configuration options for each VLL site.

Table 76-4 VLL redundancy configuration options

Option	Configuration	Description
1	SAP SDP binding	You can create an endpoint without any SAP or SDP.
2	SAP SAP	You can create an endpoint without any SAP or SDP.
3	SDP binding SDP binding	You can configure this option for switching sites only.
4	SAP Endpoint with SDP bindings (maximum of 4 spokes and 1 ICB)	Figure 76-2, “VLL redundancy configuration - option 4” (p. 2107) shows the configuration of this option on the PE nodes.
5	Endpoint with SAP/ICB SDP binding	Figure 76-3, “VLL redundancy configuration - option 5” (p. 2108) shows the configuration of this option on Node B.

Table 76-4 VLL redundancy configuration options (continued)

Option	Configuration	Description
6	Endpoint with SAP/ICB endpoint with SAP/ICB	Figure 76-4, “VLL redundancy configuration - option 6” (p. 2108) shows the configuration of this option.
7	Endpoint with SAP/ICB Endpoint with up to 4 SDP bindings	Figure 76-5, “VLL redundancy configuration - option 7” (p. 2109) shows the configuration of this option.

Figure 76-2 VLL redundancy configuration - option 4

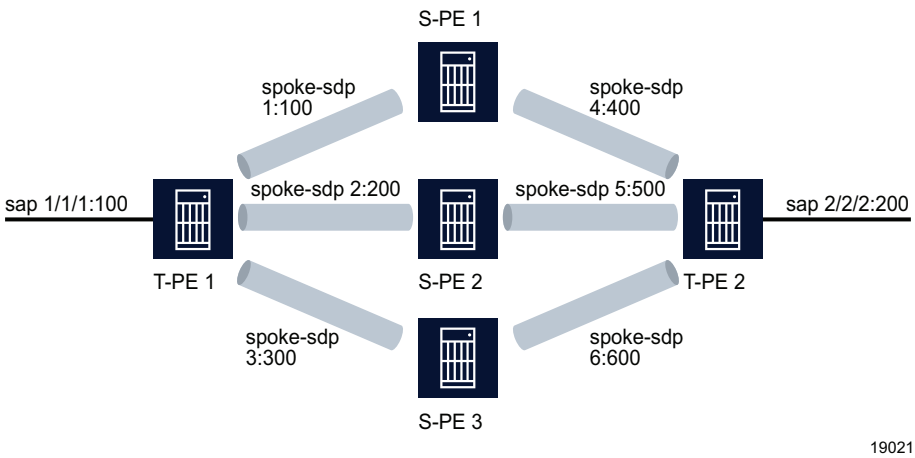

Figure 76-3 VLL redundancy configuration - option 5

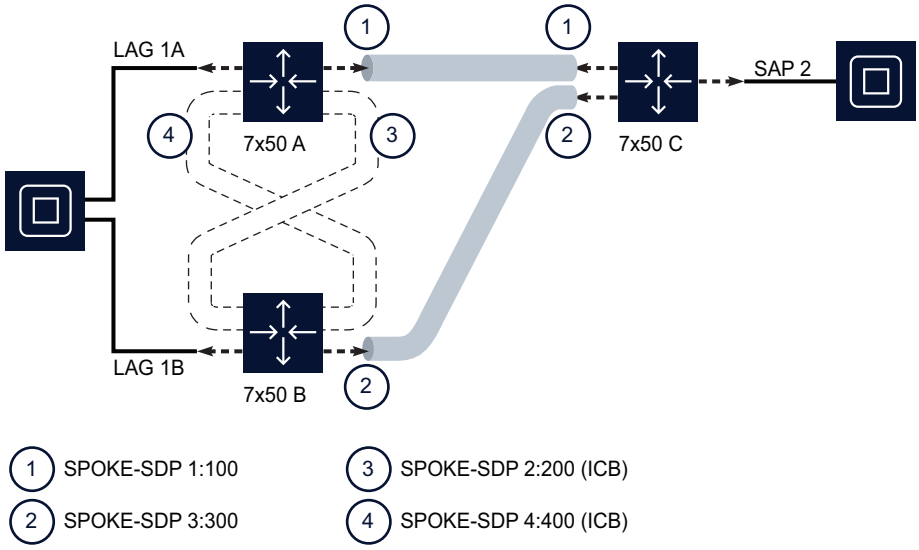

Figure 76-4 VLL redundancy configuration - option 6

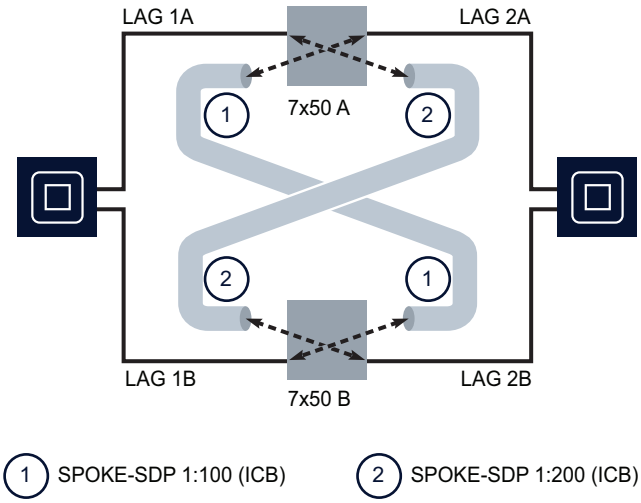

Figure 76-5 VLL redundancy configuration - option 7

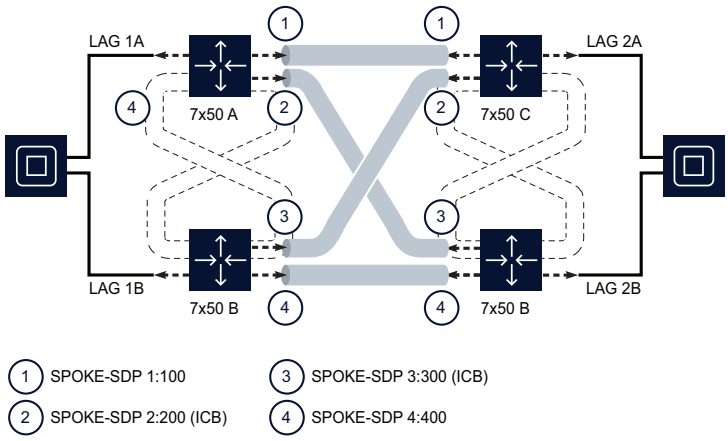

Configuration requirements

You must consider the following configuration rules when provisioning VLL services with spoke switching and redundancy:

- A VLL service can have one or more VLL instances.
- Local endpoint rules:
  - A VLL instance has a maximum of two endpoints. A terminating VLL instance has at least one access endpoint and a switching VLL instance has two network endpoints.
  - A network endpoint has a maximum of four spoke-bindings, which can include any combination of the following: a single primary spoke, one or more secondary spokes with precedence, and one ICB spoke.
  - A SAP or a binding has a maximum of one endpoint association.
  - An endpoint has a maximum of one SAP.
  - A SAP or a binding with association to an endpoint can be moved to another endpoint or removed from that endpoint.
- SAP rules:
  - A MC-LAG SAP or MC-APS SAP cannot be deleted when there is an ICB on the same endpoint.
  - SAPs cannot exist on switching sites.
  - A maximum of two SAPs can exist for each site.
  - A SAP with a non-ICB spoke cannot exist on the same endpoint.
  - Apipe and Epipe services support MC-APS.
- Spoke binding rules:
  - The SDP types (GRE, MPLS) used by the redundant spoke bindings do not have to be the same when you manually create the spoke bindings.

- Redundant configurations are not supported for S-PE because there are a maximum of two spoke bindings for a switching VLL instance.
- An ICB SDP binding should not be created on an endpoint without a MC-LAG SAP or MC-APS SAP.
- Only one ICB can exist for each endpoint.
- ICB SDP binding can only have a precedence of 4, the lowest priority.
- Only one primary spoke can exist for each endpoint.
- Spoke SDP binding cannot associate with an endpoint on a switching site.

- HSDPA offload fallback rules:
  - Apipe and Epipe services support HSDPA offload fallback on ATM interfaces.
  - The spoke SDP must be configured with primary precedence.
  - The SAP must be configured with MC-APS or ATM channel.
  - When the SAP is configured with MC-APS, the spoke SDP can be configured with ICB.

76.2.4 HSDPA Offload Resiliency

Mobile service providers deliver both voice and data services to their customers using mobile handsets. The data services provided require significantly more bandwidth than voice services. In order to minimize the operational costs (specifically, bandwidth), service providers typically separate the voice and data traffic at the mobile base station. The voice traffic may be backhauled over an ATM infrastructure, while a Metro Ethernet infrastructure (their own or third-party) is used to backhaul the data traffic. The separation of data traffic onto a separate network for backhaul is referred to as High Speed Data-Link Packet Access offload.

The HSDPA Apipe services traverse a path over the Metro Ethernet network which contains single potential points of failure that are unprotected. The ATM network can be used to provide a transient path for the data service in the event of a failure in the Metro Ethernet infrastructure, as long as the voice traffic is not impacted (data traffic is given lower QoS priority by the 7705 SAR and 7750 SR NEs). Clearly there is potential for the data service to suffer degradation (depending on the bandwidth required), until the fault in the Metro Ethernet network is resolved. However the SLA requirements for the data service are typically best effort.

The ability to switch to this alternate transient pathway for the data service is referred to as HSPDA resiliency.

HSPDA resiliency is implemented through the use of VLL Apipes on the 7705 SAR. The network architecture used in the ATM backhaul scenario is shown in Figure 76-6, “ATM-based HSDPA offload architecture - nominal operation” (p. 2111) .

Note: For HSPDA offload resiliency, the primary and secondary services must be on the same NE.

Figure 76-6 ATM-based HSDPA offload architecture - nominal operation

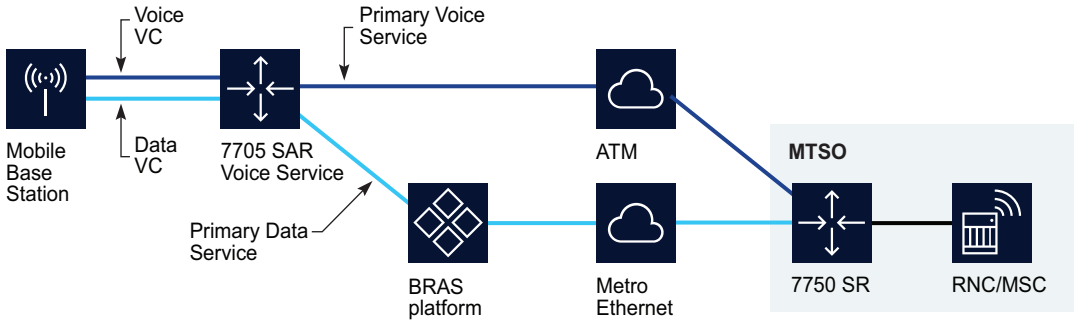

The mobile base station node separates the voice and data traffic and delivers the two different traffic types to the 7705 SAR on different VCs. Both the voice and data services use Apipe pseudo-wires to carry the traffic over the backhaul network to the 7705 SAR. The HSDPA Apipe service is paired with an OAM service (not shown in Figure 76-6, “ATM-based HSDPA offload architecture - nominal operation” (p. 2111)) which is used by the wireless infrastructure to monitor the end-to-end path between the wireless endpoints.

For the ATM voice traffic, an internal Apipe is used on the 7705 SAR to switch ATM cells between the access VC/IMA group and the network side VC/IMA group. The service does not span the radio access network between the 7705 SAR and the 7750 SR located at the Mobile Telephone Switching Office (MTSO).

Data Apipe services (HSDPA/OAM) use service tunnels based on either MPLS or GRE. GRE is generally used in the context of a third-party Metro Ethernet network, since there is a Broadband Remote Access Server (BRAS) in the path between the DSLAM and the 7750 SR located at the MTSO (BRAS platforms do not support MPLS). MPLS is typically used when the mobile provider owns the Metro Ethernet network.

When the resiliency solution described below detects a failure in the Metro Ethernet network, its protection mechanism switches the traffic associated with the data service from the path over the Metro Ethernet network to a path which traverses the ATM (shown in the following figure). The detection, switchover, and switchback mechanisms are implemented by the NFM-P.

Figure 76-7 HSDPA Offload Protection - ATM data backhaul

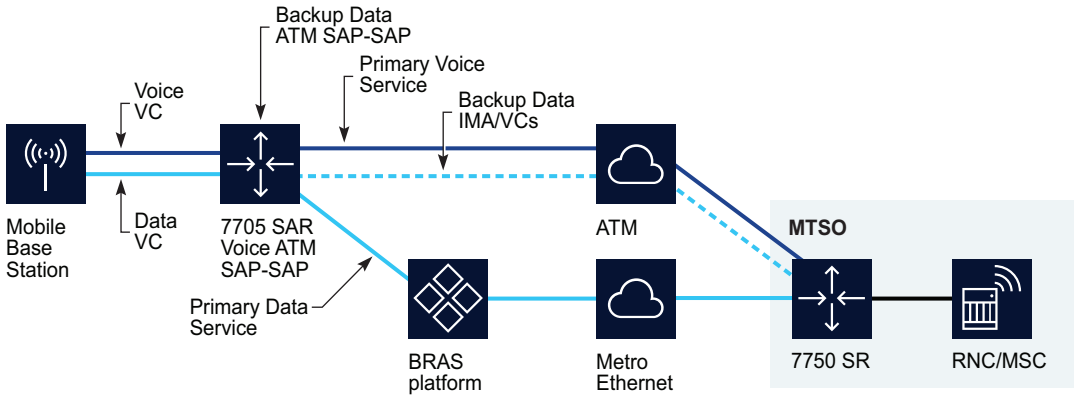

A backup (secondary) data service for an active (primary) data service is pre-configured on the 7705 SAR. For a typical 7705 SAR deployment at a cell site, two backup services are required, one for the HSDPA service and one for the OAM service. These backup services are created by the operator during the deployment of the 7705 SAR. A peer resiliency relationship between the active and backup services must also be configured.

The backup services used are internal Apipes. The switchover from primary to secondary is triggered by an event from the 7705 SAR sent to NFM-P, which indicates that the service tunnel (SDP) has failed. When the failure event is processed by NFM-P, the SAPs on the active data services are moved (along with the associated VCs) from the active service to its peer backup service and are enabled. After the re-configuration of the SAPs on the backup services completes, the traffic is moved to the backup services, which then carry the traffic over the ATM portion of the network. When the switchover is complete, an alarm is raised against the primary service indicating that it has been switched to the secondary service.

Other considerations include the following:

- If a service configured as primary is deleted from NFM-P or the CLI, then the corresponding resiliency is also deleted.
- If a service configured as secondary is deleted from NFM-P, this is blocked and a pop-up message indicates that the corresponding resiliency must first be deleted.
- If a service configured as secondary is deleted from the CLI, an alarm is raised indicating that the resiliency is misconfigured. If a secondary service is then subsequently added to the resiliency, this alarm is cleared.
- If the SAP initially configured on the primary service is deleted, the NFM-P raises an alarm. If the SAP is restored in the same service, the alarm clears.
- When the NFM-P detects that a failed SDP is up, which indicates that the primary service has recovered, the resiliency is set to the Secondary (Debounced) state and a damping timer is started. If the SDP goes down while the resiliency is in the Secondary (Debounced) state, the resiliency changes to the Secondary state and the damping timer stops. If the SDP changes state during the damping time, the value of the damping timer doubles until it rises to the maximum damping time. This prevents the service from flapping between primary and

secondary. If the SDP does not change operational state during the damping time, the resiliency is set to Primary, the SAP is moved from the secondary to the primary service, and the alarm related to the service switch clears.

- After the primary service is restored, the damping timer value is set back to its initial value. The damping timer value is not displayed. The initial damping time is 30 000 ms; the time doubles, if required, to a maximum of 480 000 ms.

76.2.5 SDP bindings bandwidth allocation

You can administratively account for the bandwidth used by VLL services inside an RSVP SDP that consists of RSVP LSPs. The SR service manager keeps track of the available bandwidth for each SDP.

When you create a service tunnel, you configure an SDP Bandwidth Booking Factor percentage, which is applied to the SDP available bandwidth. You then assign an SDP Admin Bandwidth value (in kbps) to the spoke SDP. When you bind a VLL service to this SDP, this amount of bandwidth is subtracted from the adjusted available SDP bandwidth. If you subsequently delete the VLL service binding from this SDP, this bandwidth amount is added back into the adjusted SDP available bandwidth. If you overbook the total adjusted SDP available bandwidth when adding a VLL service, a warning is issued and the binding is rejected.

This feature does not guarantee bandwidth to a VLL service, as there is no change to the data path to enforce the bandwidth of an SDP by means such as shaping or policing of the constituent RSVP LSPs. Also, this feature does not provide a CAC capability for a local VLL service which consists of a cross-connect between two SAPs.

In addition, if multipoint services such as VPLS and VPRN are using the same SDP for forwarding packets, the amount of bandwidth consumed by these services is also not accounted for. Therefore, it is advisable to dedicate an SDP for VLL services for which bandwidth reservation is required. Furthermore, VPLS and VPRN services which use separate SDPs but which forward packets over the same network port as the VLL SDP also do not have their bandwidth accounted for. This may impact the bandwidth available to the VLL services.

Auto SDP binding (for all spoke bindings or just the return binding) cannot be used when there is a bandwidth request for the binding. The converse also applies. An error message appears when saving the configuration if this conflict occurs. NFM-P checking is done for the XML API.

76.2.6 Port redundancy

CAUTION

Service Disruption

Nokia recommends that you set the OLC state of a service to Maintenance before you switch to a redundant port or channel, and reset it to the previous state after the switch is complete. See the NSP System Administrator Guide for more information.

Cpipe and Epipe services on 7450 ESS, 7705 SAR, and 7750 SR devices support port redundancy, which allows you to define a backup port or channel for a VLL SAP. The backup port or channel can be on the local site, for what is called on-node switching, or on a remote site, for what is called off-

node switching. In the event of a failure on the SAP, you can manually switch to the backup port or channel using a button on the VLL service, site, or SAP properties form, or from the port properties form.

When a SAP is switched to a backup port using off-node switching, the primary SAP on which the switching is performed is deleted along with the site and site child objects, including the associated SDP bindings. The site is subsequently created on the backup NE, as are the SAP and the other site child objects, using the same tunnel type and tunnel profile for SDP bindings as the primary site.

See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for configuration information. See 76.55 “To switch to the redundant port for a VLL SAP from an L2 access interface properties form” (p. 2199) and 76.18 “To switch to the redundant port for one or more VLL SAPs” (p. 2138) for information about switching between the redundant ports of one or multiple VLL SAPs.

Note: Ethernet OAM objects, for example, MEPs, that are associated with the protected SAP are not copied to the new SAP.

When multiple SAPs are selected, switching is attempted only for the SAPs that support port redundancy and on which port redundancy is properly configured.

Port redundancy switching fails if the SAP or port is included in one of the following redundancy objects:

- BGP-MH
- MC-APS
- MC-LAG
- MC-ring

76.2.7 Copying and moving SAPs between ports

You can copy and move SAPs between ports. See 16.6 “Copying and moving SAPs” (p. 574) in Chapter 16, “Port and channel object configuration” for more information.

76.2.8 DoS protection

To protect a VLL from a high incoming packet rate that characterizes a DoS attack, you can use the NFM-P to create DoS protection policies for VLL Epipe or Ipipe L2 access interfaces. A DoS protection policy limits the number of control-plane packets that an interface receives each second, and optionally logs a violation notification if a policy limit is exceeded. You can use the NE System Security form to view the violations for a specific NE.

You can configure a DoS protection policy to control the following on a VLL Epipe or Ipipe L2 access interface:

- the control-plane packet arrival rate per subscriber host on the interface
- the overall control-plane packet arrival rate for the interface
- whether an NE sends a notification trap if a policy limit is exceeded

Each VPLS L2 access interface on an NE that supports DoS protection is automatically assigned a default DoS protection policy. The default policy limits only the overall packet arrival rate for the interface, and cannot be deleted or modified. See the *NSP System Administrator Guide* for information about creating a DoS protection policy.

76.2.9 DDoS protection

You can configure a DDoS protection policy on a VLL Epipe or Ipipe L2 access interface. See the *NSP System Administrator Guide* for information about configuring a DDoS protection policy.

76.2.10 BGP VPWS

You can use the NFM-P to create a BGP VPWS, which is an L2 VPN service that uses BGP pseudowire signaling. BGP VPWS is implemented in the NFM-P under an Epipe site, which is configured to use BGP to establish a pseudowire with a remote Epipe site. See 76.24 “To create a BGP VPWS” (p. 2147) for information about configuring a BGP VPWS.

76.3 Sample VLL service configuration

76.3.1 Sample VLL service

The following figure shows a sample VLL service.

Figure 76-8 Sample VLL service

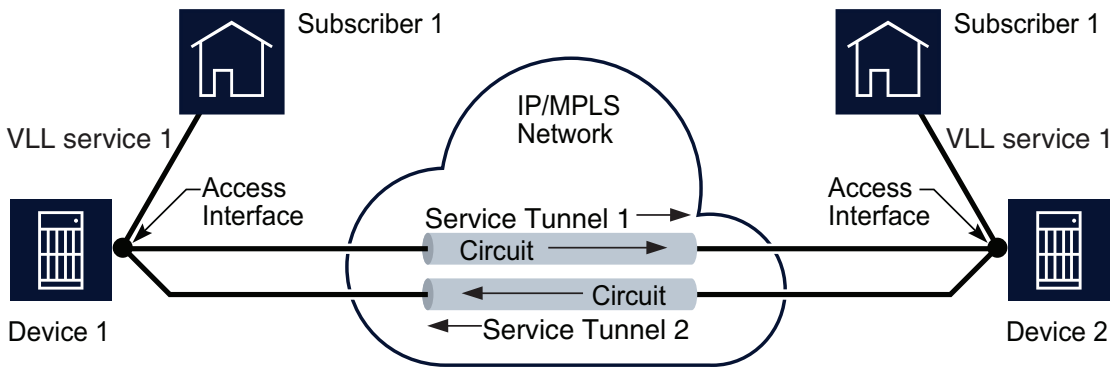

76.3.2 Configuration steps

Assuming the core IP/MPLS network and service tunnels have already been configured, the following high-level tasks are required to configure this sample VLL service.

1

Configure policies as required, for example, access ingress and egress interface, ATM, Scheduler, ACL IP and ACL MAC, Accounting, and ANCP policies.

2

Configure ports as access ports for use in the service, if applicable.

3

Configure service tunnels as required.

4

Create and configure Subscriber 1.

5

Create and configure Service 1.

VLL service management procedures

76.4 Workflow to create a VLL service

76.4.1 Overview

The following workflow lists the high-level steps required to create a VLL service. As a prerequisite for creating a VLL service, this workflow assumes the following:

- a group or customer with the required user access privileges has been configured.
- the IP or IP/MPLS core network exists.
- any required service tunnels are created including the static, dynamic or SR-TE LSP required to create the service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
- the access ports for the service are created; see Chapter 16, “Port and channel object configuration” for more information.
- any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite policies are created; see Chapter 49, “Policies overview” for more information. You do not have to create pre-defined policies if policies are created on a per-service basis.
- any required MP-BGP for PE-to-PE routing is configured; see Chapter 28, “Routing protocol configuration” for more information about protocol configuration.

76.4.2 Stages

1

Create and configure the VLL service:

1. Define the VLL service type as Epipe, Apipe, Fpipe, Hpipe, Ipipe, or Cpipe.
2. Specify the NEs (sites) used in the service based on the following topologies:
   - Traditional VLL: terminating site only
   - Switching VLL: switching site with two or more terminating sites
   - Redundant VLL: ICB spoke to MC-LAG (Epipe) or MC-APS (for Apipe and Epipe)
3. Specify the following information for redundant VLL services:
   - endpoints
   - access interfaces for VLL terminating sites
4. Specify the following information to create a VLL Apipe or Epipe HSDPA offload fallback solution:
   - endpoints
   - SAP on endpoint one with MC-APS or ATM channel
   - spoke SDP with ICB on endpoint one
   - spoke SDP with primary precedence on endpoint two
   - SAP on endpoint two with MC-APS
   - spoke SDP with ICB on endpoint two
5. Specify aggregation on a service basis, or across a card or port.

6. Specify QoS, scheduling, accounting, ANCP, MEP association and filter policies.
7. Specify the time of day suite.

2

For VLL services created on an OmniSwitch:

a. Create VLAN groups and add members to each group as required. See 75.13 “To create a VLAN group” (p. 2083) .

b. Delete a VLAN group or group member. See 75.14 “To delete a VLAN group or group member” (p. 2084) .

3

Turn up the service.

4

Perform one or more of the following as required.

a. Configure a GNE site on a VLL service. See 76.29 “To configure a GNE site on a VLL service” (p. 2155) .

b. Create a VLL L2 access interface on a terminating site. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) .

c. Create a SAP aggregation group for a 7705 SAR Apipe. See 76.19 “To create a SAP aggregation group on a 7705 SAR Apipe” (p. 2140) .

d. Configure an Epipe site for EVPN VXLAN. See 76.20 “To configure EVPN on an Epipe site” (p. 2142).

e. Create an HSDPA resiliency configuration. See 76.21 “To create an HSDPA resiliency configuration” (p. 2144) .

f. Activate and manually operate an HSDPA resiliency configuration. See 76.22 “To activate and manually operate an HSDPA resiliency configuration” (p. 2145) .

g. Run an OAM validation test for a VLL service. See 76.23 “To run an OAM validation test for a VLL service” (p. 2146) .

h. Configure a BGP VPWS. See 76.24 “To create a BGP VPWS” (p. 2147) .

5

As required, view VLL service-related information:

a. View the ECMP/LAG hashing of Epipe services to monitor the distribution of Epipe service traffic into the network. See 76.25 “To view ECMP/LAG hashing of Epipe services” (p. 2150) .

b. View the local status information associated with the VLL service. See 76.26 “To view the local PW status information for a VLL service” (p. 2151) .

c. View the peer status information associated with the VLL service. See 76.27 “To view the peer PW status information for a VLL service” (p. 2152) .

d. View the VLL service operational status. See 76.6 “To view the VLL service operational status” (p. 2120) .

e. View the service topology map associated with a VLL service. See 76.10 “To modify a VLL service using the topology view” (p. 2127) .

6

As required, modify a VLL service:

a. Using the Manage Services form. See 76.8 “To modify a VLL service” (p. 2124) .

b. Using the topology view. See 76.10 “To modify a VLL service using the topology view” (p. 2127) .

7

As required, delete a VLL service. See 76.11 “To delete a VLL service” (p. 2131) .

76.5 To create a VLL service

76.5.1 Purpose

VLL Apipe services are configurable only on the 7450 ESS in mixed mode, 7750 SR, 7950 XRS, and 7705 SAR.

VLL Hpipe services are configurable only on the 7705 SAR and 7705 SAR-M.

In order to configure a VLL Hpipe service on a supporting NE, the daughter card MDA Mode must be set to cem-fr-hdlc-ppp. For more information, see 15.78 “To configure an MDA” (p. 536) . Only ports that are configured with an Encap Type of HDLC can be mapped to an Hpipe SAP. For more information, see 16.24 “To configure Ethernet ports” (p. 599) .

VLL Cpipe services are configurable only on the 7750 SR, 7705 SAR, 7210 SAS-M, and 7450 ESS NEs. Consider the following when creating a VLL Cpipe:

- The Time Slots parameter of the DS0 channel must be configured with at least one time slot.
- Time slots are automatically configured for unstructured E1 and T1 endpoints
- The Clock Source parameter of the DS1 channel must be set to Node-Timed.

76.5.2 Steps

1

Choose Create→Service→VLL→Epipe (Ethernet to Ethernet/ATM/FR/CEM) | Apipe (ATM to ATM/FR) | Ipipe | Fpipe (FR to FR) | Hpipe (HDLC to HDLC) | Cpipe (CEM to CEM) from the NFM-P main menu. The VLL Service (Create) form opens.

2

Select a customer in the Customer panel.

3

Configure the required general parameters.

The Service ID and SVC Mgr Service ID parameters are configurable when the Auto-Assign ID parameter is disabled.

The Default VC ID parameter is configurable when the Inherit Service ID Value parameter is disabled.

4

This operation cannot be reversed.

If you will be using an MLPPP link for Epipe bandwidth management, perform the following steps to convert bandwidth units to Kbps:

1. Choose Administration→System preferences from the NFM-P main menu. The System Preferences form opens.
2. Choose the Services tab.
3. Disable Service Bandwidth Management.
4. Click Execute in the Convert to Kbps parameter.
5. Click Yes in the warning dialog box to proceed.
6. Enable Service Bandwidth Management.

5

If you enabled the Automatic SDP Binding/PBB Tunnel Creation parameter in Step 3 , select a tunnel selection profile in the Auto Tunnel Selection panel.

6

Save the changes and close the forms.

You can use the topology maps to view the service. See Chapter 4, “Topology map management” for more information about service topology maps.

END OF STEPS

76.6 To view the VLL service operational status

76.6.1 Purpose

The Operational State and State Cause indicators on the General tab of the service management form display information about service faults.

76.6.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VLL *Type_of_VLL* (Edit) form opens.

3

View the Operational State and State Cause indicators. When the Operational State is Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The Object Alarms tab is displayed.

6

Click on the Aggregated Alarms tab to view the aggregated alarms for the object. The Aggregated Alarms tab is displayed.

7

Close the VLL *Type_of_VLL* (Edit) form.

8

Close the Manage Services form.

END OF STEPS

76.7 To move a VLL service

76.7.1 Purpose

Perform this procedure to move a VLL service from one NE to another NE. The VLL service site and all of the children objects, including SAPs, SDP bindings, MEPs on the SAPs and related CFM tests, are moved to the destination site. MEPs and associated CFM tests are included in VLL move operations for L2 SAPs and SDP bindings (including return SDP bindings). Return SDP bindings (from remote site to the source site) for SDP bindings from the source site to a remote site are moved. Ingress and egress, scheduler, policer, accounting, and IP/MAC filters (ACLs) policies are moved and distributed to the destination NE.

If there are SAPs on different ports within a VLL, you must move the service port by port. When the service does not already exist on the destination NE—that is, before an move operation—the service site and children objects are moved to the destination NE. Only SAPs on the specified port are moved to the destination site. After an initial move, moving the remaining SAPs on other ports is a SAP move operation. See 16.32 “To copy or move L2 SAPs between ports” (p. 614) .

Note: Consider the following site limitations when you move a service:

- Moving services is only supported on 7750 SR, 7450 ESS, and 7950 XRS NEs.
- The NE type and major software release of the source and destination NEs must be the same.
- The chassis mode of the destination NE must be the same or higher than the chassis mode of the source NE.
- The NFM-P copies the source site and children objects and creates them on the destination NE.
- You cannot move sites that are used in a backbone configuration.
- You cannot move sites created by a RADIUS script for dynamic service.

Consider the following SAP limitations when you move a service:

- The source and destination ports must be compatible, for example, have the same encapsulation type, be HSMDA ports or not, or both be access or hybrid.
- You cannot move SAPs that are configured on PW ports and Ethernet tunnel endpoints.
- SAPs with connection profiles (ATM/VLAN) can be moved when there is at least one SAP configured on the source port with the encap value in the specified range on the source service site.
- If the source site only contains SAPs with connection profiles, the site will not be moved.
- If the source site contains SAPs associated with connection profiles on the specified port, and SAPs on a different port, or SAPs on the specified source port but with encap values out of the range, the site will not be moved.
- The ranges specified in the connection profiles (ATM/VLAN) on source SAPs must not overlap with SAPs on the destination port or fall in a range where a SAP on the destination port may exist.

  The Encap value ranges specified in the move procedure do not apply to connection profiles.
- If there is a Ethernet Ring Path Endpoint associated with the source SAP, the association is not be copied to destination SAP.

Take into account the following service tunnel considerations when you move a service:

- If the spoke SDP bindings on the source NE use an SR-ISIS or SR-OSPF service tunnel, the NFM-P looks for a matching tunnel on the destination NE. If no matching tunnel exist, then the NFM-P creates one.
- If the spoke SDP bindings on the source NE use an MPLS-TP service tunnel, the copy/move operation is allowed, but an MPLS-TP service tunnel must exist on the destination NE.
- You cannot move SDP bindings that use L2TPv3 service tunnels.
- You cannot move spoke SDP bindings if they are not manually created.

76.7.2 Steps

1

Choose Tools→Copy/Move→Service from the NFM-P main menu. The Service Copy/Move form opens.

2

Set the Action Type parameter to Move.

3

Set the Service Type parameter to a VLL service or VPLS+VLL.

Note: Only VPLS is supported for the action type Copy.

4

In the Source panel, choose the NE from which you want to move the service.

5

In the Source Port panel, choose the source port from which you want to move the service.

6

Configure required parameters in the Source panel.

7

In the Destination panel, choose the NE to which you want to move the service.

8

In the Destination Port panel, choose the port to which you want to move the service.

9

Configure the required parameters in the Execution Details panel.

The NFM-P creates SAPs on the destination port with the same encapsulation values as the source SAP if you set the Outer Encap Value Offset and Inner Encap Value Offset parameters to 0. Otherwise, the NFM-P creates SAPs on the destination port encapsulation values equal to the source encapsulation value plus the Outer Encap Value Offset or Inner Encap Value Offset parameter.

10

Click Result Export Path and specify the file name and location in which to save a text file that contains the results of service move operation.

11

Click Execute.

12

Click on the Results tab to view the results of the copy or move action.

END OF STEPS

Note:

If the new site is created successfully with all children copied from the source site, the source site is deleted.

A service move searches for all of the SAPs with specified service type on the source NE, within the encapsulation value range that you specified in Step 6 on the specified port. The NFM-P creates the associated service sites on the destination NE with the corresponding SAPs on the destination port, including all children objects, such as SDP bindings, that existed on the source site.

The source site is not deleted if the new site is created successfully with only some of the SAPs from source site, for instance, there are SAPs remaining on the source site on different ports or with outer or inner encapsulation values out of the specified range. The SAPs that are within the specified encapsulation range that you specified in Step 6 will be deleted from the source port after the move, along with all corresponding SDP bindings.

If the corresponding service site exists on the destination NE, only the SAPs within the specified encapsulation range are copied/moved to the destination site. All other children objects, such as SDP bindings, under the source site are not copied or moved to the destination site. In this case, the source site will be deleted along with all other children objects when the last SAP is moved.

If a SAP exists on the destination site with same encapsulation values as the source SAP (the source SAP encapsulation values plus the Outer Encap Value Offset or Inner Encap Value Offset parameters that you specified in Step 9), the move action fails.

If there is a failure when moving a site or any of the child objects, the move action for this site fails and nothing from this site is moved to the destination NE.

If there are multiple service sites being copied or moved along with the SAPs as specified by source port and encapsulation range, the copy or move action continues to the next site if there is an individual failure.

The NFM-P reports that the copy or move action is finished when the NFM-P database transaction has completed; deployment operations on the node may continue after the database transaction is complete.

76.8 To modify a VLL service

CAUTION

Service Disruption

Modifying parameters can be service-affecting.

76.8.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

The following tabs list the service elements that can be individually or collectively selected and configured.

Some service elements do not apply to all VLL services.

- Bandwidth tab — displays service tunnel bandwidth booking information for the service.
- Sites tab — lists the sites that are included in the service
- Endpoints tab — lists the service endpoints
- Interface tab — lists the access interfaces that are included in the service
- Spoke SDP Bindings tab — displays the spoke SDP bindings that are associated with the service
- VLAN Path Instance - list VLAN path instances associated with a Wavence service
- OAM tab
- Faults tab — displays the faults associated with the service
- Schedulers tab — configures the rate type and displays the egress aggregate rate limit

4

Modify the parameters for the service as required.

To configure items on the tabs that contain lists of service elements, select the items and click Properties.

5

Click OK. A dialog box appears.

6

Click Yes to confirm the action. The *VLL_type* Service (Edit) form closes and the Manage Services form reappears.

7

Close the Manage Services form.

END OF STEPS

76.9 To view VLL service contents

76.9.1 Purpose

Use this procedure to view various contents of a VLL L2 service and any modifications you make to it before deploying the changes.

Note: The procedure also provides information on the following policy types associated with the service:

- QoS SAP ingress and QoS SAP egress polices
- ACL IP/IPv6 filters and ACL MAC filter policies

Note: The displayed contents of a service is applicable to only one service site.

Two information views are available. The Committed info view displays the current contents of a service, and the Committed menu item is always enabled. The Modified info view allows you to review any changes you make to a service before committing them. Modified, created, and deleted attributes and objects are displayed.

76.9.2 Steps

1

Choose Manage Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search and select the service that you want to view or modify.

3

On the service navigation tree, expand the Sites icon and click on the required site. The site properties form is displayed.

4

Perform one of the following:

a. To view the currently committed service contents, go to Step 5 .

b. To view modified service contents before committing any changes, go to Step 6 .

5

Click on the Show Info and choose the Committed menu item. A Committed Values form is displayed and shows various policy attributes that have been previously configured and saved in the local policy. The CLI information of the QoS policies and ACL filters associated with the service is also displayed. The information displayed in the form is similar to the information retrieved in CLI by running the “config>service>*Service Type#* info” command, where *Service Type#* is the type of service (followed by its Service ID, for example: vll37) that you want to query.

6

Modify any service parameters or objects, as required. Otherwise go to Step 9 .

7

Click on the Show Info and choose the Modified menu item. A Modified Values form is displayed. The table lists modified, created, and deleted actions, as well as specific attributes and objects, along with their old value, new value, and tab location. The Attribute Title corresponds to the attribute or object name acted upon by your current modifications. For created objects, the values of mandatory attributes are shown in comma-separated format.

8

Select an item in the Modified Values form and then click Show on Form. The service form tab containing the changed item is displayed and the modified attribute is highlighted in blue.

9

Save your changes if required, and close the form.

END OF STEPS

76.10 To modify a VLL service using the topology view

76.10.1 Purpose

The topology view for a service provides a graphical representation of the various components and their interconnections. You can also use this view to add, modify, or just navigate to service components. This provides an alternative approach to performing these functions from the navigation tree view.

Working from the topology view can expedite the creation of the components, since many of the fields you would ordinarily have to set in the configuration forms will be automatically populated using this approach. The configuration forms can also be accessed directly at any time from this view by right-clicking on a component. This allows quick access to conduct more detailed component configuration.

76.10.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose a VLL and click Topology View. The Service Topology map opens.

The remainder of this procedure contains a number of sub-procedures describing the various components that can be created and modified from the topology view. These include:

- Creating a new site. Go to Step 4 .
- Creating site components. Go to Step 9 .
- Creating spoke SDP bindings. Go to Step 19 .

Adding a new site

4

Right-click on any blank space in the service topology map. A contextual menu is displayed. Choose the option allowing you to create a new site. This may be an Epipe, Cpipe, Ipipe, Apipe, or Fpipe site, depending on what type of VLL service you are modifying.

The Select Network Elements form appears.

5

Choose one or more sites to add to the service and click OK. The Site (Create) form for the new site is displayed. If you selected more than one site, the Site (Multiple Instances) (Create) form for the new sites is displayed.

6

Click OK. The Site (Create) (or Site (Multiple Instances) (Create)) form closes and the new site (or sites) is displayed on the map.

7

If you want to perform detailed configuration of site properties for the new site, right-click on the site icon and select Properties from the contextual menu. The Site (Edit) form opens. See 76.28 “To configure a VLL site” (p. 2154) for detailed site configuration information for the various types of VLL services.

8

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 27 to finish.

Creating site components

9

Right-click on a site icon or a SAP aggregation group icon in the service topology map. A contextual menu is displayed. Depending on the service type and the icon you clicked on, you can choose to create one of the following:

- Endpoint. Go to Step 10 .

- L2 Access Interface. Go to Step 13 .
- SAP Aggregation Group. Go to Step 16 .

10

If you choose to create an Endpoint, the Endpoint (Create) form is displayed.

11

Configure the Name parameter for the endpoint.

12

Click OK. The Endpoint (Create) form closes and the new endpoint is displayed in the topology view.

13

If you choose to create an L2 Access Interface, then the L2 Access Interface (Create) form is displayed.

14

Click on the Port tab and assign a port to the interface.

If you assign the private port of an ISA tunnel to an Epipe, an L2TPV3 tab is added. On the L2TPV3 tab, select VPRN or Base Router in the Router ID parameter and enter the Tunnel Group name.

15

Click OK. The L2 Access Interface (Create) form closes.

16

If you choose to create a SAP aggregation group, then the SAP Aggregation Group (Create) form is displayed.

17

See 76.19 “To create a SAP aggregation group on a 7705 SAR Apipe” (p. 2140) for detailed information about configuring the SAP aggregation group.

18

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 27 to finish.

Creating spoke SDP bindings

19

Choose the sites you want to connect in the service topology map and right-click on any one of them. A contextual menu is displayed.

Note: When you create a spoke binding between two sites, the order in which you select them is important. The first site you select will become the source site and the second site will become the destination site. Therefore, it is not recommended that you do a marquee-select in the topology view, since you will not be sure of this hierarchy. Instead, select the sites individually, and hold down the Shift or Ctrl key after your first selection.

20

Choose Connect from the contextual menu and choose the Create Spoke SDP Binding option.

The Spoke SDP Binding (Create) form is displayed.

Note: For this function, it is assumed that you clicked on the source site first and then held down the Shift or Ctrl key while right-clicking on the destination site to display the contextual menu.
You can also create a spoke SDP binding between a site icon and an endpoint icon, or between two endpoint icons in the topology view. Appropriate endpoints must first exist or be created to enable this.

21

Enable the Auto-Select Transport Tunnel parameter.

22

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new binding between the two sites is displayed in the topology view. See 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) for more detailed information about creating and configuring spoke SDP bindings for the various types of VLL services.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information about how to create the required tunnel. When the tunnel is created, you can repeat this sub-procedure.

23

Assuming that the spoke SDP binding was successfully created in Step 22 , select the same two sites again in the topology view, although this time do so in the opposite order that you originally selected them. This will allow you to create a spoke binding for the return tunnel.

24

Right-click on the second site you selected and choose the Create Spoke SDP Binding... option from the contextual menu. The Spoke SDP Binding (Create) form is displayed.

25

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new return binding between the two sites is displayed in the topology view.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information about how to create the required tunnel. When the tunnel is created, you can repeat this sub-procedure.

26

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 27 to finish.

27

Close the Service Topology form.

28

Close the Manage Services form.

END OF STEPS

76.11 To delete a VLL service

CAUTION

Service disruption

Deleting a service may result in a service disruption for customers.

Consider the implication of deleting the service before proceeding.

76.11.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

As required, configure the filter criteria to locate the service or range of services to be deleted. A list of services appears.

3

Choose a service or a range of services from the list.

4

Click Delete. A warning form appears. This message is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 5 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 5 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 5 .

5

For all services regardless of how their priority is configured, acknowledge the check box that prompts you confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

6

Click Yes to confirm the action. The service is deleted and removed from the list.

7

Close the Manage Services form.

END OF STEPS

76.12 To configure service tunnel required bandwidth for the service

76.12.1 Prerequisites

To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

The following prerequisites apply for use of CAC:

- An LSP must be created between the required NEs. See 31.10 “To create a static LSP” (p. 1124) or 31.11 “To create a Dynamic LSP” (p. 1126).
- If a Tunnel Selection profile is used, the Tunnel Selection profile must contain MPLS: RSVP. See 33.13 “To create a tunnel selection profile” (p. 1201).
- For CAC on a Static LSP, the LSP must be configured using the NFM-P for all hops.
  Static LSPs and hops configured in the NE do not appear complete in the NFM-P.
- For CAC on an MLPPP port, MLPPP physical links must be created. See 4.10 “To create a physical link” (p. 183).

CAUTION

Service Disruption

Bandwidth management configuration/Static LSP/Physical link is lost if the NE is unmanaged from the NFM-P.

If management is re-started, bandwidth management parameters must be configured again.

76.12.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

Click on the Bandwidth tab.

4

Configure the Bandwidth Method parameter.

5

If the Bandwidth Method parameter is set to Input Manually or Calculate from SAPs, configure LSP Path Booking and Service Path Booking parameters.

6

If the Bandwidth Method parameter is set to Input Manually, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

If the bandwidth values configured in the service and the site differ, the lower number will be used by the system.

7

Save the changes and close the forms.

END OF STEPS

76.13 To create an endpoint for a redundant VLL service

76.13.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site, right-click on Endpoints and choose Create Endpoint. The Endpoint (Create) form opens.

4

Configure the required parameters.

Enable PW Standby Signaling Master and Enable PW Standby Signaling Slave parameters are configurable on VLL Epipes and Hpipes.

Active Multipath parameter is configurable on VLL Cpipes.

5

Save the changes and close the forms.

END OF STEPS

76.14 To associate a MEP or MIP with a VLL Epipe SDP binding

76.14.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an Epipe VLL and click Properties. The Epipe Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings, right-click on the spoke SDP binding you need to configure and choose Properties. The Spoke SDP Binding (Edit) form opens.

4

Click on the OAM tab, then on the ETH-CFM tab.

5

Configure the parameters in the MIP Configuration panel.

6

Configure the parameters in the Squelch Ingress Level panel.

Levels must be assigned contiguously from Level 0. If you select a level greater than 0, then all levels lower than the one you chose will automatically be selected.

The Squelch Ingress Levels configured for all the spoke bindings on a site can be compared in one view when you click on the Spoke SDP Bindings item below a site on the navigation tree. Use the scroll bar to locate the Squelch Ingress Level columns.

7

Configure the parameters on the LMM Session Stats Collection panel as required.

8

Create one or more MEPs:

1. In the MEPs panel, Click Create. The MEP (Create) form opens.
2. Select a maintenance entity group in the Maintenance Entity Group panel.
3. Configure the required parameters in the MEPs panel.
4. Configure the required parameters in the CCM panel.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.
5. Select a MEG sub-group in the MEG Sub-Grouping panel.
6. If the MD for the MEP has a Name Type of none and the associated MEG has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable; click on the Y.1731 Tests tab and configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.
7. Click on the AIS tab and configure the required parameters.

   The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.
8. Save the changes and close the form.

9

Save the changes and close the forms.

END OF STEPS

76.15 To configure an MPLS-TP static pseudowire on a VLL spoke SDP binding

76.15.1 Prerequisites

An MPLS-TP service tunnel must be used in the SDP binding, and the Control Word parameter for pseudowire OAM must be set to Preferred. See 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) .

76.15.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an VLL service and click Properties. The VLL Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings, right-click on the spoke SDP binding you need to configure and choose Properties. The Spoke SDP Binding (Edit) form opens.

4

Click on the Control Channel tab.

5

Configure the required parameters.

6

Configure the static PW:

1. Click on the Static PW tab.
2. Click Create. The PW Path ID (Create) form opens.
3. Configure the Path AGI parameter.
4. Configure the parameters in the Source Attachment Individual Identifier panel.
5. Configure the parameters in the Target Attachment Individual Identifier panel.
6. Click OK. The PW Path ID (Create) form closes.

7

Save the changes and close the forms.

END OF STEPS

76.16 To clear BFD sessions and statistics on a VLL SDP binding

76.16.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to view the BFD session information.

4

Clear BFD sessions or statistics on the spoke SDP binding:

1. Click on the BFD tab, then on the BFD Session tab.
2. Click Clear All to clear all BFD sessions.
3. Click Clear All Statistics to clear all BFD statistics.

5

Close the forms.

END OF STEPS

76.17 To view the BFD session status on a VLL SDP binding

76.17.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and view the status of the BFD session on the spoke SDP binding:

1. Click on the BFD Session tab.
2. Choose a BFD session and click Properties. The BFD Session (View) form opens.

5

Close the forms.

END OF STEPS

76.18 To switch to the redundant port for one or more VLL SAPs

76.18.1 Prerequisites

CAUTION

Service Disruption

Nokia recommends that you set the OLC state of a service to Maintenance before you switch to a redundant port or channel, and reset it to the previous state after the switch is complete. See the NSP System Administrator Guide for more information.

Perform this procedure from a service or site properties form when port redundancy is configured on one or more a Cpipe or Epipe SAPs to:

- switch to the backup port or channel assigned to the SAP, for example, in the event of a fault on the primary port or channel
- switch back to the primary port or channel assigned to the SAP, for example, after a fault condition on the primary port or channel is cleared

Note: When multiple SAPs are selected, switching is attempted only for the SAPs that support port redundancy and on which port redundancy is properly configured.
Port redundancy switching fails if the SAP or port is included in one of the following redundancy objects:

- BGP-MH
- MC-APS
- MC-LAG
- MC-ring

76.18.2 Steps

1

Perform one of the following.

a. Choose the SAPs from a port properties form:
   1. Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.
   2. Choose Port from the Select Object Type drop-down menu.
   3. Choose the required port and then click Properties. The Physical Port (Edit) form opens.
   4. Click on the L2 Access Interfaces tab.
   5. Choose the required SAPs and then click Properties. The L2 Access Interface (Edit) form opens.
   6. Click on the Port tab.
   7. Go to Step 5 .

b. Go to Step 2 to choose the SAPs from a service or site properties form.

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Select the required VLL Cpipe or Epipe service and click Properties. The *VLL_type* Service (Edit) form opens.

4

Perform one of the following.

a. Choose the SAPs from a list of all SAPs in the service.
   1. Click on the L2 Access Interfaces tab. A list of SAPs that are configured on the service is displayed.
   2. Choose the required SAPs.

b. Choose the SAPs from a list of all of the SAPs on a site.
   1. Expand the Sites object in the service navigation tree.
   2. Click on the site object. The Site (Edit) form opens.
   3. Click on the Interfaces tab.
   4. Click on the L2 Access Interfaces tab. A list of all SAPs on the site is displayed.
   5. Choose the required SAPs.

5

Click Switch to Redundant SAP. A dialog box appears.

6

Click OK. The selected SAPs in the list are deleted, and new SAPs are created on the redundant port.

7

To confirm the redundancy status of a SAP:

1. Choose the SAP in the list and click Properties. The Port (Edit) form opens.
2. Click on the Port tab. The Manual Switch Backup State indicator displays one of the following:
   - Backup SAP in Use, if you switch from primary to backup
   - Backup Port Specified, if you switch from backup to primary

   Note:

   When the Manual Switch Backup State indicator reads Backup SAP in Use, the displayed port redundancy information is for the primary port of the SAP, and is read-only.

8

Close the forms.

END OF STEPS

76.19 To create a SAP aggregation group on a 7705 SAR Apipe

Note: SAP aggregation groups can be created only on Apipe services that are configured with the ATM VCC or ATM N:1-N>1 VC Type option. See 76.5 “To create a VLL service” (p. 2119) for more information about configuring the VC type.

A SAP aggregation group can contain up to 16 SAPs.

Alternatively, you can create SAPs for a SAP aggregation group by right-clicking on the L2 Access Interfaces entry and choosing the newly created SAP aggregation group.

The QoS and Accounting tabs are dimmed when an L2 access interface is assigned to a SAP aggregation group. The QoS and accounting policies that are assigned to the SAP aggregation group are automatically assigned to all of the SAPs in the group.

76.19.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an Apipe service with at least one 7705 SAR site and click Properties. The Apipe Service (Edit) form opens.

3

On the navigation tree, expand Site, right-click and choose Properties. The Site (Edit) form opens.

4

Click on the SAP Aggregation Group tab.

5

Click Create. The SAP Aggregation Group (Create) form opens.

6

Configure the required general parameters.

7

To configure QoS:

1. Click on the QoS tab.
2. Select an ingress policy for the SAP aggregation group in the Ingress Policy panel.
3. Select an egress policy for the SAP aggregation group in the Egress Policy panel.

8

To specify an accounting policy:

1. Click on the Accounting tab.
2. Select an accounting policy in the Accounting Policy panel.
3. Configure the Collect Accounting Statistics parameter.

9

Save the changes and close the form.

10

From the service navigation tree, expand the entries of the 7705 SAR site.

11

Right-click on the newly created SAP Aggregation Group entry and choose Create L2 Access Interface. The L2 Access Interface (Create) form opens.

12

Configure the required general parameters.

13

Click on the Port tab.

14

Select a terminating port for the L2 access interface in the Termination Port panel.

15

Configure the required parameters.

16

Click on the ATM tab and configure the required parameters.

17

Click on the QoS tab and select an ingress and egress ATM policy in the Ingress and Egress ATM Policy panels.

18

Save the changes and close the forms.

END OF STEPS

76.20 To configure EVPN on an Epipe site

76.20.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select the required Epipe and click Properties. The Epipe Service (Edit) form opens.

3

On the service tree, expand Sites and click on the site on which you want to configure a VXLAN.

4

In the General tab, configure the Source IP Address parameter in the VXLAN Tunnel End Point panel.

5

Click on the BGP tab, and then on the EVPN sub-tab.

6

In the BGP EVPN panel, click Create, or choose a BGP EVPN and click Properties. The BGP EVPN - Epipe Create/Edit form opens.

7

Configure the required parameters.

8

Configure the attachment circuits.

You can configure up to two local and two remote attachment circuits, on supporting NEs.

1. In the Local Attachment Circuit panel, click Create, or choose an item in the list and click Properties. The Local Attachment Circuit - Epipe Create/Edit form opens.
2. Configure the required parameters.
3. Save your changes and close the form.
4. In the Remote Attachment Circuit panel, click Create, or choose an item in the list and click Properties. The Remote Attachment Circuit - Epipe Create/Edit form opens.
5. Configure the required parameters.
6. Save your changes and close the form.

9

Configure the parameters in the MPLS panel, as required.

10

Configure the segment routing with IPv6.

1. In the Segment Routing V6 panel, click Create, or choose an item in the list and click Properties. The SRv6 BGP EVPN - Epipe Service Create/Edit form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

11

Configure the parameters in the Auto Bind Tunnel panel, as required.

12

Configure a network identifier for the VXLAN.

1. Click Create in the VXLAN panel. The VXLAN (Create) form opens.
2. Configure the network identifier.
3. Configure the remaining parameters as required.
4. Click OK and confirm to close the form.

13

Save and close the form.

END OF STEPS

76.21 To create an HSDPA resiliency configuration

76.21.1 Prerequisites

See 76.2.4 “HSDPA Offload Resiliency” (p. 2110) in “VLL service management overview” (p. 2100) for information about HSDPA resiliency.

See 76.5 “To create a VLL service” (p. 2119) to create the required service before configuring HSDPA Resiliency.

76.21.2 Steps

1

Choose Manage→Redundancy→HSDPA Resiliency from the NFM-P main menu. The HSDPA Resiliency Manager form opens.

2

Click Create. The HSDPA Resiliency (Create) form opens.

3

Configure the required parameters.

4

Select a site in the Site panel.

5

Select a primary service in the Primary Service panel.

6

Select a secondary service in the Secondary Service panel.

7

Save the changes and close the forms.

8

Perform one of the following:

a. Perform 76.22 “To activate and manually operate an HSDPA resiliency configuration” (p. 2145) to turn up or manually operate an HSDPA resiliency configuration.

b. Close the HSDPA Resiliency Manager form.

END OF STEPS

76.22 To activate and manually operate an HSDPA resiliency configuration

76.22.1 Prerequisites

See 76.2.4 “HSDPA Offload Resiliency” (p. 2110) for more information.

76.22.2 Steps

1

Choose Manage→Redundancy→HSDPA Resiliency from the NFM-P main menu. The HSDPA Resiliency Manager form opens.

2

Choose the required HSDPA resiliency configuration from the displayed list and click Properties. The HSDPA Resiliency (Edit) form opens with the properties of the configuration displayed on the General tab.

The Active Service field displays whether the Primary or Secondary Service is active. When you initially want to turn up a resiliency configuration, typically Primary is displayed.

3

Set the Administrative State parameter to Disabled. Activity for the resiliency configuration can only be manually switched from this form when this parameter is set to Disabled.

4

Perform one of the following:

a. Set the Administrative State parameter to Enabled. Activity automatically switches between the primary and secondary services as required when this parameter is enabled.

b. Click Turn Up to manually activate the HSDPA resiliency configuration. Do this if the Administrative State parameter is not enabled.

c. If the Primary Service is currently active, you can click Force Secondary Service Active. This manually forces the configuration to the Secondary Service (if up), regardless of state of damping timer. The Active Service field displays Secondary if the switchover is successful.

d. If the Secondary Service is currently active, you can click Force Primary Service Active. This manually forces the configuration to the Primary Service (if up), regardless of state of damping timer. The Active Service field displays Primary if the switchover is successful.

e. Click Shut Down to de-activate the HSDPA resiliency configuration.

f. Click on the Faults tab to view and address alarms related to the configuration.

5

Close the forms.

END OF STEPS

76.23 To run an OAM validation test for a VLL service

76.23.1 Prerequisites

An OAM validator test suite must be created for the tested entity. See Chapter 89, “Service Test Manager” for more information about how to create a validator test suite.

Note: As an alternative, you can also run an OAM validation test on the service by performing a One Time Validation. This is a mostly automated procedure and is described in 90.55 “To run a one-time validation test on a service” (p. 3084) .

76.23.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VLL *Type_of_VLL* (Edit) form opens.

3

Click Validate. If an OAM validator test suite is not associated to the service, a dialog box appears:

1. Click OK to associate the service with an existing OAM validator test suite.
2. Select an OAM validator test suite.

4

View the State Cause indicators. When the validation test fails, a check mark appears beside the OAM Validation Failure indicator.

5

Click on the Tests tab. The Test Suite tab is displayed.

6

Click on the Tested Entity Result tab.

7

Choose an entry and click Properties. The Tested Entity Result (Edit) form opens.

8

Click on the Results tab to display the validation test results.

9

If you need to compare two test results from the same type of test, choose the two test results and click Compare; the Difference form opens. Otherwise, go to Step 12 .

10

Compare the test results.

11

Close the Difference form.

12

Close the forms.

END OF STEPS

76.24 To create a BGP VPWS

76.24.1 Steps

1

Enable and configure BGP on the NEs involved in the BGP VPWS. See "BGP configuration workflow and procedures" (p. 914) for information about how to configure BGP.

2

Create and distribute a Service PW template policy to the NEs involved in the VPWS. See Chapter 83, "Service PW template policies" for more information about how to use Service PW template policies. You must set the VC-Type parameter to VLAN.

3

Create a routing policy statement to define the required community members. See 54.5 "To configure a routing policy statement" (p. 1745) . This defines the VSI Import/Export Routing Targets.

4

Create an Epipe service. See 76.5 "To create a VLL service" (p. 2119) for information about how to create an Epipe service.

5

Create the Epipe sites for the BGP VPWS. See 76.28 "To configure a VLL site" (p. 2154) for information about how to create an Epipe site.

Note: Epipe switching sites can have only one spoke SDP binding configured when BGP VPWS is enabled.

When BGP VPWS is enabled on an Epipe switching site, the tunnel used in the spoke SDP binding must be GRE with no signaling.
You cannot enable BGP VPWS if there are more than one spoke SDP bindings on the switching site, or if there is one or more spoke SDP bindings with a non-GRE tunnel.

6

Select one of the Epipe sites you created in Step 5 . The Site Epipe Service (Edit) form appears.

7

Click on the BGP tab.

8

Enable the Enable BGP parameter. The BGP panel is displayed.

9

Click Create. The BGP Configuration form opens.

10

Configure the Route Distinguisher parameter.

11

Click on the PW Template Binding tab and perform one of the following:

a. Select an existing PW Template Binding.

b. Add a new PW Template Binding. Go to Step 12 .

12

Create one or more PW Template Bindings:

1. Click Create. The PW Template Binding (Create) form opens.
2. Select the PW Template you created in Step 2 .
3. Select an endpoint if needed.
4. Click on the BFD tab and configure the required BFD parameters.
5. Click on the PW Template Binding Route Target tab and click Create. The PW Template Binding Route Target (Create) form opens.

   Note:

   This Import Route Target is used by the NE to decide which PW Template to use to create SDP bindings. If a far-end neighbor has a matching export target (that is, to the PW Template Import Target being defined here), then this PW Template is selected by the NE to create the pseudowire that is used to link both sites of the VPWS. If nothing is entered, and multiple PW Templates are defined, the first one found by the NE is used (most likely the one with the lowest PW Template Policy ID).

Enter the required Import Route Target in the field and click OK. The PW Template Binding Route Target (Create) form closes and the entered Import Route Target is displayed in the table on the PW Template Binding Route Target tab.

6. Click OK. The PW Template Binding (Create) form closes and the new PW Template Binding is displayed in the table on the PW Template Binding tab.
7. Select the required entry or entries from the list.

13

Click on the VSI Import Policies tab and select up to five import policies. Alternatively, you can enter the policy names manually.

14

Perform one of the following:

a. Enter the Import Route Target name manually in the provided field. The format be “target:x:y”.

b. Select an import route target. The community members you defined when creating a routing policy in Step 3 can be selected.

15

Click on the VSI Export Policies tab and select up to five export policies. Alternatively, you can enter the policy names manually.

16

Perform one of the following:

a. Enter the Export Route Target name manually in the provided field. The format must be “target:x:y”.

b. Select an export route target. The community members you defined when creating a routing policy in Step 3 can be selected.

17

Save the changes and close the BGP Configuration form.

18

Configure the parameters on the VPWS tab.

19

Create a remote VE instance, if required.

1. In the Remote VE Instances panel, click Create. The BGP VPWS Remote VE (New) form opens.
2. Configure the parameters.

20

Configure the parameters in the EVPN tab.

21

Click Apply.

22

Repeat Step 6 to Step 21 for each site in the VPWS. When the sites have been configured, endpoints and associated spoke SDP bindings are automatically created by the nodes.

END OF STEPS

76.25 To view ECMP/LAG hashing of Epipe services

76.25.1 Purpose

Perform this procedure to view the distribution of Epipe service traffic over a LAG and/or ECMP paths into the network. The Show Service Hashing menu command issues a CLI tools dump command to an NE. The CLI command can be issued from a specific LAG on an NE, or from an Epipe service site.

When you issue this command from a LAG, only the actual number of Epipe services currently using the physical port of the LAG will be displayed.

Alternatively, you can issue this command from an Epipe site that uses a port associated with a LAG. When you specify a range of Epipe Service IDs, the procedure will display information on all the Epipe services within that range that are using the same LAG as their egress port.

This procedure is only applicable to 7750 SR NEs and is not supported for PBB Epipe services.

76.25.2 Steps

1

Perform one of the following:

a. To run the Show Service Hashing command from a LAG:

1. Choose Equipment from the navigation tree view selector.
2. Navigate to the LAG icon by choosing Network→<required NE>→Logical Groups→LAGs→<required LAG>.
3. Right-click on the required LAG icon and choose Show Service Hashing from the menu.

   The Epipe LAG Show Service Hashing form opens and displays the total number of Epipe services currently using the physical port of the LAG.

   If an error is detected (for example, an invalid login), then error debugging information will be provided.

b. To run the Show Service Hashing command from an Epipe site that uses a LAG:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Configure the filter criteria. A list of services appears at the bottom of the Manage Services form.
3. Select the required Epipe service and click Properties. The Epipe Service (Edit) form opens.
4. In the navigation tree on the left side of the form, navigate to the required Epipe site.
5. Right-click on the site icon and choose Show Service Hashing from the menu. The Epipe Service Show Service Hashing form opens.
6. Enter the desired value in the End Service Id field. This is the highest Service ID in a range of up to 1024 service IDs that will be examined. It can have a value of 0 to 2147483650. The Epipe service from which you issued this command has the lowest Service ID in the range of services that will be examined. So if, for example, the Epipe has a Service ID of 100, then Service IDs up to 1124 will be examined.

   Entering a value of 0 in the End Service Id field indicates that the Service ID will not be considered. In this case, only the total number of Epipe services currently using the same physical port as the site's LAG will be displayed.
7. Click OK.

   The Epipe Service Show Service Hashing form opens and displays specific information on all Epipe services currently using the same physical port as the site's LAG.

   If an error is detected (for example, an invalid login), then error debugging information will be provided.

2

Click Close. The Show Service Hashing form closes.

END OF STEPS

76.26 To view the local PW status information for a VLL service

76.26.1 Purpose

Perform this procedure to view the state cause information for the near end of a VLL spoke SDP binding.

76.26.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose the service and click Properties. The VLL *Type_of_VLL* (Edit) form opens.

4

Click on the Spoke SDP Bindings tab.

5

Select an entry and click Properties. The Spoke SDP Binding (Edit) form opens.

6

Click on the States tab.

7

View the information in the Local State Cause panel. When there is a local fault, a check mark beside the appropriate indicator identifies the fault type.

8

Close the Spoke SDP Binding (Edit) form.

9

Close the VLL *Type_of_VLL* (Edit) form.

END OF STEPS

76.27 To view the peer PW status information for a VLL service

76.27.1 Purpose

Perform this procedure to view the state cause information for the far end of a VLL spoke SDP binding.

76.27.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose the service and click Properties. The VLL *Type_of_VLL* (Edit) form opens.

4

Click on the Spoke SDP Bindings tab.

5

Choose an entry and click Properties. The Spoke SDP Binding (Edit) form opens.

6

Click on the States tab.

7

View the peer status information in the Peer State Cause panel. When the service tunnel or peer SAP is down or partially down, a check mark beside the appropriate indicator identifies the fault type.

8

Close the Spoke SDP Binding (Edit) form.

9

Close the VLL *Type_of_VLL* (Edit) form.

END OF STEPS

VLL site management procedures

76.28 To configure a VLL site

76.28.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, right-click on *VLL_type* Service and choose Create *VLL_type* Site and select an NE, or right-click on a site and choose Properties. The *VLL_type* Site (Create|Edit) form opens.

4

Configure the required general parameters.

The Enable CE IP Address Discovery parameter is configurable only for VLL Ipipe terminating sites.

On Ipipe sites, the Stack Capability Signaling parameter is configurable only when the Enable IPv6 parameter and the Enable CE IP Address Discover parameters are enabled.

The VLL Site Type parameter is displayed as read-only for the 7950 XRS, as only switching VLL sites are supported.

You must set the VLL Site Type parameter to Terminating if the site is connected to a 7210 SAS-D, 7210 SAS-E, 7210 SAS-K, 7210 SAS-M, 7210 SAS-Mxp, 7210 SAS-S, 7210 SAS-Sx, or 7210 SAS-T.

On 7210 SAS sites:

- the SAP Type parameter is configurable when the VLL Site Type parameter is set to Terminating
- when you configure a 3-plus-tag Epipe service, you must set the SAP Type parameter to qinq-inner Tag-preserve. See “Epipe (Ethernet VLL)” (p. 2100) in “VLL service management overview” (p. 2100) . See 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) and 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for additional configuration requirements for a 3-plus-tag Epipe service.
- when you configure a VLAN range SAP on the site, you must set the SAP Type parameter to dot1q-range.

On 7210 SAS-X sites, the PBB Epipe parameter is configurable when the VLL Site Type parameter is set to Terminating.

5

To assign test generation options to the site:

1. Click on the OAM tab, then the Configuration tab.
2. Configure the required Test Generation Options parameters.

6

Save the changes and close the forms.

END OF STEPS

76.29 To configure a GNE site on a VLL service

76.29.1 Purpose

Use this procedure to add a GNE site and GNE service interfaces to an existing VLL service. This procedure applies to VLL Epipe, Apipe, Fpipe, Ipipe, and Cpipe services. This procedure also applies to Wavence VLL Epipe, Apipe, and Cpipe services.

76.29.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the component tree, right-click on *VLL_type* Service and choose Create GNE Site. The Select Network Elements - *VLL_type* Service form opens with a list of available NEs.

4

Choose a site and click OK. The GNE Site (Create) form opens.

5

Configure the required parameters.

6

Click Apply.

7

To configure an interface for the GNE site:

1. Click on the GNE Service Interfaces tab and then Click Create. The GNE Service Interface (Create) form opens.
2. Configure the parameters.
   - Name
   - Description
   - Type
3. Click on the Ports tab and then click Select beside the Interface Index field. The Select Generic NE Interface form appears.
4. Select an interface from the list and click OK. The GNE Service Interface (Create) form reappears with the interface information displayed.
5. Configure the parameters.
   - Encapsulation Type
   - Outer Encapsulation Value
   - Inner Encapsulation Value
6. Click OK. The GNE Site (Create) form reappears with the interface information displayed in the service component tree.

8

Save the changes and close the forms.

END OF STEPS

76.30 To configure service tunnel required bandwidth for the site

76.30.1 Prerequisites

To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

The following prerequisites apply for use of CAC:

- An LSP must be created between the required NEs. See 31.10 “To create a static LSP” (p. 1124) or 31.11 “To create a Dynamic LSP” (p. 1126).
- If a Tunnel Selection profile is used, the Tunnel Selection profile must contain MPLS: RSVP. See 33.13 “To create a tunnel selection profile” (p. 1201).
- For CAC on a Static LSP, the LSP must be configured using the NFM-P for all hops.
  Static LSPs and hops configured in the NE do not appear complete in the NFM-P.
- For CAC on an MLPPP port, MLPPP physical links must be created. See 4.10 “To create a physical link” (p. 183).

CAUTION

Service Disruption

Bandwidth management configuration/Static LSP/Physical link is lost if the NE is unmanaged from the NFM-P.

If management is re-started, bandwidth management parameters must be configured again.

76.30.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Sites, right-click the site on which you want configure bandwidth, and choose Properties.The VLL Site (Edit) form opens.

4

Click on the Bandwidth tab.

5

Enable the Override Service Configuration parameter.

6

Configure the Bandwidth Method parameter.

7

If the Bandwidth Method parameter is set to Input Manually, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

If the bandwidth values configured in the service and the site differ, the lower number will be used by the system.

8

Save the changes and close the forms.

END OF STEPS

76.31 To link an Epipe service to a backbone VPLS site

76.31.1 Prerequisites

For additional information on Provider Backbone Bridging, see Chapter 77, “VPLS management ” .

76.31.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

3

On the navigation tree, expand the Sites icon, right-click the site on which you want to link the VPLS site and choose Properties.The Epipe Site (Edit) form opens.

4

Click on the Backbone tab.

On the 7210 SAS-M, 7210 SAS-T, and 7210 SAS-X, the Backbone tab is configurable only if you enabled the PBB Epipe parameter in 76.28 “To configure a VLL site” (p. 2154) .

5

Select a VPLS in the Backbone VPLS Site panel.

Note: The B-Sites listed are only those that are on the same NE type as the Epipe site. See 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) for more information.
The selection of a B-Site in this step must be repeated for both Epipe sites. You must select a B-Site that is on the same NE as the Epipe site.

6

In the PBB panel, configure the required parameters.

The ISID parameter is configurable when the Inherit Service ID Value parameter is disabled.

Note: The ISID parameter should be set to the same value for both Epipe sites you create for this service.
To select a previously created Mac Name, click Select and then select a MAC Name from the list in the Select MAC Destination MAC Address Alias form. See 12.31 “To create a chassis-level PBB configuration” (p. 367) for more information about how to create a MAC Name.
The Destination MAC Address or MAC Name parameter you set for this Epipe site should be the same as the Source MAC Address of the B-Site on the other node of the Epipe service you are creating.

7

Configure the Force Q Tag Forwarding parameter.

8

Configure the Local Switching Service State parameter.

9

Save the changes and close the forms.

END OF STEPS

76.32 To associate a Facility MEP with a VLL Epipe site

76.32.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an Epipe VLL service and click Properties. The Epipe Service (Edit) form opens.

3

On the navigation tree, expand Sites, right-click the site and choose Properties. The Epipe Site (Edit) form opens.

4

Click on the OAM tab, and then on the ETH-CFM tab.

5

Configure the Tunnel Fault Notification parameter in the Facility MEPs panel.

The Tunnel Fault Notification parameter is configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation. If you are configuring a tunnel facility MEP, this parameter must be set to Accept, in order to accept the fault notification from the tunnel facility MEP.

6

Save the changes and close the forms.

END OF STEPS

76.33 To configure segment routing with IPv6 on a VLL Epipe site

76.33.1 Before you begin

The segment routing with IPv6 can be configured on 7450 ESS, 7750 SR, 7950 XRS, and IXR NEs. This configuration applies to SR and IXR from NFM-P Release 23.11 and node releases 22.2 and 22.10 onwards.

76.33.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an Epipe VLL service and click Properties. The Epipe Service (Edit) form opens.

3

On the navigation tree, expand Sites, right-click the site and choose Properties. The Epipe Site (Edit) form opens.

4

Click on the Segment Routing V6 tab.

5

Configure the required parameters in the General tab.

6

Create or select an entry in the Locator panel and click Properties. The SRv6 Function (Create/Edit) form opens.

1. Associate a locator under Locator panel.
2. Configure the parameters under End Function DX2 panel.
3. Save your changes and close the form.

7

Save your changes and close the forms.

END OF STEPS

76.34 To configure a spoke SDP binding on a VLL site

76.34.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings, right-click and choose Create SDP Spoke Binding, or expand Spoke SDP Bindings, right-click on a circuit and choose Properties. The Spoke SDP Binding (Create|Edit) form opens.

4

Specify a destination NE for the spoke SDP binding:

a. If the tunnel termination site is a managed NE, select an NE from a list of managed NEs.

b. If the tunnel termination site is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

5

Configure the required general parameters.

The VC ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The Block On Peer Fault parameter takes effect only when PW status signaling is enabled.

The Block On Peer Fault parameter is not configurable on a spoke SDP that is on an MC LAG or an endpoint.

You must configure the VLAN VC Tag parameter with the same value as the inner tag of the QinQ SAP on the service. See “Epipe (Ethernet VLL)” (p. 2100) in “VLL service management overview” (p. 2100) for more information. See 76.28 “To configure a VLL site” (p. 2154) and 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for additional configuration requirements for a 3-plus-tag Epipe service.

The SDP Admin Bandwidth parameter is not configurable if you select an MPLS-TP tunnel in Step 6 .

The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

The Override Service MTU parameter is only configurable on SDP bindings in a Epipe service. Configuring this parameter to any value other than -1 uses that value for the MTU size, instead of the value configured at the service level. When the value is -1, the service-level MTU is used.

6

To specify a transport tunnel for the spoke SDP binding:

a. Configure the transport tunnel manually by selecting a service tunnel for the SDP binding in the Tunnel panel.

Note: The NFM-P blocks the manual selection of a tunnel with insufficient bandwidth if the server configuration includes the following parameter setting:
<vll-CAC enforceTunnelBandwidth="true"/>
The parameter is set to false by default, which means that the NFM-P does not block the manual selection of a tunnel with insufficient bandwidth; instead, the “Insufficient Bandwidth To Allocate To SDP Binding” State Cause indicator is set.
To change the parameter setting, contact Nokia technical support.

b. For Apipe and Cpipe VLLs, configure an MPLS-TP transport tunnel manually by selecting a tunnel in the Tunnel panel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for more information about how to create an MPLS-TP service tunnel.

c. Let the NFM-P configure the transport tunnel automatically.

1. Enable the Auto-Select Transport Tunnel parameter.
2. Configure either the Profile Name or the Tunnel Auto-Selection Transport Preference parameter.

7

To create redundant SDP bindings, configure the endpoint in the Redundancy panel.

1. Select an endpoint in the Redundancy panel.
2. Configure the parameters.

   When the Precedence parameter is set to 0, you can configure a secondary SAP on an ATM interface, MC-APS, or APS for HSPDA offload fallback.

8

To specify a transport tunnel for the Return SDP binding:

Note: You can create a return tunnel only between sites that are in the same service. If the sites are not in the same service, the Return tab does not appear.
When the Auto Select Return Transport Tunnel parameter is enabled, NFM-P will consider and use a portion of the bandwidth you specified when setting the SDP Admin Bandwidth parameter in Step 5 .

a. Allow the NFM-P to configure the transport tunnel automatically by clicking on the Return tab, enabling the Auto Select Return Transport Tunnel parameter and selecting a tunnel selection profile beside the Profile name parameter. See 33.13 “To create a tunnel selection profile” (p. 1201) for more information about how to create a service tunnel profile.

b. Allow the NFM-P to configure the transport tunnel automatically by clicking on the Return tab, and enabling the Auto Select Return Transport Tunnel and Return Tunnel Auto-Selection Transport Preference parameters.

c. Configure the transport tunnel manually, by clicking on the Return tab, selecting a return tunnel in the Return Tunnel panel.

9

Select a return endpoint on the terminating site in the Return SDP Binding Endpoint panel.

10

Configure the parameters in the Hash Label panel, if supported.

11

Configure the Force VLAN VC Forwarding parameter in the VLAN panel.

The Force VLAN VC Forwarding parameter is configurable only on Epipe SDP bindings.

You cannot enable the Force VLAN VC Forwarding parameter if you enable the Force Q-in-Q VC Forwarding parameter in Step 12 .

12

Configure the Force Q-in-Q VC Forwarding parameter in the Q in Q panel.

The Force Q-in-Q VC Forwarding parameter is configurable only on Epipe SDP bindings.

You cannot enable the Force Q-in-Q VC Forwarding parameter if you enable the Force VLAN VC Forwarding parameter in Step 11 .

13

To configure QoS:

1. Click on the QoS tab.
2. Select a network policy and configure the required parameters in the Forwarding Plane Redirect panel.
3. Select an ingress queue group template policy in the Forwarding Plane Redirect panel.
4. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
5. Select a network policy and configure the required parameters in the Port Redirect panel.
6. Select an egress queue group template policy in the Port Redirect panel.
7. Configure the Instance ID parameter in the Port Redirect panel.

14

Click on the States tab and configure the Administrative State parameter.

15

To configure BFD:

1. Click on the BFD tab.
2. Enable the Enable BFD parameter.

3. Select a BFD template in the BFD Template panel.

16

To configure pseudowire OAM:

1. Click on the Pseudowire OAM tab.
2. Configure the Control Word parameter.

   If you are creating a spoke SDP binding using an MPLS-TP service tunnel for pseudowire static configuration, you must set the Control Word parameter to Preferred.

   For Apipes, if you set the VC Type parameter to ATM-SDU in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

   For Fpipes, if you set the VC Type parameter to FR-DLCI in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

   For Hpipes, if you set the VC Type parameter to HDLC in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

17

To assign ingress and egress filter policies to the spoke SDP binding:

1. Click on the ACL tab.
2. Select an ingress filter in the Ingress Filter panel.
3. Select an egress filter in the Egress Filter panel.
4. Select an IPv6 ingress filter in the IPv6 Ingress Filter panel.

18

To assign an accounting policy to the SDP binding:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy in the Accounting policy panel.

19

To configure security:

1. Click on the Security tab.
2. Enable the IP Src Monitoring parameter.

20

To assign an Application Profile to the spoke SDP binding:

1. Click on the Application Assurance tab.
2. Select an application profile string.

You can select only local profiles on the NE.

3. To associate a transit IP policy with an Epipe service object, select a transit IP policy in the Transit IP Policy panel.

   See 87.46 “To configure an AA transit IP policy” (p. 2860) for more information about how to configure a transit IP policy.

4. To associate an AA transit prefix policy with the service object, select a transit prefix policy in the Transit Prefix Policy panel.

   See 87.47 “To configure an AA transit prefix policy” (p. 2861) for more information about how to configure a transit prefix policy.

5. To associate an AA redundant protocol with an Epipe service object, select an AA redundant protocol in the AA Redundant Protocol panel.

21

Perform one of the following:

a. If you are configuring a VLL with spoke switching, you must configure a switching site and two or more terminating sites. Repeat 76.28 “To configure a VLL site” (p. 2154) to configure more sites, and repeat Step 1 through Step 22 of this procedure.

b. If you are configuring a PBB Epipe, you can only configure one Spoke SDP. Go to Step 22 .

22

Save the changes and close the forms.

END OF STEPS

76.35 To configure a spoke SDP binding with an L2TPv3 tunnel on a VLL Epipe site

76.35.1 Prerequisites

See 33.11 “To configure an L2TPv3 service tunnel” (p. 1199) for more information about how to create an L2TPv3 tunnel.

Note: You can configure only one spoke SDP binding on each site.

76.35.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings, right-click and choose Create SDP Spoke Binding, or expand Spoke SDP Bindings, right-click on a spoke SDP binding and choose Properties. The Spoke SDP Binding (Create|Edit) form opens.

4

Specify a destination NE for the spoke SDP binding:

a. If the destination NE is a managed NE, select an NE from a list of managed NEs.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

5

Configure the required general parameters.

The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

6

Configure the L2TPv3 transport tunnel.

1. Select an L2TPv3 transport tunnel in the Tunnel panel. See 33.11 “To configure an L2TPv3 service tunnel” (p. 1199) for more information about how to create an L2TPv3 tunnel.
2. Configure the parameters in the L2TPv3 Cookie panel.

7

To configure QoS:

1. Click on the QoS tab.
2. Select a network policy and configure the required parameters in the Forwarding Plane Redirect panel.
3. Select an ingress queue group template policy in the Forwarding Plane Redirect panel.
4. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
5. Select a network policy and configure the required parameters in the Port Redirect panel.
6. Select an egress queue group template policy in the Port Redirect panel.
7. Configure the Instance ID parameter in the Port Redirect panel.

8

Click on the States tab and configure the Administrative State parameter.

9

To configure pseudowire OAM:

1. Click on the Pseudowire OAM tab.

2. Configure the Control Word parameter.

   If you are creating a spoke SDP binding using an MPLS-TP service tunnel for pseudowire static configuration, you must set the Control Word parameter to Preferred.

   For Apipes, if you set the VC Type parameter to ATM-SDU in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

   For Fpipes, if you set the VC Type parameter to FR-DLCI in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

   For Hpipes, if you set the VC Type parameter to HDLC in Step 3 of 76.5 “To create a VLL service” (p. 2119) , you must set the Control Word parameter to Preferred.

10

To assign ingress and egress filter policies to the spoke SDP binding:

1. Click on the ACL tab.
2. Select an ingress filter in the Ingress Filter panel.
3. Select an egress filter in the Egress Filter panel.
4. Select an IPv6 ingress filter in the IPv6 Ingress Filter panel.

11

To assign an accounting policy to the SDP binding:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy in the Accounting policy panel.

12

To configure security:

1. Click on the Security tab.
2. Enable the IP Src Monitoring parameter.

13

To assign an Application Profile to the spoke SDP binding:

1. Click on the Application Assurance tab.
2. Select an application profile string.

   You can select only local profiles on the NE.
3. To associate an AA transit prefix policy with the service object, select a transit prefix policy in the Transit Prefix Policy panel.

   See 87.47 “To configure an AA transit prefix policy” (p. 2861) for more information about how to configure a transit prefix policy.

14

Save the changes and close the forms.

END OF STEPS

76.36 To create a spoke SDP FEC binding on a VLL Epipe site

76.36.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP FECs, right-click and choose Create Spoke SDP FEC. The Spoke SDP FEC (Create) form opens.

4

Configure the required general parameters.

5

In the Source panel, select an SAII address.

6

In the Target panel, configure the Auto Config parameter or select a TAII address.

7

In the General Parameters panel:

1. Configure the parameters.
2. Select a path.

8

In the Redundancy panel:

1. Select an endpoint.
2. Configure the parameters.

9

In the Pseudowire Signaling panel, configure the Enable PW Standby Signaling Slave parameter.

10

Save the changes and close the forms.

END OF STEPS

76.37 To configure an Epipe site for BGP multi-homing

76.37.1 Prerequisites

Configure an Epipe site for BGP multi-homing to allow for multi-homed connectivity to G.8031 Ethernet tunnels between two CEs. G.8031 resiliency is used on the edges and then pseudowires are allowed to reach the far end, where G.8031 can again be used in a multi-chassis mode.

G.8031 is created on two CEs which are multi-homed to two PEs. Using BGP and CFM, PE routers enable or disable the CE-PE SAPs. The core is IP/MPLS and Epipe is used for transport between the PEs. The Ethernet tunnels are created on the CEs and the Epipe service on the CEs uses these tunnels.

76.37.2 Steps

1

Before you configure a site for BGP multi-homing, you must either:

a. Enable BGP on the routing instance of each NE in the Epipe BGP multi-homing configuration. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.

b. Configure global-level BGP on each NE in the Epipe BGP multi-homing configuration. See 28.31 “To configure global-level BGP” (p. 918) for more information. The following items are required:

- Enable the L2 VPN parameter on:
  - the VPN tab of the BGP site form
  - the VPN tab of the BGP Peer Group form, which contains the peers that are involved in the BGP multi-homing
- Create a peer group under BGP. This peer group is used to collectively define the peers involved in the Epipe BGP multi-homing configuration.
- Create the required peers under the peer group. These peers are the NEs involved in the Epipe BGP multi-homing configuration.

Note: For optimal processing while a BGP multi-homing site is activated or de-activated, or the system is rebooted, you should also:

- Enable the L2 VPN parameter in the Rapid Update Address Family panel on the BGP site VPN tab.

- Enable the Enable Rapid Withdrawal parameter on the BGP site Behavior tab.

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

4

On the service navigation tree, right-click on Epipe Service and choose Create Epipe Site, or right-click on a site and choose Properties. The Epipe Site (Create|Edit) form opens.

Note: You can see a list of all current Epipe multi-homing service sites by viewing the BGP Multi-homing tab on the service configuration form.

5

Configure the general BGP parameters:

1. Click on the BGP tab, then on the General tab, and enable the Enable BGP parameter.
2. Configure the Route Distinguisher parameter in the RD panel. Enabling the Show Format checkbox displays a list of acceptable formats and value ranges.

   Note:

   An RD or RT configured under the BGP of an Epipe site cannot be removed as long as there is a multi-homing site configured that uses these and whose administration state is up.
3. Configure the Import Route Target and Export Route Target parameters in the RT panel. Enabling the Show Format checkbox displays a list of acceptable formats and value ranges.

   You can also configure the Export Route Target to be the same as the Import Route Target by enabling the Same as Import checkbox.

6

To configure multi-homing:

1. Click on the Multi-homing tab.
2. Click Create. The BGP Multi-homing Site (Create) form opens.
3. Configure the required parameters.

   You cannot turn up a BGP multi-homing site without specifying a Multi-homing ID. In addition, the Multi-homing ID of the site can only be removed at a later time when the site is shutdown.
4. Select a SAP for the BGP multi-homing in the Enable BGP Multi-homing to SAP panel.

You cannot turn up a BGP multi-homing site without specifying a SAP. The SAP configured under this site can be part of an endpoint, but the only other object allowed in the endpoint is an ICB spoke-SDP. In addition, an assigned SAP can only be removed at a later time when the BGP multi-homing site is shutdown.

5. Configure the parameters in the Timer panel.

   If you enable either the Use Node Level Boot Timer and/or Use Node Level Site Activation Timer parameters, then the associated Boot Timer (seconds) and/or Activation Timer (seconds) parameters are not configurable. These parameter values are inherited from the network element configuration.

   The Boot Timer (seconds) and Activation Timer (seconds) parameters can be configured for an NE on the BGP Multi-homing sub-tab under the Redundancy tab in the Network Element (Edit) form. See 12.5 “To modify NE properties” (p. 343) for more information about how to change device properties.

   If you enable the Use Node Level Site-Down Minimum Timer parameter, the Site-Down Minimum Timer (seconds) parameter and Site-Down Minimum Timer Remaining (seconds) parameters are ignored and the values configured on the NE are used.

6. Configure the Administrative State parameter in the States panel.
7. Save the changes and close the form.

7

Save the changes and close the forms.

END OF STEPS

76.38 To enable the automatic selection of an RD on a VLL Epipe site

76.38.1 Prerequisites

Since an RD must be unique on each PE in the network, you can allocate either a route distinguisher that you manually select or an NE-selected route distinguisher for each service. When you configure an auto-RD on a VLL Epipe site, a Type-1 RD is automatically allocated by the NE based on the community range that you configure.

76.38.2 Steps

1

Before you configure a site for auto-RD selection, you must:

1. Enable BGP on the routing instance of the NE. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.
2. Enable the automatic selection of an RD on the NE and specify the community range. See 12.11 “To enable the automatic selection of an RD on an NE” (p. 349) .

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

4

On the navigation tree, right-click on Epipe Service and choose Create Epipe Site, or right-click on a site and choose Properties. The Epipe Site (Create|Edit) form opens.

5

Click on the BGP tab, then on the General tab, and enable the Enable BGP parameter.

6

Configure the Auto Route Distinguisher parameter in the RD panel.

The Operational RD is displayed after you apply the changes to the site and to the service.

7

Save the changes and close the forms.

END OF STEPS

76.39 To view the last cleared BFD statistics and sessions on a VLL site

76.39.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, click on the site on which you want to view the last cleared BFD statistics and sessions.

4

Click on the Clear Status tab and view the five last cleared BFD statistics and the five last cleared BFD sessions.

5

Close the forms.

END OF STEPS

VLL access interface management procedures

76.40 To create a VLL L2 access interface on a terminating site

76.40.1 Prerequisites

If you are configuring access dual-homing with local switching over PBB tunnels, you must configure two L2 access interfaces. The L2 access interfaces can be on LAGs that participate in the MC LAG, or they can be just regular or Ethernet tunnel access interfaces.

If the Use SAP Backbone MAC Address parameter is enabled on both B-Sites:

- MC-LAG access interfaces for the PBB Epipe local switching must be on IOM3 MDAs for the feature to work properly. There is no NFM-P or CLI validation.
- Other access interfaces may be on IOM1 or IOM2 MDAs.

If you configure an L2 access interface on a PBB Epipe site you cannot configure a spoke SDP binding on the same site. They are mutually exclusive.

When you configure a 3-plus-tag Epipe service, only ports with QinQ encapsulation as one endpoint can be selected. When the other endpoint is an L2 access interface, you must choose a port with dot1q or QinQ encapsulation as the other L2 access interface. Both endpoints must be on the same site. See “Epipe (Ethernet VLL)” (p. 2100) in “VLL service management overview” (p. 2100) for more information. See 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) and 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) for additional configuration requirements for a 3-plus-tag Epipe service.

The parameters displayed may vary depending on the NE type or release.

76.40.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site, right-click on L2 Access Interfaces and choose Create L2 Access Interface. The L2 Access Interface (Create) form opens.

For information on adding a GNE service interface to a VLL service, see 76.29 “To configure a GNE site on a VLL service” (p. 2155) .

4

Configure the required parameters.

The ATM Connection Type parameter is configurable only in a VLL Apipe service that has the VC Type parameter set to ATM-cell.

The MC Ring Node and Hold Service Up On SAP Failure parameters are configurable only in a VLL Epipe service.

The Tunnel Fault Notification parameter is configurable on L2 access interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation. If you are configuring a tunnel facility MEP, this parameter must be set to Accept, in order to receive the fault notification from the tunnel facility MEP.

The Enable AIS parameter is configurable only in a VLL Epipe service where the device has ports configured in access or hybrid mode with QinQ encapsulation to forward AIS frames generated by a tunnel facility MEP.

5

If you are creating a redundant L2 access interface, select an endpoint in the Redundancy panel.

6

CAUTION

Service Disruption

The creation of a SAP that uses the same port and encapsulation values as an existing inactive MSAP fails under the following conditions:

- If you try to use the NFM-P to create a SAP, the configuration fails and the NFM-P displays an error message.
- If you use a CLI to create a SAP in a service other than the service that contains the MSAP, the configuration succeeds but the MSAP is inactive until the regular SAP is deleted. Although the NFM-P displays the SAP and MSAP, the MSAP remains inactive and consumes resources.
- If you use a CLI to create a SAP in the service that contains the MSAP, the SAP creation fails.

Nokia recommends that you delete an inactive MSAP from the NFM-P if you need to create a regular SAP on the same port using the same encapsulation values.

Configure the port:

1. Click on the Port tab.
2. Select a port in the Terminating Port panel.

   Only ports in access or hybrid mode are listed. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. The port is then listed when you click Search.

On Epipe L2 access interfaces, only PW ports that are bound to a service tunnel are listed. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for information about how to create a service tunnel with a PW port binding.

On Epipe L2 access interfaces, SAP types X.0 and X.* as well as SAP types *.null and *.* can be configured on the same Q in Q port. You must enable the Q in Q Untagged Sap parameter on the NE which allows the creation of the following default SAP types:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- The SAP type *.* functions as a default SAP for double-tagged frames on a Q in Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

See 12.23 “To enable a Q in Q untagged SAP on an NE” (p. 360) for more information.

When you configure a 3-plus-tag Epipe service, you must choose a port with QinQ encapsulation as one endpoint (SAP). When the other endpoint is an L2 access interface, you must choose a port with dot1q or QinQ encapsulation as the other L2 access interface. Both endpoints must be on the same site. See “Epipe (Ethernet VLL)” (p. 2100) in “VLL service management overview” (p. 2100) for more information. See 76.28 “To configure a VLL site” (p. 2154) and 76.34 “To configure a spoke SDP binding on a VLL site” (p. 2161) for additional configuration requirements for a 3-plus-tag Epipe service.

If you are configuring a port as a VLAN range SAP, you must choose a port with dot1q encapsulation type.

If you choose an Ethernet tunnel endpoint, the form is refreshed and an Ethernet Tunnel tab is added.

If you choose a port that is associated with a 7705 SAR MW link, the form is refreshed and a Microwave tab appears. See 15.38 “To configure a 7705 SAR MW link” (p. 499) and 15.39 “To configure a 7705 SAR MW link member” (p. 501) for more information about configuring a 7705 SAR microwave link and MW link member. If you are configuring microwave compression, see 76.51 “To configure microwave compression on an MW link SAP on a VLL L2 access interface” (p. 2195) .

If you are configuring a 7705 SAR-8 or 7705 SAR-18 that has SCADA configured, you can add SCADA branches as SAPs to an Epipe or Cpipe service site. See 15.36 “To configure SCADA on a 7705 SAR” (p. 495) for more information about configuring SCADA on a 7705 SAR.

3. Configure the required parameters.

   When the selected port uses dot1q encapsulation, you can enable the Auto-Assign ID check box to have the Outer Encapsulation Value parameter automatically assigned. If you choose this, the system assigns the lowest unused encapsulation value.

   Note:

   You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter in the User Preferences form. Choose Application→ User Preferences from the main menu.

The Inner Encapsulation Value parameter is configurable only when the port encapsulation type is QinQ.

When you configure a 3-plus-tag Epipe service with L2 access interfaces as both endpoints, you must observe the following requirements:

- One endpoint must be a QinQ SAP
- The other endpoint must be a dot1q SAP or a QinQ SAP.
- When the other endpoint is a dot1q SAP, the outer encapsulation value for the dot1q SAP must match the inner encapsulation value of the QinQ SAP that shares the L2 connection.
- When the other endpoint is a QinQ SAP, the inner encapsulation values for both QinQ SAPs on the service must match.

See "Epipe (Ethernet VLL)" (p. 2100) in this chapter for more information. See 76.28 "To configure a VLL site" (p. 2154) and 76.34 "To configure a spoke SDP binding on a VLL site" (p. 2161) for additional configuration requirements for a 3-plus-tag Epipe service

The Outer Encapsulation Value (VPI) parameter is configurable only when the port encapsulation type is ATM.

The Inner Encapsulation Value (VCI) parameter is configurable only when the port encapsulation type is ATM.

The LLF Enabled parameter is configurable only for Epipe L2 access interfaces on ports with Null encapsulation.

If the port you have chosen is an Ethernet Tunnel Endpoint, you will be able to set the Outer Encapsulation Value to 8191. This automatically enables the Ethernet Tunnel Endpoint Control SAP parameter.

For Apipes, the Encapsulation Value parameter and the Encapsulation Value parameter are configurable only when the VC Type parameter is set to ATM-cell or ATM-VPC in Step 3 of 76.5 "To create a VLL service" (p. 2119) and the ATM Connection Type parameter is set to PVT in Step 4 of this procedure.

For Apipes, the LLF Enabled parameter is configurable only for SAPs with "Port" ATM Connection Type and on a clear channel under 4 port OC3-STM1 ASAP MDA.

For Apipes, the Connection Profile ID parameter is configurable only when the Signalling VC Type Override parameter is set to ATM-VCC or None. The Connection Profile ID parameter can only be associated with an Apipe when the VC Type parameter is set to ATM-cell.

4. If you are creating an Epipe and use dot1q or QinQ encapsulation, you can enable ingress VLAN translation. Configure the required parameters in the VLAN Translation panel.
5. For QTag Manipulation on the 7250 IXR, configure the parameters in the QTag Manipulation panel.

   For the Ingress Action parameter, when you select the Push Outer or Replace Outer option, the Outer Tag parameter is available. You must configure the Outer Tag parameter.

   If required, click Reset to restore the default values.
6. For VLAN ranges on 7210 SAS Epipes, enable the VLAN Range SAP parameter and select a VLAN range.

   You can configure the VLAN Range SAP parameter only when the SAP Type parameter is set to dot1q-range. See Step 4 in 76.28 "To configure a VLL site" (p. 2154) .

See Chapter 63, “Connection profile policies” for more information about VLAN ranges for the 7210 SAS. See 63.4 “To configure a VLAN range for a 7210 SAS VPLS or VLL Epipe service” (p. 1825) for information about how to create a connection profile with a VLAN range.

7

If a VLAN Connection Profile policy was created and you want to use it for an Epipe service, configure the parameters in the VLAN Connection Profile panel on the Port tab. See Chapter 62, “VLAN Connection Profile policies” for details.

8

To configure port redundancy for a Cpipe or Epipe service:

Port redundancy is configurable only on a 7705 SAR or 7750 SR.

1. Enable the Enable Port Redundancy parameter on the Port tab.
2. Configure the Off-Node Protection parameter in the Backup Port Information panel.
3. If you enable the Off-Node Protection parameter, select the site that contains the backup port.

   Note:

   The site that you choose must be the same type and release of device as the site that you are configuring.
4. Select the redundant port.

   Note:

   Only ports with the same encapsulation type as the currently assigned port are listed.
5. Configure the required parameters.
6. If you are creating a Cpipe L2 access interface, save the changes and close the form.

9

Configure the Ethernet Tunnel Endpoint Control SAP parameter.

Note: Enabling the Ethernet Tunnel Endpoint Control SAP parameter creates the control L2 Access Interface (also known as a Control SAP). It also automatically sets the value of the Outer Encapsulation Value parameter to 8191.
If you are currently creating a same-fate SAP, the Ethernet Tunnel Endpoint Control SAP parameter must not be enabled.

10

To configure policy overrides for Epipe or Ipipe L2 access interfaces:

Note: The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.

- If you configured a non-HSMDA port, the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.

- If you configured an HSMDA port, the Access Ingress Queues, Access Egress HSMDA Queues, and Ingress Policer tabs are active.

1. Click on the Override Policy Items tab.
2. Set the policy overrides, as described in 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) .

   To configure meter overrides on a 7210 SAS, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) . To configure queue overrides on a 7210 SAS, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

11

To specify Ipipe functionality for Ipipe L2 access interfaces:

1. Click on the IPIPE tab.
2. Configure the required parameters.

   The MAC Address parameter, MAC Refresh Interval parameter, and Use Broadcast MAC Address parameters are configurable only when the port encapsulation type is dot1q, QinQ, or Null.

12

To specify Epipe functionality for Epipe L2 access interfaces:

1. Click on the Epipe tab.
2. Configure the required parameters.
3. Click on the following tabs to view or edit information:
   - General tab—displays the general Epipe properties
   - Port
   - OAM tab—create and execute service-specific diagnostic tests
   - QoS—configure the required parameters and select an Egress Remark policy, Egress VLAN QoS policy, and a Shared Policer policy, as required
   - Schedulers tab
   - ACL
   - Accounting
   - Virtual Port Name—configure the required parameters to assign a virtual port to the interface.
   - Security—configure the required parameters and select a DDoS protection policy
   - TCA
   - Deployment
   - Faults—displays the faults associated with the Epipe L2 access interface

13

To assign an Application Profile to the spoke SDP binding:

1. Click on the Application Assurance tab.
2. Select an application profile string.

   You can select only local profiles on the NE.
3. To associate a transit IP policy with an Epipe service object, select a transit IP policy in the Transit IP Policy panel.

   See 87.46 “To configure an AA transit IP policy” (p. 2860) for more information about how to configure a transit IP policy.
4. To associate an AA transit prefix policy with the service object, select a transit prefix policy in the Transit Prefix Policy panel.

   See 87.47 “To configure an AA transit prefix policy” (p. 2861) for more information about how to configure a transit prefix policy.
5. To associate an AA redundant protocol with an Epipe service object, select an AA redundant protocol in the AA Redundant Protocol panel.

14

Save the changes and close the forms.

Note: If you are configuring access dual-homing with local switching over PBB tunnels, you must configure two L2 access interfaces. The L2 access interfaces must be on LAGs that participate in the MC LAG. See Chapter 43, “MC LAG groups” for information about MC LAGs.

END OF STEPS

76.41 To configure LAG per-link hashing on a VLL Epipe or Ipipe L2 access interface

76.41.1 Prerequisites

You can configure weighted per-link hashing on an Epipe or Ipipe L2 access interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member.

76.41.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL Epipe or Ipipe and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces.

4

Right-click on the L2 access interface you want to modify and choose Properties. The VLL L2 Access Interface (Edit) form opens.

5

Click on the LAG Per Link Hash tab.

6

Configure the Class and Weight parameters.

7

Save your changes and close the forms.

END OF STEPS

76.42 To assign ingress and egress QoS policies to a VLL L2 access interface

76.42.1 Before you begin

The available panels and parameters vary depending on the NE, chassis type, and release.

If you are assigning QoS policies to an access interface on a 7210 SAS NE, see 76.43 “To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site” (p. 2184).

76.42.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface, and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the QoS tab.

5

Configure the Ingress Match QinQ Dot1P parameter.

The Ingress Match QinQ Dot1P parameter is configurable only on Epipe L2 access interfaces. The Ingress Match QinQ Dot1P parameter is not configurable if the port encapsulation type is ATM or FR.

6

Select an ingress QoS policy in the Ingress Policy panel.

Note: For Epipe and Ipipe VLL L2 access interfaces, if you select an access ingress policy which has a forwarding class mapped to an ingress queue group, you must ensure that the port you are configuring for the Epipe or Ipipe L2 access interface has the access ingress queue group with the same name created on it.
See 16.24 “To configure Ethernet ports” (p. 599) in Chapter 16, “Port and channel object configuration” for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

7

Select an ingress queue group template policy in the Forwarding Plane Redirect panel.

8

Configure the Instance ID parameter in the Forwarding Plane Redirect panel.

Note: The Instance ID parameter cannot be configured for Ipipe on Ethernet ports.

9

Configure the Egress Mark QinQ Top Bits Only parameter.

The Egress Mark QinQ Top Bits Only parameter is configurable only on Epipe L2 access interfaces.

10

Configure the required parameters in the Aggregate Rate Limit panel.

11

Select an egress policy in the Egress Policy panel.

Note: For Epipe and Ipipe VLL L2 access interfaces, if you select an access egress policy which has a forwarding class mapped to an egress queue group, you must ensure that the port you are configuring for the Epipe or Ipipe L2 egress interface has the access egress queue group with the same name created on it.
See 16.24 “To configure Ethernet ports” (p. 599) in Chapter 16, “Port and channel object configuration” for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

12

Select an egress queue group policy in the Port Redirect panel.

Note: Selecting an Egress Queue Group Template Policy permits the redirection of Ethernet traffic packets to a queue ID specified in the egress port queue group of the SAP. The following properties and restrictions apply:

- If an Egress Queue Group Template Policy is specified here, the policy must have port redirection enabled.
- You cannot use policy-based redirection with the queue group when the SAP has SAP-based redirection enabled.
- Port access egress redirection is only supported on Ethernet/LAG ports. It is not supported on SAPs bound on non-Ethernet, Eth-tunnel, or CCAG ports.
- Supported ports include access, hybrid, and HSMDA.
- Queue groups can be applied to SAPs that incorporate LAGs. The LAGs can include port members from just a single card or from multiple cards.
- If you edit a LAG incorporated by the SAP, you cannot remove the last LAG member if a queue group reference exists to the containing SAP.
- You cannot add a secondary LAG member that has a queue group mismatch with primary LAG member.

13

If you are configuring an L2 access interface for a 7705 SAR, or if the port you selected in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is not an HSMDA port, then go to Step 21 .

14

Configure the Packet Byte Offset (bytes) parameter. You must enable the associated Override parameter.

15

Select a WRR policy in the Egress HSMDA Override panel.

16

In the Shaper Group panel, select a Shaper Group for the access ingress port or access egress port .

17

In the IXR Specific panel, select an Egress Remark policy and Egress VLAN QoS policy, as required. See 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634) and 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594).

18

Select an HSMDA egress secondary shaper policy in the Egress HSMDA Override panel.

19

Configure the Use Shared Queue parameter.

20

Select an HS secondary shaper in the HS Overrides panel, if required.

21

Save the changes and close the forms.

END OF STEPS

76.43 To assign ingress and egress QoS policies to a VLL L2 access interface on a 7210 SAS site

Note: You can assign policies and configure parameters for 7210 SAS NEs only when the system resource profile is appropriately configured for the device. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” .

The available parameters and policies vary depending on the device type and chassis variant. The configurations that are supported on the site NE are shown on the form.

76.43.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface, and choose Properties. The VLL L2 Access Interface (Edit) form opens.

If you are configuring an L2 Uplink SAP on a 7210 SAS-K site to preserve the dot1p values of ingress packets, expand to the Uplink SAP, open the properties form, and perform Step 4 and Step 10.

4

Click on the QoS tab and expand the 7210/7250 Specific panel.

5

Select a SAP Access Ingress policy in the Ingress Policy panel.

When you assign an ingress policy, the Color Mode parameter setting for the meters in the policy must coincide with the Enable DEI parameter setting on the physical port. When Enable

DEI is selected, the Color Mode for meters must be set to Color Aware. When Enable DEI is not selected, the Color Mode for meters must be set to Color Blind. See 16.24 “To configure Ethernet ports” (p. 599) and 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .

If you are assigning a 7210, 7250, and 1830 SAP Access Ingress policy that contains a 7210 FC Meter Map policy, you must select the Enable Table Classification parameter. See 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544).

To support H-metering, you must choose an ingress policy with all meter rate modes set to trTCM (RFC 4115).

For 7250 IXR sites, the selected SAP Access Ingress policy must contain an assigned 7250 Ingress CoS policy.

6

Select an egress policy in the Egress Policy panel.

See 50.32 “To configure a 7210 SAP access egress policy” (p. 1558) for more information about 7210 SAP access egress policies.

SAP-based remarking values are defined in a 7210 SAP access egress policy. To enable SAP-based remarking on the 7210 SAS-X, you must enable the SAP QoS Marking parameter on the port; see 16.24 “To configure Ethernet ports” (p. 599) .

7

To enable table-based color-aware ingress classification, select the Enable Table Classification parameter. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

8

Select an Egress Remarking policy in the Egress Remark Policy panel.

9

Configure the required parameters in the Aggregate Rate Limit panel.

You can configure the Ingress Meter Rate (kbps) and Ingress Meter Burst parameters only after SAP creation.

You can configure the Ingress Meter parameter only during SAP creation. The parameter must be set to true to support H-metering.

You can configure the Egress Meter Rate and Egress Meter Burst parameters only when resources are allocated to the SAP Egress Aggregate Meter parameter in the system resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380). To allocate resources on the 7210 SAS-R, configure the Egress SAP Aggregate Meter parameter in the system resource profile policy assigned to the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).

You must also enable port-based scheduling on 7210 SAS-Mxp and 7210 SAS-R NEs; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384).

The Enable Egress Meter Stats parameter is available when a value is configured for the Egress Meter Rate parameter.

10

For interfaces on 7210 SAS-K sites, configure the required parameters in the Egress Dot1p Remarking panel.

For access ports, configured parameters take effect only when the Remarking parameter is set to true in the 7210 SAP Access Egress policy assigned to the interface; see Step 6.

For L2 Uplink ports, configured parameters take effect only when the Remarking parameter is set to true in the 7210 and 1830 Network policy assigned to the port; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636).

11

Save the changes and close the forms.

END OF STEPS

76.44 To configure scheduling on a VLL L2 access interface

76.44.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the Schedulers tab.

Note: The Schedulers tab is configurable only if a port is assigned to the interface.

5

To configure egress scheduling on a 7210 SAS site, enable and configure the required parameters in the Egress Aggregate Rate Limit panel and go to Step 9 .

6

To configure scheduling on a 7705 SAR site:

Note: For the 7705 SAR, scheduler behavior is determined by the scheduler mode, which is 4-Priority by default. You can only configure the Egress and Ingress Aggregate Rate Limit parameters when the Scheduler Mode parameter is set to 16-Priority. You can set

the Scheduler Mode to 16-Priority only when the port is on an MDA that supports 16-Priority. See the 7705 SAR documentation for more information.

If you change the Scheduler Mode parameter from 16-Priority to 4-Priority, the NFM-P automatically restores the default settings for the Egress Aggregate Rate Limit and Ingress Aggregate Rate Limit panels when you click Apply or OK.

1. In the Egress Scheduler panel, configure the Scheduler Mode parameter.
2. In the Ingress Scheduler panel, configure the Scheduler Mode parameter.
3. If you set the Scheduler Mode parameter to 16-Priority in the Egress Scheduler panel or Ingress Scheduler panel, configure the required parameters in the Egress Aggregate Rate Limit panel and Ingress Aggregate Rate Limit panel. Otherwise, save the changes and close the forms.

7

To specify that an aggregation scheduler policy is not applied to the interface:

1. Set the Aggregation parameter to Off.

   Note:

   The Aggregation parameter is not configurable if the port you selected in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is an HSMDA port.
2. Configure the required parameters.

   Note:

   The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

   The Frame-Based Accounting parameter is not configurable if the port you selected in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is an HSMDA port.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.
3. Select an ingress scheduler in the Ingress Scheduler panel.
4. If you are configuring this interface for an Ipipe service then go to 6 .
5. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
6. If the port you selected in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is an HSMDA port, save the changes and close the forms.
7. Select an egress scheduler in the Egress Scheduler panel.
8. If you are configuring this interface for an Ipipe service, save the changes and close the forms.
9. Select an egress policer control policy in the Egress Policer Control Policy panel.
10. Save the changes and close the forms.

8

To specify that an aggregation scheduler policy is applied to the interface:

1. Set the Aggregation parameter to On.

   Note:

   You cannot specify an access scheduler policy if the port you selected in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is an HSMDA port. Go to Step 9 .

2. Select an aggregation scheduler in the Aggregation Scheduler panel.

9

Save the changes and close the forms.

END OF STEPS

76.45 To assign ingress and egress ACL filters to the VLL L2 access interface

Note: IPv6 ACL filters are not supported on the 7705 SAR.
ACL filters are not supported for CPipe L2 access interfaces.

76.45.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the ACL tab.

Note: When you configure ACL filters on a 7210 SAS NE, you must configure the system resource profile appropriately. See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) for more information.

5

Select the required ACL filter policies.

6

Save the changes and close the forms.

END OF STEPS

76.46 To assign an accounting policy to a VLL L2 access interface

76.46.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the Accounting tab.

5

Select an accounting policy.

6

Configure the required parameters.

The Collect Egress Queue Statistics parameter can be configured only during VLL L2 access interface creation. See 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174).

7

Save the changes and close the forms.

END OF STEPS

76.47 To configure Ethernet loopback for a VLL Epipe L2 access interface on a 7705 SAR

76.47.1 Purpose

Perform this procedure to configure OAM loopback for an Epipe SAP on a 7705 SAR NE. You can perform individual loopback tests, or enable loopback as part of Y.1564 test configuration. See

89.7 “STM Y.1564 test configuration” (p. 2924) for more information about Y.1564 tests.

76.47.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an Epipe service and click Properties. The Epipe Service (Edit) form opens.

3

On the navigation tree, expand the 7705 SAR site, then right-click on the required L2 access interface and choose Properties. The properties form for the SAP opens.

4

Click on the Ethernet Loopback tab.

5

Configure the required parameters.

If you are configuring Ethernet loopback for Y.1564 tests, set the Mode parameter to Internal, and enable the Mac Swap parameter.

6

Save the changes and close the forms.

END OF STEPS

76.48 To assign a time of day suite to the VLL L2 access interface

76.48.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the TOD Suite tab.

5

Select a time of day suite in the Time Of Day Suite panel.

Note: You cannot assign a ToD suite to a L2 access interface if accounting statistics collection is enabled on the L2 access interface. You must disable the Collect Accounting Statistics parameter in Step 6 in 76.46 “To assign an accounting policy to a VLL L2 access interface” (p. 2189) .
SapEgrQosPlcyStats and SapIngQosPlcyStats statistics are only collected if a Time Of Day Suite is applied on the SAP.

6

Save the changes and close the forms.

End of steps

76.49 To assign a DoS or DDoS protection policy to a VLL L2 access interface or SDP binding

76.49.1 Prerequisite

This procedure applies to VLL Epipe or Ipipe services.

Note: A default DoS protection policy is automatically assigned to the interface.

76.49.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces|Spoke SDP Bindings, right-click on the L2 access interface or spoke SDP binding and choose Properties. The VLL L2 Access Interface (Edit) form opens or Spoke SDP Binding (Edit) opens.

4

Click on the Security tab.

5

Select a DoS protection policy in the NE DoS Protection panel.

6

Configure the MAC Monitoring parameter.

7

Click OK. The VLL L2 Access Interface (Edit) form or Spoke SDP Binding (Edit) closes and a dialog box appears.

8

Save the changes and close the forms.

END OF STEPS

76.50 To create MIPs and MEPs on an Epipe or Apipe L2 access interface

76.50.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL service and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand the Sites icon and click the required site. The site's properties form is displayed.

4

Expand the required site's icon on the navigation tree to show the L2 Access Interfaces.

5

Click the required SAP and choose Properties. The L2 Access Interface (Edit) form opens.

6

Click on the OAM tab, then on the Configuration sub-tab.

7

Configure the required Test Generation Options parameters.

Note: You can propagate test generation role settings to all MEPs on a SAP or service site by clicking Propagate to MEPs.

8

Click on the ETH-CFM sub-tab.

9

To add a MIP on a 7210 SAS NE, configure the MIP and MIP MAC Address parameters in the MIP Configuration panel and then go to Step 13 .

If you set the Initial MHF-Creation parameter to default during MEG configuration (see 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109)), you must set the MIP parameter to Enabled.

If you set the Initial MHF-Creation parameter to explicit during MEG configuration (see 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109)), set the MIP parameter to Disabled.

You can only configure the MIP MAC Address parameter when the MIP parameter is set to Enabled.

10

To create MIPs:

1. Click Create in the MIP Configurations panel. The MIP Configuration (Create) form opens.
2. Configure the parameters, then click OK. The MIP Configuration (Create) form closes

   The Primary VLAN Enable, VLAN ID, and MAC Address parameters are configurable only when a port is assigned to the interface.
3. Repeat the previous two steps to create additional MIPs, as required.
4. Click Apply in the L2 Access Interface (Edit) form. The default MD objects are then automatically created on the service site.
5. Go to the service site's properties form and click the OAM tab and then on the ETH-CFM sub-tab.
6. Scroll down to the Default Domain panel and multi-select all required entries, then click Properties. The CFM Vlan Default Domain (Multiple Instances) form opens.
7. Configure the parameters for the selected entries and click OK. The CFM Vlan Default Domain (Multiple Instances) form closes.
8. Click Apply and then return to the L2 Access Interface (Edit) form.
9. Scroll down to the MIPs panel and click Resync MIPs. The default Up and Down MIPs that were automatically created by the NE are displayed in the list.

11

For 7705 SAR NEs, configure the Hold MEPs Up When SAP is Down parameter in the Hold MEP Up on SAP Failure panel.

Note: You can configure the Hold MEPs Up When SAP is Down parameter on only one SAP per node.

12

To enable a tunnel facility MEP on the L2 access interface, set the Tunnel Fault Notification parameter in the Facility MEPs panel to Accept.

Tunnel Fault Notification is configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation.

13

To create a MEP:

1. Click Create in the MEPs panel. The MEP (Create) form opens.
2. Select a MEG in the Maintenance Entity Group panel.
3. Configure the parameters in the MEP panel.

   The Type and Interface Type parameters are automatically populated based on whether the MEP is created on a SAP, SDP binding, or Ethernet Tunnel Path Endpoint.

   For 7210 SAS-K NEs, to enable the Primary VLAN Enable parameter, you must enable the Primary VLAN parameter in the system resource profile for the device; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380).
4. Configure the parameters in the CCM panel.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.
5. Select a MEG sub-group in the MEG Sub-Grouping panel.
6. Configure the required Test Generation Options parameters.

   Note: You can only configure Test Generation Options parameters for Epipe access interfaces.

   The Test Generation Options parameters are only displayed when you set the Direction parameter to Up.
7. If the MD for the MEP has a Maintenance Domain Name Type of none and the associated MEG has a Maintenance Entity group Name Type of icc-based, then the Y.1731 Tests and AIS tabs are configurable. Click on the Y.1731 Tests tab.
8. Configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.
9. Click on the AIS tab and configure the required parameters.

   The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.
10. Select a MEG sub-group in the MEG Sub-Grouping panel.
11. Click OK. The MEP (Create) form closes.

14

Configure the parameters in the Squelch Ingress Level panel.

15

Configure the parameters on the LMM Session Stats Collection panel as required.

16

Save the changes and close the forms.

END OF STEPS

76.51 To configure microwave compression on an MW link SAP on a VLL L2 access interface

Note: You can configure microwave compression only on an MW link SAP associated with a 7705 SAR-8 or 7705 SAR-18.

Microwave compression cannot be configured on an MW link SAP when the Outer Encapsulation Value parameter of the port in Step 6 of 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) is set to 0 or 1.

See 15.38 “To configure a 7705 SAR MW link” (p. 499) and 15.39 “To configure a 7705 SAR MW link member” (p. 501) for more information about configuring a 7705 SAR MW link and MW link member.

76.51.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the Microwave tab.

5

Enable the Compression parameter.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

76.52 To configure an Ethernet tunnel on a VLL L2 access interface

Note: You can only configure an Ethernet tunnel if you are creating a same-fate SAP or a control/data SAP.

76.52.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the Ethernet Tunnel tab.

5

If you are configuring a fate-sharing Ethernet Tunnel Endpoint SAP, or same-fate SAP, go to Step 6 . Otherwise, go to Step 8 .

6

Click Create. The Ethernet Tunnel (Create) form opens.

7

Configure the required parameters.

8

Save the changes and close the forms.

END OF STEPS

76.53 To assign an ANCP policy to a VLL L2 access interface

76.53.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the ANCP Static Map tab.

5

Click Create. The ANCP Static Map (Create) form opens.

6

Configure the ANCP String parameter.

7

Select an ANCP policy.

8

Save the changes and close the forms.

END OF STEPS

76.54 To specify the CEM functionality for an Epipe or Cpipe L2 access interface with CEM encapsulation

Note: Consider the following when creating a VLL Epipe or Cpipe L2 access interface with CEM encapsulation:

- The Time Slots parameter of the DS0 channel must be configured with at least one time slot.
- Time slots are automatically configured for unstructured E1 and T1 endpoints
- The Clock Source parameter of the DS1 channel must be set to Node-Timed.

76.54.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces, right-click on the L2 access interface and choose Properties. The VLL L2 Access Interface (Edit) form opens.

4

Click on the CEM EPipe tab or the CEM CPipe tab.

5

Configure the required general parameters.

6

For Cpipe VLL L2 access interfaces, configure the required parameters in the VLL Specifics panel.

7

For Cpipe VLL L2 access interfaces, configure the Control parameter in the Asymmetric Delay panel.

You must enable the Control parameter before you can configure the Samples parameter and the Repeat Period parameter.

You must enable the Control parameter on both ends of the Cpipe VLL to ensure symmetric delay in the jitter buffer.

8

Save the changes and close the forms.

END OF STEPS

76.55 To switch to the redundant port for a VLL SAP from an L2 access interface properties form

76.55.1 Prerequisites

CAUTION

Service Disruption

Nokia recommends that you set the OLC state of a service to Maintenance before you switch to a redundant port or channel. See the NSP System Administrator Guide for more information.

Then reset the port to the previous state after the switch is complete.

Perform this procedure from a port properties form when port redundancy is configured on one or more a Cpipe or Epipe SAPs to:

- switch to the backup port or channel assigned to the SAP, for example, in the event of a fault on the primary port or channel
- switch back to the primary port or channel assigned to the SAP, for example, after a fault condition on the primary is cleared

Note: When multiple SAPs are selected, switching is attempted only for the SAPs that support port redundancy and on which port redundancy is properly configured.
Port redundancy switching fails if the SAP or port is included in one of the following redundancy objects:

- BGP-MH
- MC-APS
- MC-LAG
- MC-ring

76.55.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VLL and click Properties. The *VLL_type* Service (Edit) form opens.

3

On the navigation tree, expand Site→L2 Access Interfaces→L2 access interface→SAP, right-click on SAP and choose Properties. The Port (Edit) form opens.

4

Click on the Port tab.

5

Click Switch to Redundant SAP. A dialog box appears.

6

Click OK. The Port (Edit) form closes and a dialog box appears indicating that the SAP object on the current site is removed.

7

Click OK.

8

Verify the following:

- For on-node switching, the SAP is deleted from the primary port, and a new SAP is created on the site.
- For off-node switching, the SAP and the parent site are deleted, and a new site, SAP, and spoke SDP bindings are created for the service.

9

Choose the SAP in the navigation tree. The Port (Edit) form opens.

10

Click on the Port tab. The Manual Switch Backup State indicator displays one of the following:

- Backup SAP in Use, if you switch from primary to backup
- Backup Port Specified, if you switch from backup to primary

Note: When the Manual Switch Backup State indicator reads Backup SAP in Use, the displayed port redundancy information is for the primary port of the SAP, and is read-only.

11

Close the forms.

END OF STEPS

76.56 To configure FPE association on a VLL Epipe site

76.56.1 Purpose

Use this procedure to configure PW Port FPE on an Epipe switching site. This will allow the system to extend forwarding path processing for pre-defined applications, for example, MPLS over PW.

When this procedure is completed, an internal spoke SDP is automatically created.

76.56.2 Steps

1

Perform one of the following:

a. To create an FPE bound to a PXC path on a single port, go to step Step 2.

b. To create an FPE bound to a PXC path on a LAG, go to step Step 4 .

2

Create a PXC path:

1. Create a PXC port, see 16.59 “To configure PXC loopback ports” (p. 652).
2. On the PXC port, right-click on subport b and choose Properties. The Port Cross Connect SubPort (Edit) form opens.
3. Choose the Egress Scheduling Virtual Port tab.
4. Click Create and create an egress scheduling virtual port with the IP address of the NE.
5. Complete the changes and close the form.

3

Create an FPE on the NE, with PW Port enabled. In the Path panel, enter the PXC ID configured in the previous step. See 12.41 “To create an FPE” (p. 374). Go to step Step 6.

4

Create a LAG path:

1. Create two PXC ports, see 16.59 “To configure PXC loopback ports” (p. 652).
2. Create two LAGs. The LAGs must include the PXC subports of the PXCs created in the previous step, and the subports in each LAG must have the same direction (a or b). For example, LAG 1 can contain the a subports of the two PXCs and LAG 2 can contain the b subports. See 13.16 “To create a LAG” (p. 431)

5

Create an FPE on the NE, with PW Port enabled. In the Path panel, enter the LAG IDs configured in the previous step. See 12.41 “To create an FPE” (p. 374).

Note: The LAG entered as the FPE Xa LAG is not required to be the LAG including the direction a subports.

6

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

7

Choose a VLL service and click Properties. The Epipe Service (Edit) form opens.

8

On the navigation tree, right-click on Epipe Service and choose Create Epipe Site, or right-click on a site and choose Properties. The Epipe Site (Create|Edit) form opens.

9

Click on the PW Port FPE tab.

10

Enable PW Port FPE, and add all applicable configurations.

11

Complete the configuration and close the form.

The PW Port association appears in the PW Port FPE tab of the FPE properties.

END OF STEPS

77 VPLS management

77.1 Overview

77.1.1 Purpose

VPLS is a class of virtual private network multipoint L2 service that allows multiple customer sites to be connected in a single bridged domain contained within the service provider-managed IP/MPLS network. Customer sites in the VPLS appear to be on the same LAN, even if the sites are geographically dispersed.

VPLS offers the following advantages:

- Ethernet interfaces on the host access side simplify provisioning.
- All routers in the VPLS are part of the same LAN, which simplifies IP addressing and allows customers to control and simplify their routing strategies.
- VPLS is protocol independent, which means there is no L2 protocol conversion between LAN and WAN technologies.

You can use HVPLS to eliminate the need for a full mesh of virtual circuits between devices in the VPLS. See 77.2.1 “HVPLS” (p. 2210) in this section for more information.

A VPLS can span a single site or multiple sites. A VPLS that spans a single site is called a local VPLS. In a local VPLS, customer data enters the service through multiple access interfaces on a single PE device. No circuit provisioning is required for the local VPLS.

A VPLS that spans multiple sites is called a distributed VPLS. In a distributed VPLS, customer data enters the service using two or more interfaces on different PE devices. The VPLS is transported by service circuits over an IP/MPLS provider core network carried by service tunnels. Service tunnels are created using GRE or MPLS LSPs.

The NFM-P supports end-to-end VPLS configuration using tabbed configuration forms with an embedded navigation tree.

The General tab of the NFM-P service management form displays useful information about the operational state of the service and its sites through the Aggregated Operational State and State Cause indicators.

The Aggregated Operational State indicator has four possible values: Up, Down, Partially Down, and Unknown. The value is derived from the operational states of the sites that are part of the service, as follows:

- Up—all sites are operationally up
- Partially Down—at least one site is operationally down
- Down—all sites are operationally down
- Unknown—the service has no provisioned sites

When the Aggregated Service Site Operational State is Partially Down or Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P operator. Alarms can be viewed on the Faults page.

When you use the NFM-P to create or discover a service, the NFM-P assigns a default tier value to the service. The Service Tier parameter value is relevant only in the context of composite service topology map views. See Chapter 85, “Composite service management” for more information about the hierarchical organization of composite services.

Common to all services, such as VPLS, are policies that are assigned to the service. Policies are defined at a global level and can then be applied to components of the service, such as interfaces and circuits, when the service is configured or modified.

The following policies are common to all services:

- QoS policies to define ingress classification, policing, shaping, and marking on the ingress side of the interface. QoS policies are configured using the Access Ingress Policy Manager and the Access Egress Policy Manager.
- Policer control policies to control access ingress policers and access egress policers under a common hierarchy. Policer control policies are configured using the Policer Control Policy Manager.
- Scheduling policies to define hierarchical rate limiting and scheduling to govern the scheduling of queues. Scheduler policies are configured using the Scheduler Policy Manager.
- Port scheduler policies define hierarchical bandwidth allocation and scheduling at the egress port level. Port scheduler policies are configured using the Port Scheduler Policy and HSMDA Scheduler Policy forms.
- Filter policies to control network traffic into or out of an interface or circuit based on IP or MAC matching criteria. Filter policies are configured using the ACL IP Filter Manager and the ACL MAC Filter Manager.
- Accounting policies to count the traffic on a service to ensure proper billing and enforcement of SLAs. Accounting policies are configured using the Accounting Policy Manager.
- ANCP policies provide status and control information based on port-up and port-down messages and current line rate changes between the edge device and the access node. ANCP policies are configured using the Manage Subscriber Policies form.
- Time of day suites specify time and day restriction policies that are assigned to QoS policies and schedulers, ACL filters, and aggregation schedulers. Time of day suites and time range policies are configured using the Time of Day Suite form and Time Range form, respectively.

See Chapter 49, “Policies overview” for more information about policies.

Packets that arrive at an edge device are associated with a VPLS based on the access interface on which they arrive. An access interface is uniquely identified using the following parameters:

- physical Ethernet port or POS port and channel
- encapsulation type
- encapsulation identifier (if required)

If there are service issues, the service provider can use OAM tools to troubleshoot service and network transport issues, and ensure problems are handled properly through the physical and logical network.

To provide a VPLS over an MPLS infrastructure, the device is configured to provide bridging and replication for each VPLS. The routers that are part of the VPLS are connected by MPLS LSPs. Multiple VPLS can use the same set of service tunnels. Multiple service tunnels can rely on multiple LSPs. The signaling is specified in sets of ingress and egress VC labels for each VPLS.

The following additional features are configured for the VPLS:

- MAC learning for the access ports and tunnels, including filtering based on MAC addresses on a per SAP basis
- MAC learning protection on SAPs to prevent DoS attacks from sourcing
- rate limiting of broadcast, destination unknown, and multicast traffic on a per access port basis
- FIB for each VPLS, including FIB size limits, static MAC addresses, alarms, and discarding unknown locations
- optional support for spanning tree for loop detection
- GSMP for each VPLS
- L2 management interfaces for each VPLS

You can run an OAM Validation test suite for the service by clicking on the Validate button. If the Validate button is not visible, click on More Actions and choose Validate. Alternatively, you can also perform a One Time Validation. If a check mark appears beside the OAM Validation Failed state cause indicator, the test has failed. In addition, the Tested Entity Result tab on the Tests tab displays detailed information about the OAM test result. OAM validation tests are not supported for HVPLS.

77.1.2 Contents

VPLS management overview

77.2 VPLS management overview

77.2.1 HVPLS

A hierarchical VPLS is created by enhancing the VPLS core mesh with a spoke SDP binding that is connected to another site in the same VPLS, a site in another VPLS, or a VLL site.

HVPLS offers the following advantages:

- reduces the complexity of mesh configuration
- decreases the amount of signaling of routes between devices

When traffic arrives at an access-spoke circuit, it acts like a bridge port where flooded traffic received on the access spoke is replicated to all other spokes, meshes, or SAPs but it is not transmitted on the port where it is received.

Sample configuration

The following figure shows a sample HVPLS with a mesh and spoke configuration. Spokes 50 and 51 are unidirectional access-spoke circuits bound to service tunnels. The access-spoke circuits exist within the context of a VLL service or VPLS that is interconnected to the original, fully meshed VPLS. Alternatively, the access-spoke circuit can provide interconnectivity to a service site.

Figure 77-1 HVPLS configuration

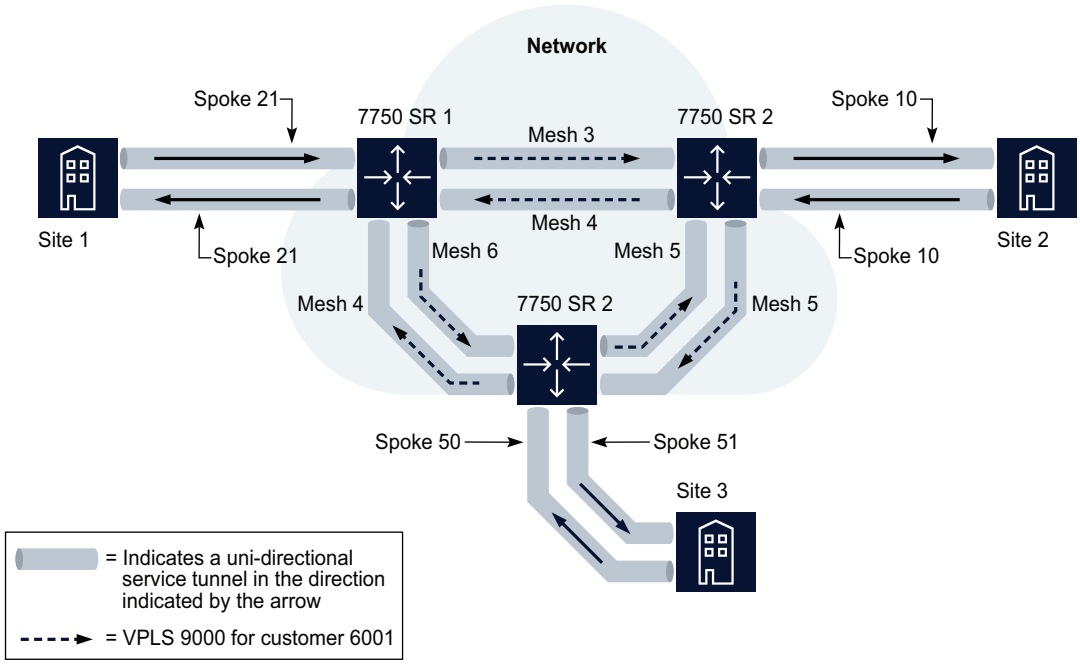

17438

The NFM-P supports the following HVPLS interconnectivity using spoke SDPs:

- VPLS to VPLS
- intra-VPLS
- VPLS to VLL

77.2.2 MVPLS

VPLS topology loops can occur if either of the following is true:

- Two VPLS are connected by redundant spoke SDPs.
- A CE NE is connected to a VPLS with redundant L2 access interfaces.

To remove topology loops, RSTP must be enabled on the redundant spoke SDPs or L2 access interfaces to block some of them from passing traffic. This requires the creation of an MVPLS.

MSTP is an extension of RSTP which allows VLANs to be grouped into spanning tree instances. Each instance has an independent spanning tree topology. MSTP can be run in an MVPLS to provide multiple forwarding paths for data traffic, which allows load balancing and reduces the number of spanning tree instances required to support a large number of VLANs. An MST region comprises a set of interconnected switches that have the same MST configuration. Each region can be configured with up to 16 MST instances. The instance with ID 0 is an internal spanning tree that runs an MST region and sends and receives BPDUs. All other spanning tree instance information is encapsulated within MSTP BPDUs.

An MVPLS is created to run RSTP or MSTP and manage traffic on the associated VPLS. An MVPLS contains sites, spoke SDP bindings, mesh SDP bindings, and L2 access interfaces. The MVPLS spoke SDP bindings and L2 access interfaces are configured to manage the associated VPLS spoke SDP bindings and L2 access interfaces.

In the case of spoke redundancy, the MVPLS runs RSTP on the redundant spoke SDPs and associates the resultant traffic-blocking actions with all VPLS that use the same spoke SDPs.

MVPLS traffic blocking can also be used on the access side to manage redundant L2 access interface connections. A VLAN ID range is specified for each MVPLS L2 access interface which identifies the VC IDs of the managed VPLS L2 access interfaces.

RSTP is enabled by default on an MVPLS. When the Admin state of an MVPLS is down, all managed L2 access interfaces and spoke SDPs in the associated VPLS are disabled. However, if the Admin state of individual L2 access interfaces or spoke SDPs of an MVPLS are down, then the managed VPLS L2 access interfaces or spoke SDPs are not affected by traffic blocking.

A common MVPLS situation occurs when two VPLS are connected by redundant spoke SDPs. If traffic is not blocked on one of the redundant spoke SDPs, then a loop results. To remove the loop, RSTP must be run on the spoke SDPs that form the loop to block one of the redundant spoke SDPs. Blocking is accomplished by creating an MVPLS on each side of the redundant spoke SDPs and creating a composite MVPLS to connect the MVPLS.

Another common MVPLS situation occurs when an access switch with many VLANs is redundantly connected to two other bridges, on which each uplink carries half the VLANs. MSTP allows you to

build multiple spanning trees over VLAN trunks and to group and associate the VLANs to spanning tree instances, each with a different port instance cost and port instance priority.

The following figure shows an example of a composite MVPLS that is composed of MVPLS 1, MVPLS 2, and spoke SDPs.

Figure 77-2 Composite MVPLS

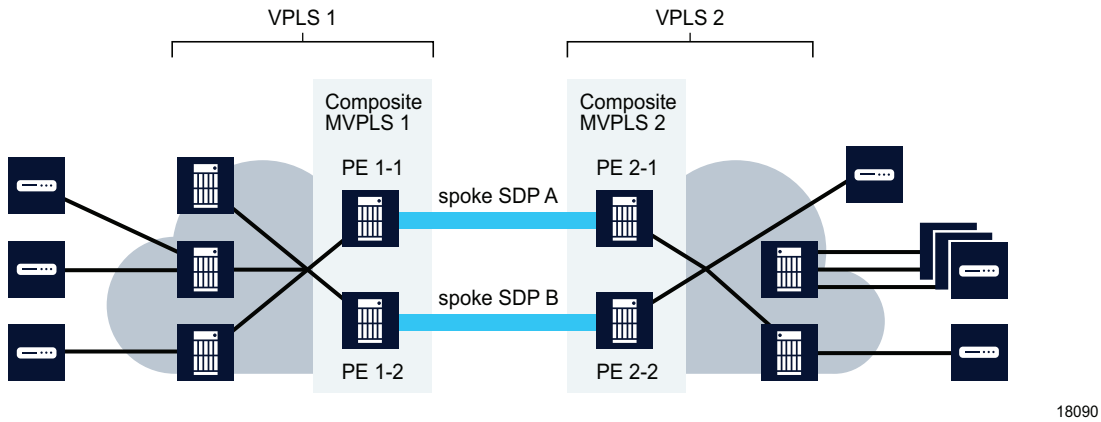

Another scenario occurs when multiple L2 access interfaces from a VPLS are connected to the same customer edge equipment. In this case, a single MVPLS must be created with L2 access interfaces defined to manage the traffic on the associated VPLS redundant L2 access interfaces.

77.2.3 Dual homing for VPLS

A VPLS can be configured for dual homing through the use of redundant spoke SDPs. NFM-P handles the redundant spoke SDPs by grouping them together to form an endpoint object. The redundant spoke SDPs provide active and standby pseudowires for the service. This spoke SDP access arrangement allows data flow control and management support without requiring STP, which cannot be enabled on a spoke SDP binding that is under an endpoint. For VPLS, you can associate only spoke SDP bindings with an endpoint, and each endpoint can be associated with a maximum of two spoke SDP bindings.

Sample configurations

The following figure shows a simple dual homing configuration.

Figure 77-3 MTU redundant access to VPLS

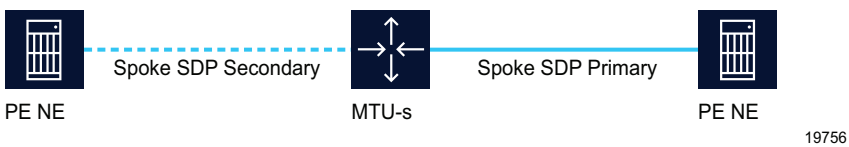

VPLS dual homing provides the ability to have an NE deployed as an MTU-s with links to multiple PE NEs without requiring an MVPLS.

Note: You cannot create a VPLS endpoint on a site that has an active or inactive MC ring SAP. See Chapter 45, “MC ring groups” for more information.

You cannot create an endpoint in an MVPLS.

In this example, the MTU-s has spoke SDPs to two PE devices. One is designated as the primary spoke and the other as the secondary, or standby spoke, based on a precedence value specified for each spoke. The standby spoke is in a blocking state when the primary spoke is available. If the primary spoke becomes unavailable, the MTU-s immediately switches the traffic to the standby spoke. You can configure the service to revert back to the original configuration, after a specified delay, when the primary spoke is again available. Forced manual switchover is also supported.

You can configure a MAC flush to speed the convergence during a switchover. The PE devices that receive the MAC flush remove each MAC address that is associated with the affected VPLS instance and forward the MAC flush to the other PE devices in the VPLS. The following figure shows a dual-homed VPLS for BTV distribution.

Figure 77-4 BTV distribution in redundant VPLS architecture

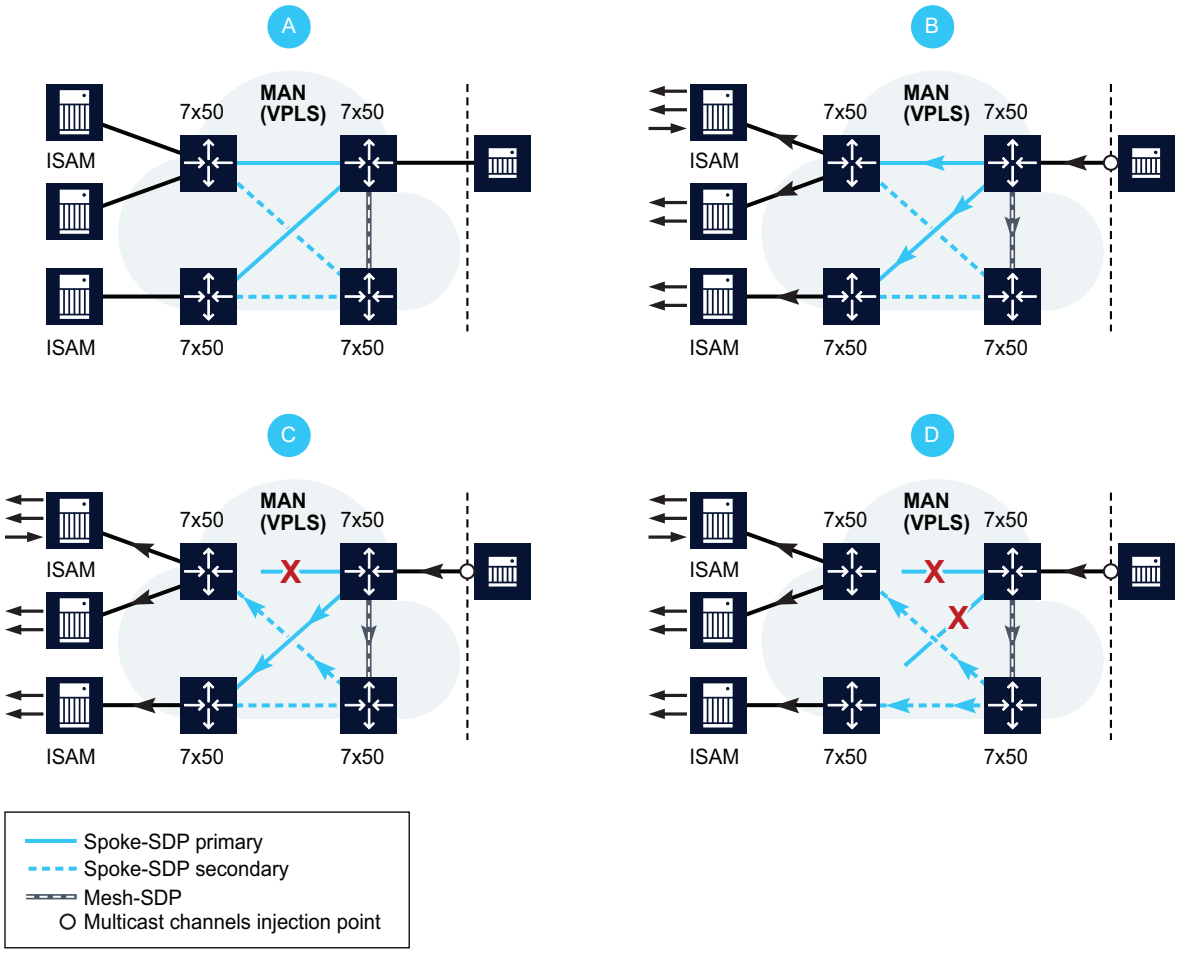

In the nominal operating mode shown in panel A, the edge router (grey icon) has been configured to statically join all multicast channels and inject them into the aggregation network. The access layer 7x50 unit (directly connected to ISAM) is dual-homed into two aggregation layer 7x50s edge routers (larger icons) using primary and secondary spoke SDPs. A mesh SDP interconnects both aggregation nodes. Injected BTV traffic from the edge router is broadcast on the primary spoke SDPs to the connected MTU devices (panel B).

A copy of the channels is also sent on the mesh SDP to the peer aggregation node, which also replicates the traffic to the connected spoke SDPs (aggregation layer nodes are not aware of primary/secondary spoke selection done by the MTU layer devices). The MTUs only receive traffic from the primary spoke SDP. Traffic received on the secondary spoke SDP is blocked. In the event of a link failure (panel C) or MDA failure (panel D), the MTU switches over to the secondary spoke SDP and immediately start receiving traffic from it instead of the primary spoke SDP.

Composite services also support VPLS with redundant spoke SDP bindings to VLL services. The following figure shows a VPLS and VLL combination example that provides an E2E redundant path.

Figure 77-5 VPLS and VLL combination to provide E2E redundant path

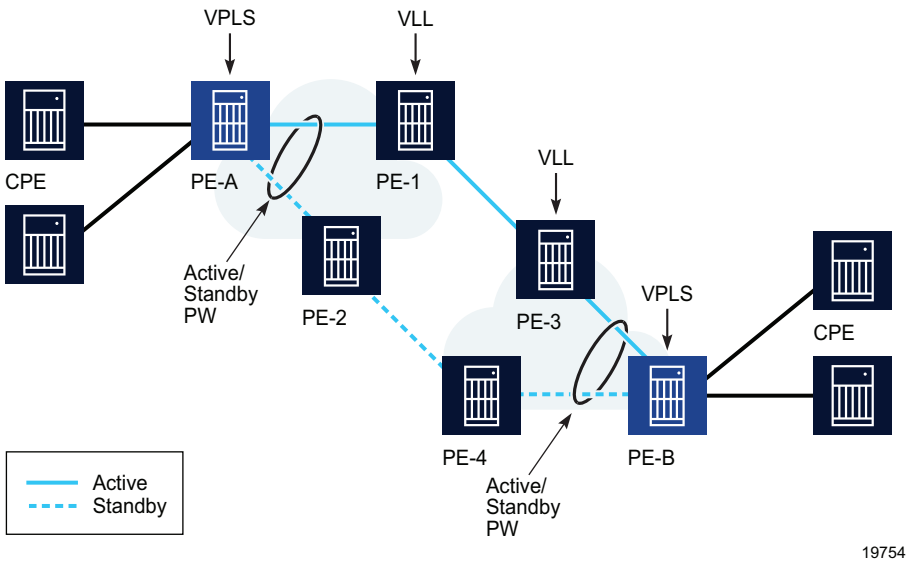

77.2.4 Provider Backbone Bridging in VPLS

Provider Backbone Bridging (PBB) is a technology configuration employed in next-generation networks that utilize carrier-grade Ethernet as the transport architecture. It addresses the potentially enormous increase in MAC addresses stored in the router lookup databases by encapsulating the customer frame in a Provider Ethernet header. The Customer MAC address (C-MAC) is then only dealt with by lower tier (or satellite) H-VPLS PEs. The core H-VPLS PEs only need to handle the MAC addresses (B-MAC) of the backbone provider, which are substantially less in number. For this reason, the technique is also referred to as MAC-in-MAC encapsulation.

IEEE 802.1ah defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs). PBB is defined in IEEE as a connectionless technology based on multipoint VLAN tunnels. MSTP is used as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks.

A Provider Backbone Bridged Network is a Virtual Bridged Local Area Network that comprises Backbone Edge Bridges (BEBs) and Backbone Core Bridges (BCBs) under the administrative control of a single backbone provider. Each BEB provides interfaces that encapsulate or verify the encapsulation of customer frames, and then relay those frames across the backbone. The term may also mean a provider who is purchasing a service from another provider and using either a PBN or PBBN internally.

Backbone VLANs are used to create multipoint trunks in the backbone. The B-VLAN determines the route the frames take and limits broadcasting within the backbone. The B-TAG is added to the

frame at the Customer Backbone Port (CBP). The selection of B-VLAN used to form the B-TAG is determined by the configuration of the CBP service instance table. This table maps ISIDs to B-VIDs and is created as part of service provisioning.

Backbone VLAN Connectivity

The backbone provider can use and configure MSTP to provide a number of independent spanning tree active topologies and can assign each B-VLAN to one of these active topologies to best use the resources in the network. MVRP, running in the context of each spanning tree active topology, configures the extent of each B-VLAN to the subset of that active topology necessary to support connectivity between the customer points of attachment to the instance of MAC service provided, and can reconfigure that connectivity as required if the spanning tree active topology changes. The operation of MSTP within a backbone provider's network is independent of the operation of any spanning tree protocol within attached provider or customer networks. This is achieved by removing all MSTP BPDUs received or to be transmitted at the service access interfaces. The operation of MVRP within a PBBN is independent of the operation of any configuration protocol within attached customer networks.

SR PBB implementation

The IEEE PBB model is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of Customer/Provider Bridge (QinQ) domain (that is, MACs, and VLANs) to the provider backbone (that is, B-MACs, and B-VLANs). The I-component contains the boundary between the customer and backbone MAC domains. PBB encapsulates the customer payload in a provider backbone Ethernet header, which allows the C-MACs to be hidden from the core PEs. A special Group MAC is used for the Backbone Destination MAC when the customer frame type is either unicast, multicast, broadcast, or unknown.

The SR PBB solution can be summarized as follows:

- Two VPLS variants are employed, namely B-VPLS and I-VPLS, functioning as the B-type BEB, and the I-type BEB respectively. A B-VPLS instance (a service instance within a service router) and its corresponding I-VPLS instances must be co-located in a service router. From a network-wide perspective, a B-VPLS comprises multiple Backbone Virtual Switch Instances (B-VSIs).

  Note:

  For the description in this section and in the procedures in this chapter, B-VSI is used for a single B-type BEB instance, and is referred to as a B-Site. Similarly, I-VSI is used for a single I-type BEB instance, and is referred to as an I-Site.
- mB-VPLS and mI-VPLS are also available to provide loop avoidance for B- and I-VPLS in the same way as m-VPLS and regular VPLSs operate.
- An NFM-P VPLS can include regular sites and either B-Sites or I-Sites. The service should not contain both B-Sites and I-Sites. When a VPLS has at least one B-Site it becomes a B-VPLS. When a VPLS has at least one I-Site, it becomes an I-VPLS.

  Note:

  Regular sites, B-Sites, and I-Sites cannot change their type after they are created.
- An I-Site can be bound to one B-Site, but a B-Site can be used by multiple I-VPLSs.

B-VPLS and I-VPLS instances

PBB processing may be seen as a chain of two linked VPLS contexts, namely B-VPLS and I-VPLS.

Their characteristics are summarized in the sections that follow.

I-VPLS

I-VPLS is the abbreviated form for Service Instance ID (ISID) VPLS. An I-VPLS instance on a service site is referred to as an I-Site.

The following are I-VPLS and I-Site characteristics:

- An I-Site operates using customer addressing and maps the C-MACs to B-MACs.
- You can select one B-Site to associate with an I-Site; the B-Site must be on the same PE NE as the I-Site.
- I-Sites support only spoke SDP bindings and not mesh SDP bindings.
- An I-Site L2 access interface, or I-SAP, can co-exist on a port with regular L2 access interfaces or subscriber management M-L2 access interfaces. The existing port encapsulation is supported. An encapsulation tag that is used for service selection on an I-SAP is removed before the PBB encapsulation is added. The appropriate encapsulation tags are added at the remote PBB PE when sending the packet out on the egress access interface.
- An I-Site can be connected to one or more regular VPLS sites. A regular (network level) VPLS can have a mix of regular sites and I-Sites. The I-Sites of such services are responsible for the mapping of C-MACs to B-MACs. The regular sites of the service function as normal (bridge).
- ISID is a 24-bit field that carries the service instance identifier associated with this frame. It is used at the destination PE as a demultiplexer field, a function similar to a VC label. Default to service ID only works if the service ID is within the ISID range. For a service with service ID larger than 16 777 215, the ISID value must be specified.
- The ISID must be unique on one router.
- The Provider MSTP support in an M-VPLS is in the I-VPLS space.
- The I-Site MTU must be at least 18 bytes smaller than the B-Site MTU to which it is bound.
- If a VPLS has an attached I-Site, the Include I-Site(s) indicator on the General tab of the VPLS configuration form is selected.
- IGMP snooping can be configured for I-Sites and I-L2 Access Interfaces, except on routed I-VPLS services.
- PIM snooping can be configured for I-Sites and SAP and Spoke SDP bindings of I-VPLS services.

Backbone-VPLS (B-VPLS)

Multiple L2 services can use a single B-VPLS. Ordinarily, a pair of SDP bindings (in opposite directions) provide either point-to-point connection between two sites of a service (as SDP bindings) or between different services (as a service connector). However, a B-VPLS provides a multipoint connection between sites of a service or for multiple services.

The following are properties and characteristics of a B-VPLS and B-Sites:

- The B-VPLS operates using the provider or backbone addressing (B-MACs).
- The B-VPLS provides backbone tunneling for one or multiple I-VPLSs.

- The B-VPLS accepts mesh or spoke SDP bindings, thereby providing both routing and MAC hiding using PBB/PW encapsulation.
- The B-VPLS accepts access interfaces using PBB encapsulation for tunneling through an Ethernet-only network.
- A regular (network level) VPLS can have a mix of regular sites and B-Sites to function as a B-VPLS (that is, operating using B-MACs).
- The backbone's Source MAC address can be configured on a B-Site. All the I-Sites provisioned under this B-Site shares the provisioned values. By default, this is a loopback chassis MAC address. It must be a unicast MAC address.
- A B-Site site can not be deleted until all its I-Site associations are removed.
- A B-Site can have both spoke and mesh SDP bindings and only an MPLS type of tunnel can be used (including LDP SDP). This also applies for a regular pseudowire, where the outgoing PBB frame on a B-SDP (that is, a B-PW) contains a B-VID qtag only if the PW type is Ethernet VLAN. Alternatively, if the pseudowire type is Ethernet, the B-VID qtag is stripped before the frame goes out.
- Only Null, dot1q, and QinQ encapsulation types can be used by a B-Site L2 access interface. These access interfaces must use PBB encapsulation and have the following properties:
  - Ethernet dot1q is applicable to the bulk of PBB use cases, such as one B-VID.
  - Ethernet Null is supported for direct connection between neighboring I-VPLS, for example, when no B-VID is required and all traffic is sent to or from local I-VPLS.
  - There is no requirement for a PBB SAP type for PBB on the B-VPLS SAPs. Only the B-VID is used for tunnel delimitation on the port.
  - The default access interface type is blocked for the B-L2 access interface.
  - The following rules apply to the SAP processing of PBB frames:

    > For transit frames (frames not destined to a local MAC), there is no need to process the I-tag component of the PBB frames. Regular Ethernet SAP processing is applied to the backbone header (B-MACs and B-VID).

    > If a local I-VPLS instance is associated with the B-VPLS, then local frames (frames originated or terminated on local I-VPLSs) are PBB encapsulated and de-encapsulated using the pbb-etype provisioned under:

    related port->I-VPLS->root pbb component

    (listed in decreasing order of precedence, where the related port is highest in the order).
- If a VPLS has an attached B-Site, the Include B-Site(s) indicator on the General tab of the VPLS configuration form is selected

Service topology map views

Service and composite service topology maps support PBB.

The service map shows different types of sites with various icons for I-Sites, B-Sites, and Epipe PBB sites. With an I-VPLS bound to B-VPLS, the map shows the PBB backbone network as a cloud. The bindings between I-VPLS and B-VPLS are shown as a binding link.

MRP and MMRP support

The Multiple Registration Protocol allows participants in an MRP application to register Group MAC addresses with other participants in a Bridged LAN. An MRP participant may transmit and receive

MRP PDUs. For the PBB implementation, the MRP parameters can be configured at the service site level, on the access interface, or the SDP binding.

If MRP is enabled on the NE, the NFM-P MMRP application automatically advertises the presence of the Group B-MAC address on the active B-VPLS virtual links—that is, on the B-Site spoke bindings or the B-L2 access interfaces. You can view the MMRP entries advertised and/or received on the Forwarding Control→MMRP Entries tabs of the B-Site spoke bindings or the B-L2 access interfaces. All of the MMRP entries may also be viewed together at the I-Site level.

Epipe service with PBB

A PBB tunnel may be linked to an Epipe and to a B-VPLS. MAC switching and learning is not required for the point-to-point service, since all packets ingressing the Epipe access interface are PBB encapsulated and then forwarded to the PBB tunnel for the backbone destination MAC address. Similarly, all the packets ingressing the B-VPLS and destined for the ISID are PBB de-encapsulated and then forwarded to the Epipe access interface. A fully-specified backbone destination address must be provisioned for each PBB Epipe instance to be used for each incoming frame on the related Epipe L2 access interface. If the backbone destination address is not found in the B-VPLS forwarding database, then packets may be flooded through the B-VPLSs.

To enable an Epipe service with PBB, the Epipe site is configured as a PBB site and functions similarly to a VPLS I-Site. This can only be specified during the site's creation. On the Epipe PBB site, you select a B-Site that acts as the tunnel for the Epipe. You also specify the destination B-MAC address of the remote PE where the other site is located. The following figure shows a simplified view of this configuration.

Figure 77-6 Epipe service link to a B-VPLS

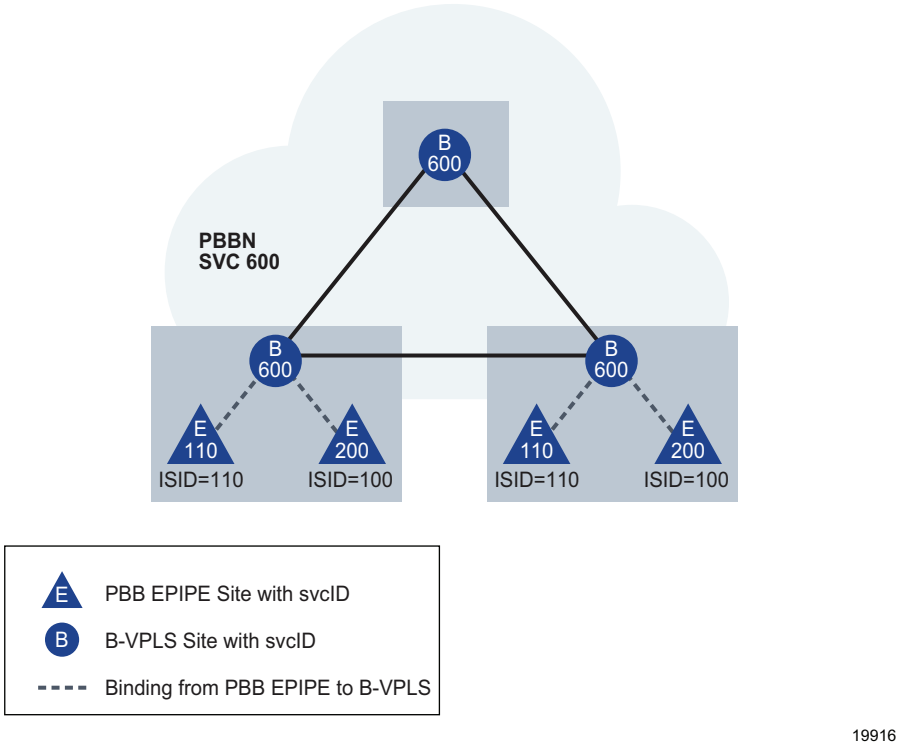

See the *7750 SR OS Services Guide* for more information about PBB.

77.2.5 EVPN

The VPLS solution has some limitations: MAC address learning is limited in the data plane, and active-active multi-homing is not available. Ethernet VPN addresses these limitations.

EVPNs are used to connect data centres. An EVPN is a network of PE and CE devices, connected using a layer 2 overlay over a layer 3 network. An EVPN instance is configured on a PE router. The PE router can be the VPLS interface of a VPRN.

In an EVPN, MAC learning between PEs occurs in the control plane. The control plane for advertising MAC reachability information is Multiprotocol BGP (MP-BGP). PEs learn MAC addresses from the CEs that are connected to them and advertise them to other PEs in the control plane using MP-BGP. Control plane learning allows EVPN to support different data plane encapsulation technologies between PEs. The PEs may be connected to each other by an MPLS LSP infrastructure or an IP infrastructure. If an IP or MPLS infrastructure is in use, IP/GRE or MPLS tunneling can be used between the PE devices.

In addition to MAC learning within the control plane, BGP-EVPN is capable of efficient multi-destination traffic delivery and active-active multi-homing.

The data plane is configured in VPLS and Epipe. The following data planes are supported:

- EVPN-VXLAN:

  With EVPN-VXLAN integration the VPLS site performs data plane learning on the traffic received from the VXLAN and implements EVPN to distribute the client MAC addresses learned into BGP. VXLAN or Ethernet frames are encapsulated with MPLS when sending the packets over the MPLS core and with the VXLAN tunnel header when sending the packets over the VXLAN network.

  Depending on the NE and release, up to two VXLAN instances can be created on a VPLS. If two instances are created, one sends data to the DC LAN and one toward the WAN.

  EVPN tunneling with EVPN-VXLAN:

  Each VXLAN can be identified using a VNI in the VPLS. The VNI needs to be the same for the VPLS to communicate. Each VPLS is associated with a VPRN for L3 connectivity. If an EVPN tunnel is configured in an IRB backhaul R-VPLS, there is no need to provision the IRB IPv4 addresses on the VPRN.

- EVPN-MPLS:

  MPLS edge switches (MES) can be PEs in an EVPN. The MES devices are interconnected using LSPs. EVPN-MPLS integration allows the use of MPLS functionality for routing with EVPN control plane learning.

- EVPN-VPWS:

  EVPN with VPWS offers the benefits of EVPN in a point-to-point implementation. An EVPN-VPWS network offers single-active or active-active multihoming.

- PBB-EVPN:

  With PBB-EVPN, instead of sending the customer MAC addresses as control plane learning, the backbone MAC addresses are distributed in the EVPN. This simplifies control plane learning in the EVPN and allows for increased efficiency in a large network.

- ETree:

  An ETree service is a type of Ethernet service that is based on a rooted-multipoint Ethernet virtual connection. In an ETree, each attachment circuit is designated as either a root or a leaf AC. A leaf AC can send or receive traffic only from a root. A root can send traffic to another root or to a leaf. If a root or leaf tag specified in a SAP, it will replace the outer VLAN tag for data routing.

Ethernet segments

An ES is a group of ports on an NE that are part of the same redundancy group, and are identified by a unique Ethernet Segment Identifier. They can be associated with ports, LAGs, SDPs, or VXLANs. To support multi-homing, the ESI should be the same between two PEs.

The use of an ES composed of physical links satisfies the redundancy requirements for CEs that are directly connected to the ES PEs by a port, LAG, or SDP, however it does not work when there is an aggregation network between the CEs and the ES PEs, and different ESs must be defined in the same port or LAG.

The concept of the physical links in an ES is extended to Ethernet Virtual Circuits where many of such EVCs can be aggregated on a single physical External Network-to-Network Interface (ENNI).

An ES that consists of a set of EVCs is referred to as a virtual ES (vES).

A vES can be associated with:

- Qtag-ranges on dot1q ports or LAGs
- S-tag-ranges on QinQ ports or LAGs
- C-tag-ranges per S-tag on QinQ ports or LAGs
- VC-ID ranges on SDPs
- Service ranges on VXLANs
  For service ranges to be created, the Network Interconnect VXLAN must be set to 1.

Interconnect Ethernet segments

An I-ES is a virtual ES, that is, an ES that consists of a set of Ethernet Virtual Circuits instead of a set of physical links. An I-ES allows DC GWs with two BGP instances to handle VXLAN access networks. I-ESs support the multi-homing functions of BGP MPLS EVPN.

An I-ES is used to connect a VXLAN to MPLS. Both services will be up at the same time with two different BGP instances.

77.2.6 SPB in VPLS

SPB simplifies network creation and configuration because service provisioning is required only at the network edge. SPB uses IS-IS to dynamically build the topology between NEs. The NFM-P supports SPB on B-VPLS services.

SPB allows shortest path forwarding in a mesh network by using multiple equal cost paths. In addition, SPB supports a link state spanning tree. SPB is appropriate for larger layer 2 topologies, with faster convergence than STP variant protocols for the same size networks, and improved use of the mesh topology.

The main components of B-VPLS SPB are:

- control B-VPLS
- user B-VPLS

The control B-VPLS defines the SPB IS-IS instance and is associated with an ECT algorithm and a set of SAPs or spoke SDPs for link connectivity. The user B-VPLS is a logical B-VPLS subset of the control B-VPLS. User B-VPLSs are associated with the forwarding ID, which is a logical backbone VLAN ID, that you define in the control B-VPLS. The user B-VPLS forwarding ID does not have to match the control B-VPLS forwarding ID. The user B-VPLS must be a subset of the control B-VPLS, or else a black hole can occur.

You can perform the following OAM diagnostic tests on SPB-enabled B-VPLSs:

- CFM link trace
- CFM loopback

Note: SPB does not support MEPs on SAPs and spoke SDP bindings. You cannot create an SPB instance on a control or user B-VPLS if any MEPs or MIPs are enabled on a SAP or

spoke SDP binding. Virtual MEPs are supported. You cannot configure MIPs on SPB-enabled SAP and spoke SDP bindings.

Static B-MACs and ISIDs

An SPB interface on a SAP or SDP can have static B-MACs and static ISIDs that are not part of the SPB network or region. Static B-MACs and ISIDs allow SPB networks to interface with other PBB networks that use other control planes. Static B-MACs allow remote PBB Epipes to connect to SPB managed networks. Static ISIDs allow I-VPLS services to connect to non-SPB I-VPLS services. You can define an ISID policy to use the default multicast tree and to suppress the advertisement of ISIDs in SPB when I-VPLS or static ISIDs are used for unicast services. See 77.27 “To create a static ISID range on a VPLS B-L2 access interface or spoke SDP binding” (p. 2288), 77.65 “To create a static B-MAC on a B-VPLS site” (p. 2329), and 77.66 “To create an ISID policy on a control or user B-VPLS site” (p. 2330) for more information.

SPB B-VPLS audit

You can perform an RCA audit on an SPB-enabled B-VPLS to verify whether:

- all user B-VPLS SAPs and spoke SDP bindings are fate-shared with a control B-VPLS. The following are the fate-sharing rules:
  - a user B-VPLS Dot1Q SAP fate-shares with the control B-VPLS SAP if they are in the same physical port
  - a user B-VPLS QinQ SAP fate-shares with a control B-VPLS SAP if they are in the same physical port and have the same top tag
  - a user B-VPLS spoke SDP binding fate-shares with a control B-VPLS spoke SDP binding if they are in the same SDP
- user B-VPLS instances belonging to the same service on different sites must have same forwarding ID
- the ECT-to-forwarding ID mapping is consistent on all the sites for the active B-VPLS services
- the ISID policy (Control/User) of all sites in the B-VPLS is consistent. All sites in this service must have the same ISID configuration if the Use Default Multicast Tree and No Advertise Local parameters are enabled.
- the ISID policy (Control/User) of all sites in the B-VPLS is configured. ISID ranges that are defined by one service site are be defined by all service sites of this service.

77.2.7 BGP Auto Discovery

BGP Auto Discovery enables a VPLS PE router to discover other PE routers that are part of the same VPLS domain. This allows each PE's configuration to consist only of the identity of the VPLS instance established on a specific PE, and not the identity of every other PE in that VPLS instance. If you need to change the VPLS topology, only the affected PE's configuration needs to change. Other PEs automatically discover the change using MP-BGP and adapt themselves accordingly. In contrast, if the BGP AD functionality is not used, you must then explicitly configure each PE router with the identities of all the other PEs in a specific VPLS.

You must assign a single, globally unique VPLS-ID to each VPLS (that is, the same value for all sites to the same VPLS across the entire network). The VPLS-ID eliminates the possibility of a collision between VPLSs belonging to different service providers.

There is also a globally unique Route Distinguisher (RD) ID associated with a VPLS. Each site also needs a unique ID that is a BGP NLRI. The PE address does not have to be globally routable, but it must be unique within the VPLS. The PE ID can be the PE router ID, for operational convenience.

Each site must also be associated with one or more RT Extended Communities, and the RTs control the distribution of NLRIs.

Each PE distributes the NLRI for each of its sites, with itself as the BGP next hop, and with the appropriate RT for each NLRI. A PE with a specific RT imports all NLRIs that have that same RT (and learns the other PEs addresses through their next hops). H-VPLS can be configured by using multiple RTs.

In summary, the BGP advertisement for a specific site in a PE includes:

- An NLRI. This is the VSI-ID.
- A BGP next hop equal to the loopback address of the PE.
- An extended community attribute containing the VPLS-ID.
- An extended community attribute containing one or more RTs.

Targeted-LDP [(T-)LDP] signaling is set up for the point-to-point PWs between sites using the selected (T-)LDP sessions corresponding to the remote PE(s) that have been recently added to their list.

To auto-create a spoke SDP binding, you must create a PW template policy and distribute the policy to the NEs. The NFM-P policy distribution mechanism is used to send out the template and maintain consistency in the network. The template selection is at the PE level, not at the service level, because not all PEs are capable of supporting BGP AD and some VPLS sites may not support BGP discovery.

Consider the following with regard to tunnel creation:

- If you plan to use BGP AD for all or part of the VPLS, you must not enable automatic mesh SDP binding creation.
- A provisioned PW to a specific remote PE takes precedence over one that is auto-discovered using BGP AD. In other words, if there is an existing SDP binding available, the router selects this existing binding and does not automatically create a new one.
- When the NFM-P auto-tunnel creation function is being used for non-BGP AD VPLSs, the automatically-created components are excluded from discovery. Also, you cannot select an automatically-created SDP or SDP binding when you create a service tunnel that is to be used as part of a BGP AD VPLS.

An NFM-P-managed VPLS or H-VPLS may consist of various site types located on different PEs, and which in turn, may be of differing versions. Due to these possibilities, some restrictions apply in terms of using Auto Discovery. For example, an M-VPLS site cannot be in the same service with a regular site. However, BGP AD-enabled sites and regular sites can be components of the same VPLS.

77.2.8 BGP VPLS

BGP VPLS is an extension of the VPLS concept. When configured as a BGP VPLS, such a service can interconnect with another BGP VPLS across different VPLS domains.

The control plane of the BGP VPLS provides auto-discovery and signaling capability. In this context, auto-discovery is a means for a PE router to discover other remote PE routers that are members of a given VPLS. The signaling function enables a PE router to know which pseudowire label a given remote PE router will use when sending the data to the local PE router. The BGP VPLS control plane carries sufficient information to provide the auto-discovery and signaling functions concurrently.

Some of the major features of the Nokia BGP VPLS solution include:

- The data plane is identical with the BGP AD (LDP VPLS) solution. For example, VPLS instances are interconnected via a pseudowire mesh. Split horizon groups may be used for loop avoidance between the pseudowires.
- Addressing is based on a two-byte VE ID assigned to the VPLS instance.
- The target VPLS instance is identified by the Route Target (RT) contained in the MP-BGP advertisement (extended community attribute).
- Auto-discovery is MP-BGP based.
- Pseudowire label signaling is MP-BGP based. As a result, the BGP NLRI content also includes label related information such as block offset, block size, label base, and so forth,

The Nokia BGP VPLS solution is compliant with RFC 4761.

See the *7750 SR OS Services Guide* for more information regarding BGP VPLS.

77.2.9 BGP VPLS Multi-homing

The NFM-P allows BGP VPLS multi-homing to be established for CEs and access PEs by first configuring the multi-homing VPLS sites and then assigning the same multi-homing site ID to each.

The following figure shows an example of this approach, where a VPLS contains certain CEs that are multi-homed to pairs of VPLS PEs.

Figure 77-7 CE Multi-homing in VPLS

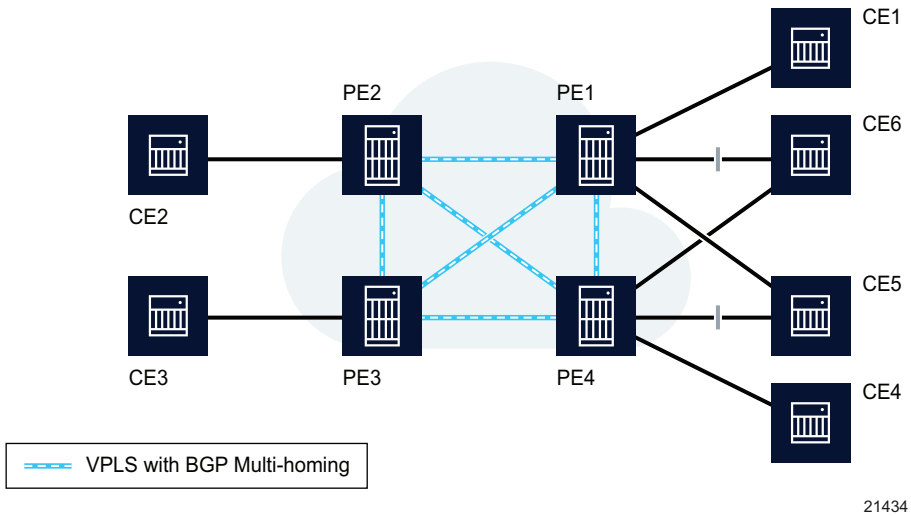

The BGP/LDP-signaled PW infrastructure (shown as the cloud at the center) is used to interconnect the VSIs between PEs. In this example, CE5 and CE6 are dual-homed to PE1 and PE4. To avoid loops, only one SAP must be active at any point in time between any multi-homed CE (such as CE5 or CE6) and its pair of connected PEs (such as PE1 and PE4). The others are blocked. Service providers use their MP-BGP on PE1 and PE4 to control the activation of the SAPs connected to the same customer site.

Other CE topologies (for instance, square connectivity) where, for example, CE1 and CE4 are part of the same customer site (and are themselves interconnected) are also supported.

When multi-homing a VPLS site using BGP (potentially into different autonomous systems), the PE routers (for example, PE1 and PE4) that are connected to the same customer site (for example, CE5) are configured with the same multi-homing ID. In this way, a loop-free topology is constructed using a routing mechanism such as BGP path selection. When a BGP speaker receives two equivalent NLRIs, it applies standard path selection criteria such as local preference and AS path length to determine which NLRI to choose.

Two VPLS NLRIs are considered equivalent from a path selection perspective if the following are identical:

- Route distinguisher
- Multi-homing ID

77.2.10 MVR on VPLS

MVR on VPLS is a bandwidth optimization method for applications on a broadband services network. At the port level, MVR allows a VPLS end user to subscribe or unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances without requiring that the stream be part of the customer VPLS.

MVR on VPLS is a mechanism through which the supporting devices are able to participate in a multicast distribution system. Separate, dedicated VLANs must be constructed specifically for multicast traffic distribution.

MVR assumes that hosts join and leave multicast streams by sending IGMP join and leave messages. The IGMP join and leave messages are sent inside the VPLS to which the host port is assigned. The multicast VPLS is shared in the network while the hosts remain in separate VLANs. For example, two user VPLS that are bound to the same MVR VPLS cannot exchange any information, but the same multicast service can still be provided to them.

An MVR VPLS is a VPLS that is responsible for sending multicast traffic through the network. An MVR VPLS is associated with a multicast package policy and has MVR-enabled sites. The MVR VPLS is configured to distribute certain multicast streams. An MVR VPLS can also be configured as a user VPLS to receive multicast traffic.

A user VPLS is a VPLS that contains SAPs that can receive multicast traffic from an MVR VPLS. Each SAP must be configured individually to use a specific MVR VPLS. Any VPLS, including an MVR VPLS, can be used as a multicast receiver for an MVR VPLS. IGMP and/or MLD snooping must be enabled on each site.

The following figure shows an example of MVR on VPLS. MVR reacts only to join and leave IGMP messages from the multicast groups configured for the MVR VPLS with which the user VPLS is

associated. Join and leave messages from all other multicast groups are managed by IGMP and/or MLD snooping. Therefore, several MVR VPLS instances can be configured, each with its own set of multicast channels.

Figure 77-8 MVR on VPLS

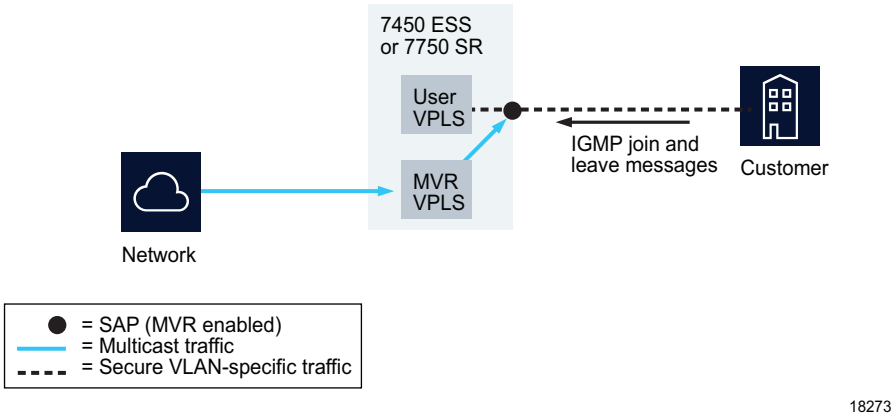

In some situations, such as when a host is connected to a 7301 ASAM, the multicast traffic cannot be sent from the MVR VPLS to the VLAN on which the IGMP message was received (standard MVR behavior) but to another VLAN. This configuration is known as MVR by proxy. The 7450 ESS, 7750 SR, and 7950 XRS allow multicast traffic to be sent to a SAP other than the SAP from which the IGMP message originated. When configuring MVR by proxy, you must indicate the MVR VPLS on which the multicast channel is available and the SAP to which the multicast traffic must be copied. The following figure shows an example of MVR by proxy.

Figure 77-9 MVR by proxy

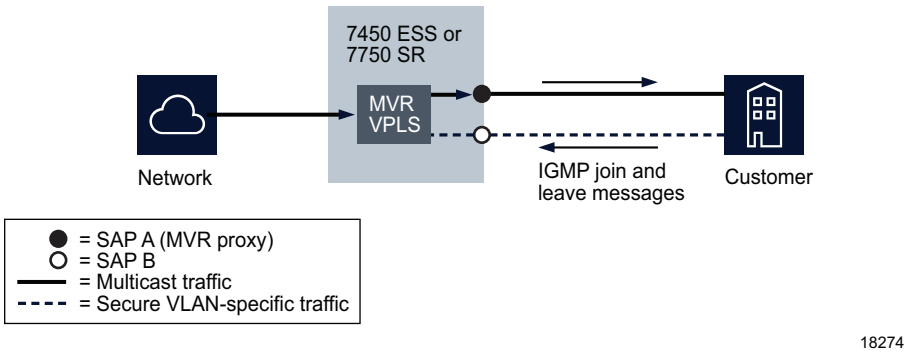

Configuring MVR on VPLS using the NFM-P involves the following steps.

1. Configure PIM before configuring MVR. See Chapter 28, “Routing protocol configuration” for more information.

2. Create a multicast package policy for all supporting NEs in an MVR VPLS. See Chapter 52, “Multicast policies” for more information.

   Alternatively, you can configure NEs individually in an MVR VPLS using a routing policy statement. See 54.5 “To configure a routing policy statement” (p. 1745) for more information.

3. Create the MVR VPLS. See 77.48 “To configure MVR for a VPLS site” (p. 2308) for more information.
   - Associate the multicast package policy with the service, which indicates that a VPLS is an MVR VPLS.
   - Specify the VPLS sites that are the sources of the multicast groups. MVR must be enabled on each site.
   - Configure SAPs only if you are configuring the MVR VPLS to be a user VPLS as well. See 77.5 “To create a VPLS” (p. 2249) for more information about configuring a user VPLS.

4. Create the user VPLS. See 77.5 “To create a VPLS” (p. 2249) for more information.
   - The SAPs are standard host access points.
   - Configure IGMP and/or MLD snooping and MVR on each site.
   - Associate the MVR VPLS with each SAP for which access to multicast traffic is needed. This association means that IGMP requests received at that SAP for a multicast group are fulfilled as long as the multicast group being requested is included in the multicast package policy of the MVR VPLS. After the SAPs in a user VPLS have been associated with a specific MVR VPLS, the SAP becomes known to the MVR VPLS.

5. If using MVR by proxy, configure a VPLS SAP that is to act as the MVR proxy. See 77.48 “To configure MVR for a VPLS site” (p. 2308) for more information.

In a situation in which an MVR VPLS has VPLS sites that do not support MVR, the following conditions apply:

- The ability to configure MVR is not available for the VPLS sites and SAPs of a device that does not support MVR.
- Multicast package policies are not distributed to the devices that do not support MVR.

77.2.11 GSMP group on VPLS

The edge devices determine the circuit that opens an ANCP session. ANCP provides status and control information such as current line rate and port-up and port-down messages to the edge devices. The edge devices perform the following functions:

- adjust the H-QoS subscriber scheduler with the correct rate
- raise an alarm when the rate goes below a set threshold
- send DSL line OAM commands to complete OAM tests

A GSMP group is created under the GSMP tab of the VPLS site form. Multiple groups can be defined and different ANCP capabilities can be associated with different groups. A neighbor can be defined in a GSMP group. Multiple neighbors can be configured for each group.

Note: A GSMP group must be configured on VPLS, MVPLS or VPRN for an ANCP session to open.

77.2.12 L2 management interfaces on VPLS

L2 management interfaces act as a host. L2 management interfaces are created the same way out-of-band interfaces are created on VPLS. L2 management interfaces are used for CPM protocols such as telnet, SSH, SNMP, ping, and ANCP.

CPM filtering is used to limit access to L2 management interfaces.

77.2.13 Routed VPLS

A routed VPLS connector joins an L3 access interface within an IES or VPRN service context to a VPLS service on the same site. When an IES or VPRN IP interface is bound to a VPLS site name, the site name cannot be bound to another IP interface. While an IES or VPRN IP interface can only be bound to a single VPLS site, the service context containing the IP interface can have other IP interfaces bound to other VPLS sites. Both the IES or VPRN IP interface and VPLS site must be located on the same NE.

If a VPLS site name does not exist within the system, the binding between the IP interface and the VPLS site remains operationally down until a VPLS site name is assigned to the VPLS site. When an IP interface is bound to a VPLS site, the operational state of the RVPLS binding is dependent upon the operational state of the VPLS site, and whether the IP interface binding is enabled on the VPLS site.

This functionality is limited to supporting devices.

Note: You can create and manage a routed VPLS connector from the navigation tree on the Composite Service (Edit) form or from an IES or VPRN access interface form, routed VPLS path.

The routed VPLS binding will not be operationally up until the Enable IP interface binding parameter is set to true and the VPLS site is operationally up. See 77.33 “To configure a VPLS site” (p. 2294) for more information.

You can add a routed VPLS L3 access interface to an IGMP interface on an IES or VPRN service.

For supported 7210 SAS-K nodes, RVPLS uses an IES, VPRN L3 access interface, or spoke SDP binding on the same NE as the RVPLS connector.

77.2.14 FIBs

The FIB is the set of information that represents the best forwarding information for a destination. A FIB entry is analogous to a static MAC address, and every computer and network node has a MAC address that is hardware-encoded. In NFM-P, static MAC addresses can be also created on VPLS endpoint objects, access interfaces, and service circuits.

The edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. Devices perform MAC-address learning to reduce the amount of unknown destination MAC address flooding. The edge devices learn the source MAC addresses of the traffic arriving on their access and network ports. You can also specify and manage static MAC addresses using the FIB entries table.

Each device maintains a FIB for each VPLS instance. Learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating sites using the service. Unknown destination packets (i.e., the destination MAC

address has not been learned) are forwarded on all LSPs to the participating devices for that service until the router responds and the MAC address is learned by the device associated with that service.

Note: Each VPLS FIB entry consumes system resources. The devices allow you to set the maximum number of MAC entries allowed in a VPLS instance to prevent a VPLS instance from consuming a disproportionate amount of resources.

The size of the VPLS FIB can be configured with a low watermark and a high watermark expressed as a percentage of the total FIB size limit. If the actual FIB size grows above the configured high watermark percentage, an alarm is generated. If the FIB size falls below the configured low watermark percentage, the alarm is cleared.

77.2.15 MAC learning

Like an L2 device, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). In each VPLS, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the FIB. A local MAC address is a MAC address associated with an access interface, because it ingressed on a SAP. A remote MAC address is a MAC address received using a service tunnel from another device that is part of the VPLS.

Unknown MAC discard is a feature which discards all packets ingressing the service where the destination MAC address is not in the FIB. The normal behavior is to flood these packets to all end points in the service.

MAC learning can also occur for Split Horizon Groups (SHG) but with accompanying risks. For example, in an L2 environment of an SHG, hosts (among whom mutual communication is disallowed) can launch DoS attacks by sending a flood of packets that source an uplink MAC address to unprotected customer SAPs.

This situation can be managed by controlling MAC learning on SAPs and SDPs. When a frame arrives at a protected SAP or SDP, the MAC is applied to its learning table; when a frame arrives at an unprotected customer SAP or SAP containing the address of a protected source MAC address, the system can be configured to take one of the following actions:

- The frame is immediately dropped and not learned by the unprotected SAP. As a result, the unprotected SAP does not know the MAC address of the uplink and, therefore, cannot use it to flood packets to other SAPs in the SHG.
- The frame is not dropped, but an alarm is generated.
- The SAP is placed in an operational Down state.

In addition, you can specify automatic population of the MAC protect list with source MAC addresses learned on the SAP or SDP.

Auto Learn Mac Protect and Restrict Protected Source functionalities are supported on SAPs, spoke and mesh SDP bindings, split horizon groups, and endpoints under a VPLS service. The functionality is also supported on Service PW-Template policies, in order to allow its use with BGP auto discovery for spoke SDP bindings and split horizon groups.

77.2.16 MAC move

A sustained high MAC re-learn rate can be a sign of a loop somewhere in the VPLS topology. Typically, STP detects loops in the topology, but for those networks that do not run STP, the MAC move feature is an alternative way to protect the network against loops.

When enabled on a VPLS, MAC move monitors the re-learn rate of each MAC. If the rate exceeds the configured allowed limit, it disables the SAP on which the source MAC last arrived. The SAP can be disabled permanently or for a length of time that grows linearly with the number of times the SAP is disabled.

There is also the option of marking a SAP as non-blockable, which means that when the re-learn rate exceeds the limit, another SAP—one that is blockable—is disabled instead. When the MAC move parameter is set to blockable, ports can be blocked in a specific order depending on the number of times the re-learn rate exceeds the configured threshold period.

MAC move is configurable on VPLS SAPs and VPLS spoke SDPs. Blocking information for an object is displayed on the MAC move configuration form for the object. This information includes:

- the number of MAC learning retries that remain before blocking occurs
- the time that remains before the blocked object is unblocked
- the order of blocking, starting with tertiary, secondary and primary

77.2.17 MAC rewrite

With Layer 2 policy-based redirects, packets may be discarded by the far-end router when the packets arrive. This can occur if the far-end IP interface has a different MAC address than the IP interface that is reachable via normal forwarding.

To avoid the discards, you can specify a destination MAC overwrite address for VPLS and MPLS SAPs to override the address specified in the device routing table. When enabled, all Unicast packets have their destination MAC rewritten to an operator-configured value that is expected to be that of the far-end IP interface. Multicast and broadcast packets are unaffected. This feature is not supported on B-VPLS, I-VPLS, Capture SAP, or Routed VPLS. See 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340)

77.2.18 Flooding

Traffic that is normally flooded throughout the VPLS can be rate limited on SAP ingress through the use of Service Ingress QoS Policies. In a Service Ingress QoS Policy, individual queues can be defined per forwarding class to provide shaping of broadcast traffic, MAC multicast traffic and unknown destination MAC traffic. You can also specify how to classify frames.

Multiple services and service types can be configured on a port. VPLS spanning tree protocols are configured on a per-service site basis, not a per-port basis, thus, multiple instances of STP per site are supported. Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service.

The flooding mechanism and the way that the Interior Gateway Protocol (IGP) operates ensure that no packets are duplicated on any interface. If SAPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and duplicate packets can traverse the network. The STP is designed to prevent multiple SAPs from forwarding a packet into the VPLS

77.2.19 Spanning tree protocols

The NFM-P supports RSTP for VPLS instances and maintains support for legacy STP implementations. STP on the 7750 SR and 7950 XRS incorporates an optimized and compatible implementation of IEEE 802.1D which attempts to eliminate STP blocking of links in the core of the VPLS. STP on the 7210 SAS and 7450 ESS is used to guarantee that service tunnels are not blocked in any circumstance while not imposing artificial restrictions on the placement of the root bridge in the network. To provide this support, all mesh service tunnels are configured as root ports or designated ports.

RSTP, which is the default STP mode managed by the NFM-P, is compliant with IEEE standard 802.1D-2004. Other available STP types include an RSTP variant with 802.1w-2001 backward compatibility, an STP variant that is compliant with 802.1w-2001, and MSTP, an STP variant that is compliant with 802.1s-2002.

The NFM-P verifies STP parameters that are configured within each VPLS instance. However, it does not check the compatibility of STP configurations between interconnected VPLS instances.

77.2.20 IGMP snooping

IGMP snooping allows a device to snoop packets sent between IP multicast devices and IP multicast hosts to learn the IP multicast group membership. The device checks the IGMP packets for the group registration information, and configures multicasting accordingly.

Without IGMP snooping, multicast traffic is forwarded to all ports, which is the same as broadcast traffic. IGMP snooping ensures that multicast traffic is forwarded only to ports that are members of a specific multicast group to reduce the amount of multicast traffic that passes through the device.

You can enable IGMP snooping for VPLS sites, access interfaces, and spoke SDPs.

Note: IGMP snooping is not supported when MAC subnetting is enabled.

77.2.21 MLD snooping

Multicast Listener Discovery snooping is essentially the IPv6 version of IGMP snooping. The guidelines and procedures are very similar to IGMP snooping.

When MLD snooping is not enabled, L2 switches treat multicast traffic like an unknown MAC address or broadcast frame, that is, the frame is flooded out on every port of a VLAN. When MLD snooping is enabled, switches snoop the frame's L3 header for more efficient switching. In the context of IP multicast, only hosts that have expressed interest in receiving packets for the multicast groups have the frames forwarded to them.

The 7x50 router allows the enabling of MLD snooping for VPLS. A database of group members (per VPLS instance) is built by listening to MLD queries and reports from each SAP and SDP of the instance. These reports are forwarded to the multicast routers, if any are present.

MLD snooping is not supported when MAC subnetting is enabled.

Consider the following:

- Multicast groups can be learned (using the destination IP addresses of multicast packets) through MLD snooping or by static configuration at the port.
- The Fast leave feature modifies the membership leave mechanism by terminating the session

immediately, rather than issuing a group-specific query to check if other members are still present on the network. Therefore, if a port (SAP or SDP) is configured for Fast leave, the session is terminated immediately without checking if the port also has other hosts subscribed to that same multicast group.

- A multicast router retains a list of multicast group memberships for each attached network. Therefore, a multicast router can assume the role of querier or listener. However, there can be only one querier per physical network.
- MLD snooping statistics are collected for each NE port, SAP, or SDP binding, and are viewable from the Statistics tab of a VPLS site properties form.
- The NFM-P supports MLDv1 and MLDv2.
- The NFM-P supports MVR.
- MLD snooping can co-exist with IGMP snooping and PIM snooping.
- MAC-based forwarding entries can be built using MLD snooping results.

77.2.22 PIM snooping

Protocol Independent Multicast snooping for VPLS allows a VPLS PE router to build multicast states by snooping PIM protocol packets that are sent over the VPLS. The VPLS PE then forwards multicast traffic based on the multicast states.

When all receivers in a VPLS are IP multicast routers running PIM, multicast forwarding in the VPLS can be efficient when PIM snooping for VPLS is enabled. After PIM snooping is turned up at the service level, all sites have PIM snooping configured and are set to the down state, by default. If any site does not support this feature, the NFM-P treats it as having PIM snooping turned off.

Since PIM snooping operates on PIM Hello packets as well as Join/Prune packets, PIM neighbors of the current router are learned by snooping on Hello packets. Therefore, in a meshed VPLS, every node learns from every other node's Hello packet and considers that node as a neighbor.

VPLS PE routers only snoop on PIM Hello and Join/Prune packets; they do not generate PIM messages on their own. Therefore, when PIM snooping is configured at the service level, all CE routers must have Join/Prune suppression disabled. If a VPLS PE router detects a condition where Join/Prune suppression is not disabled on one or more CE routers, the PE router puts PIM snooping into a non-operational state for the entire service. A trap on the PE is generated to report this condition and an alarm is raised to the NFM-P operator. To bring PIM snooping back to the operational state, PIM snooping must be disabled and then re-enabled.

Since PIM uses state refreshes, VPLS PE routers may not learn multicast states from all the CE routers, if PIM snooping was just enabled or the all snooping state was just cleared, until the next refresh.

To avoid traffic interruption, PIM snooping should hold up its operations for a period of time (60 seconds, if default timer is used). During this period of time, multicast traffic is flooded in the VPLS just like snooping was not enabled. The NFM-P should have this hold-up timer configurable on VPLS properties panel having range 0-120 with 60 secs as default.

A variety of statistics are gathered for PIM snooping operation and are available for viewing and analysis in the PIM Snooping tab of a VPLS configuration form under the Multicast tab.

PIM snooping is not supported when MAC subnetting is enabled.

77.2.23 Split horizon groups

SHGs control traffic that flows through SAPs or spoke SDPs for a VPLS site. SHGs prevent a packet received on a SAP within the group from being propagated to other members of the group.

SHGs are defined when you create or modify a VPLS site. You can create multiple SHGs for a VPLS site. SHGs can support a mix of spoke SDPs and SAPs. When you create SAPs or spoke SDPs they can be associated with an SHG.

Users can:

- configure, modify, or delete SHGs on a VPLS site
- associate SAPs or spoke SDPs with SHGs

77.2.24 Residential split horizon groups

RSHGs are SHGs with the Residential parameter enabled. SAPs that are associated with RSHGs are called lightweight SAPs. RSHGs use dual-pass queue optimization and do not support downstream broadcast or multicast traffic.

Users can:

- configure, modify, or delete RSHGs on a VPLS site
- associate SAPs or spoke SDPs with RSHGs

Note: If a SAP or spoke SDP is associated with an RSHG, then the following apply:

- MAC pinning is enabled by default and cannot be disabled for the interface
- IGMP snooping, MLD snooping, and MVR are not configurable for the interface

77.2.25 QTag Manipulation

The NFM-P supports QTag Manipulation on supporting NEs. QTag Manipulation allows you to define actions for inner and outer VLAN IDs, and assign tag values where required. For example, you can configure QinQ Tunneling, VLAN Translation, or other manipulations for L2 traffic that ingresses and egresses the interface. QTag Manipulation is configured on the Port tab of the L2 Access Interface properties form. See the NE documentation for more information about QTag Manipulation actions.

77.2.26 Default SAPs

The NFM-P allows you to create default SAPs that you can use in an L2-based service to perform management tasks or to deliver a specific class of end-user service. One default SAP can be defined on any dot1q encapsulated Ethernet port. Up to four default SAPs can be defined on any Q in Q encapsulated Ethernet port.

You can create a default SAP by specifying an outer encapsulation value of 4095 or *. If the XML API is used, the outer encapsulation value is always 4095. You can also enable the Enable Q in Q Untagged Sap parameter on an NE which allows the creation of the following default SAPs for a Q in Q Ethernet port:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.

- The SAP type *.* functions as a default SAP for double-tagged frames on a Q in Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

The following figure shows a default SAP used as a dedicated management port.

Figure 77-10 Default SAP as a dedicated management port

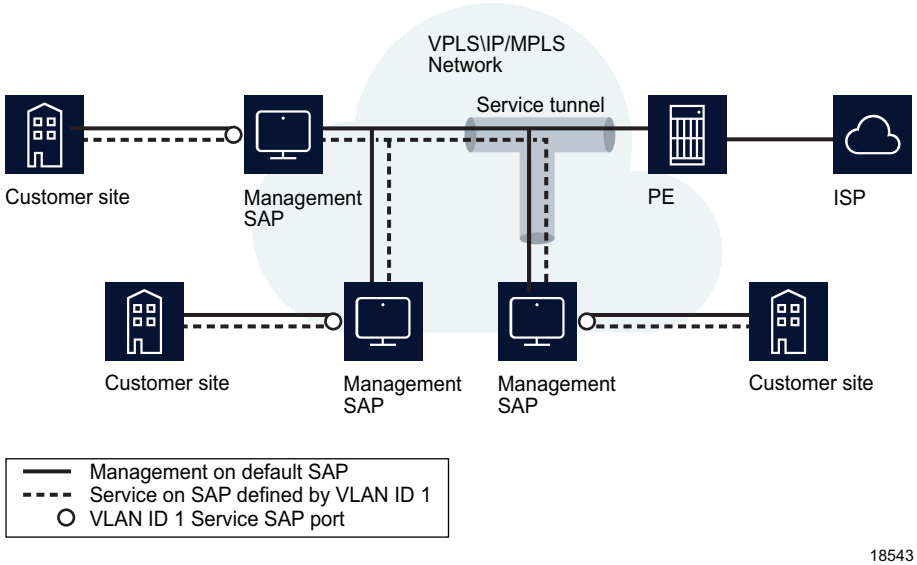

The dotted line shows that a business customer uses an entire port to access an L2 service using a VLAN tag ID 1, which is transparent to the network service provider. The service provider has assigned a default SAP for management of the customer network, as shown by the solid line. The customer uses one SAP for service delivery, and the service provider uses the default SAP to manage the customer network.

Default SAPs can also be used to differentiate one class of service from another on a single port. In the following figure the service provider can deliver aggregated high-speed Internet services on a single default SAP for multiple hosts while applying tags that assign one VLAN per host. At the same time, each of the remaining SAPs on the same port can be deployed for higher-level services, such as triple play delivery, in which multiple or individual hosts are assigned to a single SAP. Using such differentiated SAPs is efficient and significantly increases network scalability, because of the reduced number of SAPs allocated to various customers.

Figure 77-11 Default SAP to differentiate subscriber services

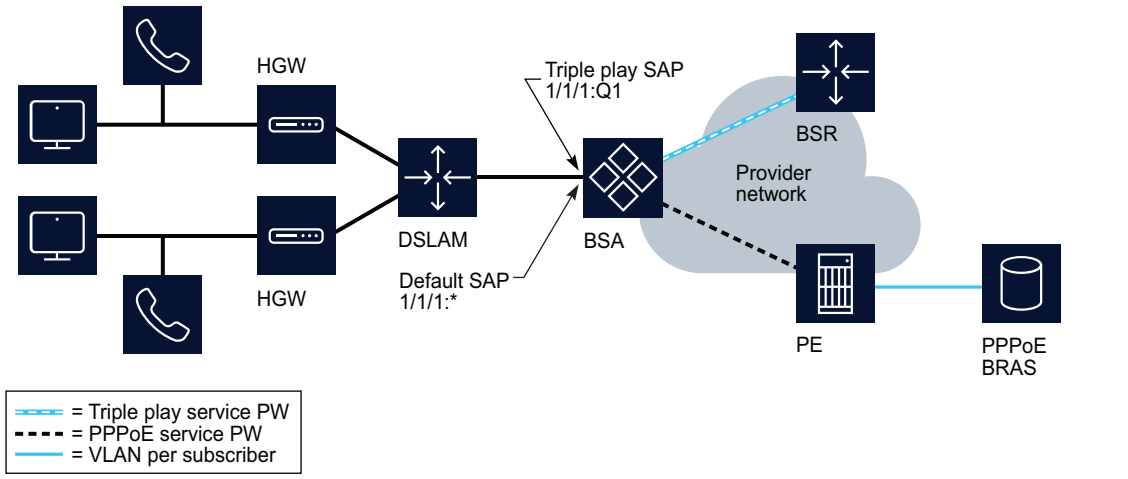

77.2.27 Layer 2 protocol tunneling termination

L2PT allows service providers to preserve the VLAN and Layer 2 protocol configurations of individual customers without impacting the traffic of other customers across the core network. L2PT termination allows Layer 2 PDUs to be transparently tunneled across the core network, avoiding interaction between the network provider and customer protocols.

Transparent L2PT is performed on the ingress side of every SAP or spoke SDP of the PE routers configured with L2PT termination. L2PT tunnels PDUs by overwriting the destination MAC address in an Ethernet frame using the multicast MAC address 01-00-0c-cd-cd-d0. The Ethernet frame is then transparently tunneled over the core network to a peer PE router. The peer PE router at the egress side of the tunnel restores the MAC address and the L2 protocol so that packets are forwarded to all ports in the same VLAN.

L2PT termination can only be enabled if STP is disabled on the VPLS.

77.2.28 BPDU translation

VPLS networks typically interconnect customer sites that use different access technologies, such as Ethernet and bridge-encapsulated ATM PVCs. Because of this, BPDU translation may be necessary to provide end-to-end interconnectivity.

If BPDU translation is enabled on a SAP or spoke SDP, the device intercepts all BPDUs and performs the required translation.

BPDU translation can be enabled only on a SAP or spoke SDP binding if STP is disabled on the VPLS.

77.2.29 DoS protection

To protect a VPLS from a high incoming packet rate that characterizes a DoS attack, you can use the NFM-P to create DoS protection policies for the VPLS L2 access interfaces. A DoS protection policy limits the number of control-plane packets that an interface receives each second, and optionally logs a violation notification if a policy limit is exceeded. You can use the NE System Security form to view the violations for a specific NE.

You can configure a DoS protection policy to control the following on a VPLS L2 access interface:

- the control-plane packet arrival rate per subscriber host on the interface
- the overall control-plane packet arrival rate for the interface
- whether an NE sends a notification trap if a policy limit is exceeded

Each VPLS L2 access interface on an NE that supports DoS protection is automatically assigned a default DoS protection policy. The default policy limits only the overall packet arrival rate for the interface, and cannot be deleted or modified. See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for information about creating a DoS protection policy.

77.2.30 DDoS protection on access interfaces

You can configure a DDoS protection policy on a VPLS L2 access interface, I-VPLS I-L2 access interface, MVPLS L2 access interface, or I-MVPLS I-L2 access interface. See the procedure to configure an NE DDoS protection policy in the *NSP System Administrator Guide* for more information.

77.2.31 Copying and moving SAPs between ports

You can copy and move SAPs between ports. See 16.6 “Copying and moving SAPs” (p. 574) in Chapter 16, “Port and channel object configuration” for more information.

77.3 Sample VPLS configuration

77.3.1 Sample VPLS

The following figure shows a sample VPLS configuration.

Figure 77-12 Sample VPLS

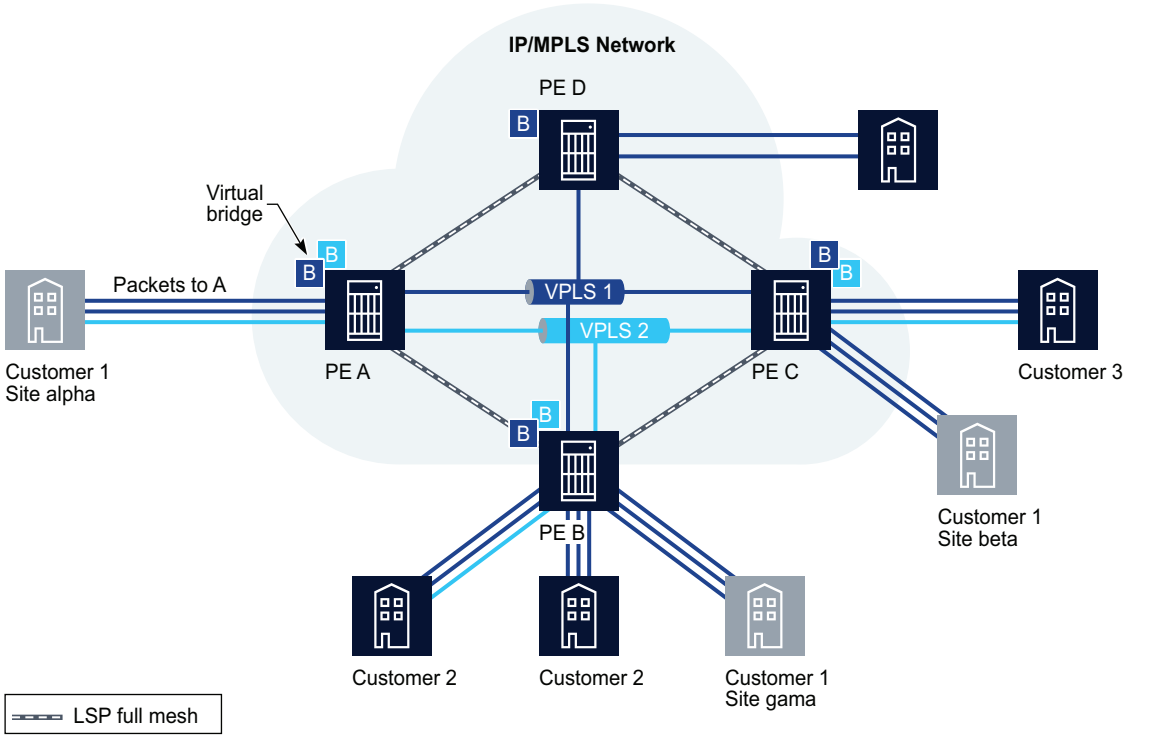

VPLS 1 is a distributed service, which consists of customer 1 connected to PE A, PE B, and PE C, at sites alpha, gamma, and beta, respectively. All three customer sites belong to VPLS 1.

In the following example, Customer 1 wants to send data from site alpha to site beta.

Customer 1 packets arriving at PE A are associated to the appropriate VPLS 1 for that customer, based on the combination of the access port and the dot1q (VLAN ID) in the packet. PE A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the access port on which it was received.

PE A is sending the packets to PE C. The destination MAC address in the packet is looked up in the FIB table of PE A for the VPLS instance, as shown in the following figure.

Figure 77-13 Packet forwarding by ingress router PE A

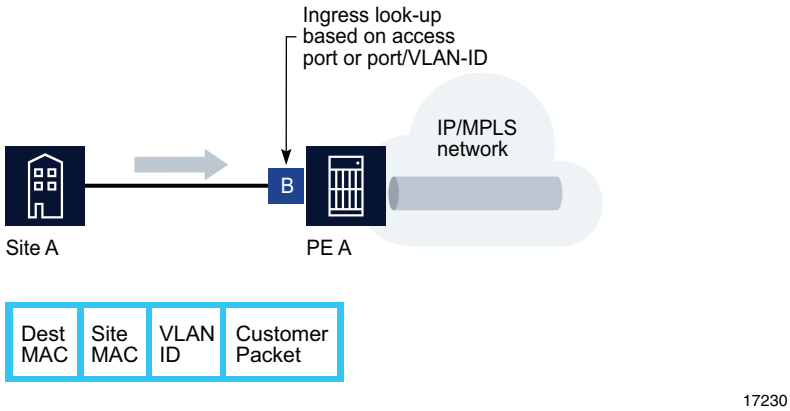

The destination MAC address in the FIB table has one of two values:

- known
- unknown

If the destination MAC address is known by PE A, an existing entry in the FIB table identifies the far-end PE C and the service VC label (VLAN ID) used to send the packets from PE A to PE C. PE A chooses a transport LSP to send the packets to PE C. The packets from the customer 1 site alpha to site beta are sent on the LSP after the VC label is removed and the transport label is added to the packet, as shown in Figure 77-13, “Packet forwarding by ingress router PE A” (p. 2239) .

If the destination MAC address is not known by PE A, PE A floods packets to both PE B and PE C, which are both part of VPLS 1. PE A uses the VC labels (VLAN IDs) that PE B and PE C previously signaled for this VPLS 1.

As shown in the following figure , the packets from PE A are transported across the core IP/MPLS network.

Figure 77-14 Packet forwarding from PE A across the core IP/MPLS network

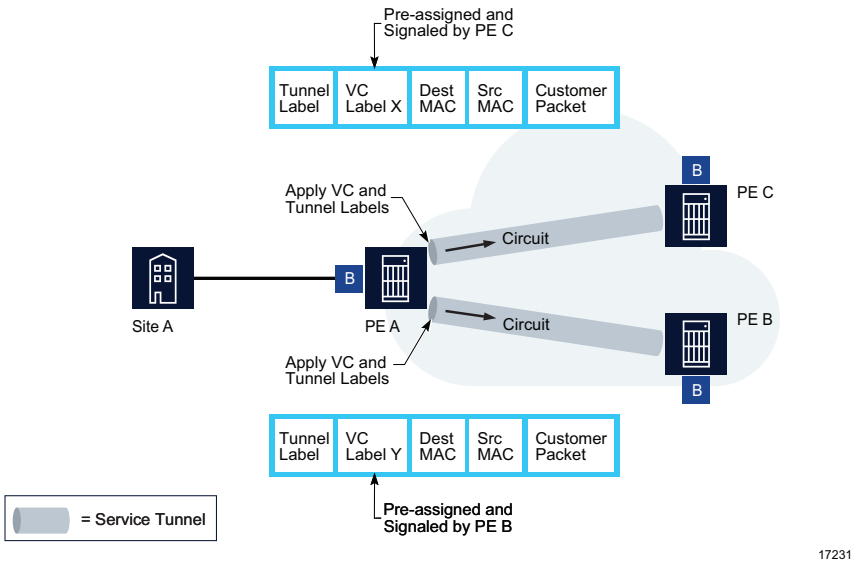

The core routers are LSRs that switch the packets towards their destination based on the tunnel label, also called a transport label. The core routers are not aware that the packets belong to a VPLS.

When the packets from PE A arrive at PE C, PE C removes the tunnel label to reveal the VC label that associates the packets with VPLS 1, as shown in the following figure.

Figure 77-15 Packet forwarding by egress router PE C

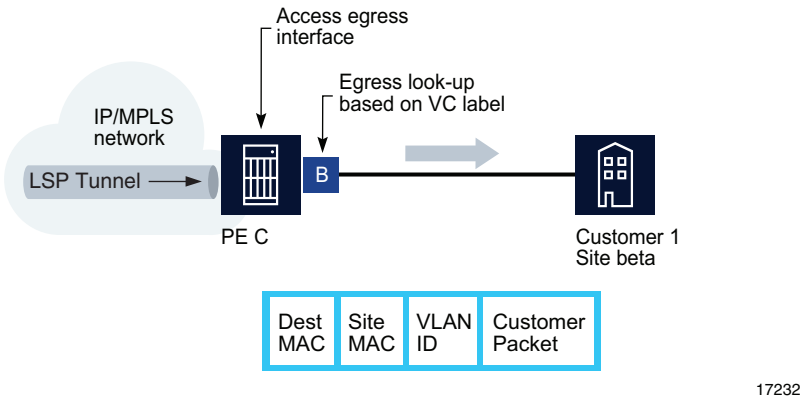

PE C learns the source MAC address of PE A and creates an entry in its FIB table that associates the MAC address and the VC label with PE A. The destination MAC address is looked up in the FIB table. It has one of two values:

- known
- unknown

If the destination MAC address is known by PE C, an existing entry in the FIB table identifies the local access (egress SAP) port used by VPLS 1 site beta and the service VC label (VLAN ID) that needs to be added to send the packets from PE C to customer 1.

If the destination MAC address is not known by PE C, PE C floods packets to all local access ports that belong to VPLS 1.

77.3.2 Configuration steps

Assuming the core IP/MPLS network has already been configured, the following high-level tasks are required to configure this sample VPLS.

1

Configure policies as required, for example, access ingress and access egress, Scheduler, ACL MAC, Accounting, ANCP, DoS protection, and DDoS protection.

2

Configure ports as access ports for use in the service.

3

Configure service tunnels, as required.

4

Create and configure Customer 1.

5

Create and configure VPLS 1.

6

Create, update, or configure additional sites or L2 access interfaces for the VPLS.

7

For an HVPLS, use spoke SDP bindings to connect VPLS sites to other sites in the same VPLS, to sites in a different VPLS, or to VLL service sites.

8

Create a single MVPLS.

9

Create a composite MVPLS to manage traffic blocking for multiple VPLS.

VPLS management procedures

77.4 Workflow to create a VPLS

77.4.1 Overview

The following workflow lists the high-level steps required to create a VPLS. As a prerequisite for creating a VPLS, this workflow assumes the following:

- a group or customer with the required user access privileges has been set up.
- the IP or IP/MPLS core network exists.
- any required service tunnels are created including the static, dynamic or SR-TE LSP required to create the service tunnel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
- the access ports for the service are created. See Chapter 16, “Port and channel object configuration” for more information.
- any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite policies are created. See Chapter 49, “Policies overview” for more information. You do not have to create pre-defined policies if policies are created on a per-service basis.
- any required MP-BGP for PE-to-PE routing is configured. See Chapter 28, “Routing protocol configuration” for more information about protocol configuration.

77.4.2 Stages

1

Create and configure a VPLS. See 77.5 “To create a VPLS” (p. 2249) .

1. Define the service type as VPLS.
2. Ensure the LSP network is configured when the transport mechanism is MPLS.
3. Configure HVPLS spoke redundancy.

2

Specify the sites for the service. See 77.33 “To configure a VPLS site” (p. 2294) .

1. As required, perform the following for the VPLS site. See 77.35 “To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site” (p. 2296) .
   - specify STP parameters
   - create MSTP instances and associate VLANs
   - specify protected MAC addresses
2. Configure a default gateway for the site. See 77.37 “To configure a default gateway for a VPLS site” (p. 2299) .
3. Configure ingress multicast forwarding for the site. See 77.38 “To configure ingress multicast forwarding on a VPLS site” (p. 2299) .
4. Configure a provider tunnel for the site. See 77.39 “To configure a provider tunnel for a VPLS site” (p. 2300) .

5. Specify the service tunnel required bandwidth for the site. See 77.40 “To configure service tunnel required bandwidth for a VPLS site” (p. 2301) .
6. Configure IGMP snooping on the site. See 77.41 “To configure IGMP snooping on a VPLS site” (p. 2302) .
7. Configure PIM snooping on the site. See 77.42 “To configure PIM snooping on a VPLS site” (p. 2303) .
8. Configure the split horizon group parameters. See 77.44 “To configure an SHG on a VPLS site” (p. 2304) .
9. Configure the EVPN gateway parameters. See 77.45 “To configure an EVPN gateway on a VPLS site” (p. 2305) .
10. Configure MVR. See 77.48 “To configure MVR for a VPLS site” (p. 2308) .
11. Configure GSMP group and GSMP group neighbor parameters. See 77.49 “To configure a GSMP group on a VPLS site” (p. 2309) .
12. Create L2 management interfaces on the site. See 77.50 “To configure L2 management interfaces on a VPLS site” (p. 2310) .
13. Configure MLD snooping on the site. See 77.51 “To configure MLD snooping on a VPLS site” (p. 2311) .
14. Associate a virtual MEP to the site. See 77.52 “To create a Virtual MEP on a VPLS site” (p. 2312) .
15. Configure MVR for MLD on the site. See 77.53 “To configure MVR for MLD on a VPLS site” (p. 2313) .
16. Configure IGMP host tracking on the site. See 77.54 “To configure IGMP host tracking on a VPLS site” (p. 2314) .
17. Configure WLAN GW L2 wholesale forwarding on a site. See 77.55 “To configure WLAN GW L2 wholesale forwarding on a VPLS site” (p. 2315) .
18. Configure a non-system IP termination for VXLAN. See 77.56 “To configure a non-system IP address VXLAN termination” (p. 2316).
19. Configure EVPN VXLAN on a site. See 77.57 “To configure EVPN on a VPLS site” (p. 2317) .
20. Configure PBB-EVPN on the site. See 77.59 “To configure PBB-EVPN on a VPLS site” (p. 2320) .
21. Configure a black hole MAC address on the site. See 77.60 “To configure a black hole MAC address on a VPLS site” (p. 2321) .
22. Enable SPB on a control B-VPLS site. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) .
23. Connect an Ethernet ring to the service. See 77.22 “To connect a G.8032 Ethernet ring to a VPLS” (p. 2280) .

3

Create L2 access interfaces for the VPLS sites, as required. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) .

1. Specify aggregation on a service basis, or across a card or port.
2. Configure MSAP policies. Create an MSAP Policy to control how the parameters are applied to an MSAP when it is automatically created. See 64.10 “To configure an MSAP policy” (p. 1852) for more information.
3. Configure the SAP Sub Type as Regular or Capture, as required. Regular is the default value used for the creation of a SAP and Capture is the value used to enable the automatic creation of an MSAP. That is, you create a Capture SAP to enable the creation of an MSAP. See 74.26 “To configure a capture SAP” (p. 2049) for more information about creating a Capture SAP.
4. Assign QoS, scheduling, accounting, and filter policies. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) .
5. Specify MAC ACL filters. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) .
6. Assign a time of day suite. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) .
7. Configure queue override parameters. See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) .
8. Configure the BPDU Termination, STP, and FIB parameters. See 77.72 “To configure BPDU Termination, STP, and FIB parameters for the VPLS L2 access interface” (p. 2348) .
9. Assign a DoS protection policy or DDoS protection policy. See 77.73 “To assign a DoS protection policy or DDoS protection policy to the VPLS L2 access interface” (p. 2350) .
10. Configure subscriber management. See 77.74 “To configure residential subscriber management for the VPLS L2 access interface” (p. 2351) .
11. Configure Ethernet tunnels. See 77.75 “To configure an Ethernet tunnel on a VPLS L2 access interface” (p. 2352) .
12. Configure a redundant VLAN range. See 77.76 “To configure a redundant VLAN range on a VPLS L2 access interface” (p. 2353) .
13. Configure IGMP snooping parameters. See 77.77 “To configure IGMP snooping for a VPLS L2 access interface” (p. 2354) .
14. Configure ARP hosts. See 77.78 “To configure the ARP host for the VPLS L2 access interface” (p. 2355) .
15. Configure DHCP relay parameters. See 77.79 “To configure DHCP for the VPLS L2 access interface” (p. 2356) .
16. Configure MVR. See 77.80 “To configure MVR for a VPLS L2 access interface” (p. 2357) .
17. Configure anti-spoofing parameters. See 77.81 “To configure anti-spoofing filters for a VPLS L2 access interface” (p. 2358) .
18. Configure DHCPv6 snooping parameters. See 77.90 “To configure DHCPv6 snooping for a VPLS or MVPLS L2 access interface” (p. 2381) .

19. Create a MIP and a MEP on the interface. See 77.82 “To create MIPs and MEPs on a VPLS L2 access interface” (p. 2360) .
20. Assign an ANCP policy. See 77.83 “To assign an ANCP policy to a VPLS L2 access interface” (p. 2362) .
21. Configure PIM snooping parameters. See 77.84 “To configure PIM snooping on a VPLS L2 access interface” (p. 2363) .
22. Configure MLD snooping parameters. See 77.85 “To configure MLD snooping for a VPLS L2 access interface” (p. 2364) .
23. Configure MVR for MLD. See 77.86 “To configure MVR (MLD) for a VPLS L2 access interface” (p. 2365) .
24. Configure custom object attributes for AA reporting. See 77.23 “To configure custom object attributes for AA reporting” (p. 2282) .

4

Create endpoints on the site for redundant configuration. See 77.43 “To create an endpoint for redundancy (dual homing) on a VPLS site” (p. 2303) .

5

Create a spoke SDP binding under any endpoints you created on the site. See 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

6

Create a mesh SDP binding for the VPLS site. See 77.91 “To create a VPLS or MVPLS mesh SDP binding” (p. 2382) .

7

Create a spoke SDP binding for the VPLS site. See 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

8

Configure the following on the SDP binding, as required.

1. Assign a DoS protection policy. See 77.94 “To assign a DoS protection policy to a VPLS SDP binding” (p. 2393) .
2. Configure the MFIB, STP, and FIB parameters. See 77.105 “To configure learning protection parameters on a VPLS SDP binding” (p. 2403) .
3. Configure DHCP for the SDB binding. See 77.95 “To configure DHCP for the VPLS SDP binding” (p. 2394) .
4. Configure IGMP snooping parameters. See 77.96 “To configure IGMP snooping for the VPLS SDP binding” (p. 2395) .
5. Configure DHCPv6 snooping parameters. See 77.108 “To configure DHCPv6 snooping for the VPLS or MVPLS SDP binding” (p. 2407) .

6. Create a MIP on the SDP binding. See 77.98 “To create a MIP on a VPLS SDP binding” (p. 2396) .
7. Create a MEP on the SDP binding. See 77.99 “To create a MEP on a VPLS SDP binding” (p. 2398) .
8. Configure MLD snooping parameters. See 77.100 “To configure MLD Snooping for the VPLS SDP binding” (p. 2399) .
9. Configure PIM snooping parameters. See 77.104 “To configure PIM snooping for a VPLS spoke SDP binding” (p. 2403) .
10. Configure custom object attributes for AA reporting. See 77.106 “To configure custom object attributes for AA reporting for a spoke SDP binding” (p. 2405) .

9

Configure a site for BGP AD or BGP VPLS, as required. See 77.110 “To configure a site for BGP AD or BGP VPLS” (p. 2408) .

1. Create a routing policy statement to define the required community members. See 54.5 “To configure a routing policy statement” (p. 1745) .
2. Enable BGP on the routing instance of each NE in the VPLS or BGP VPLS. See 28.29 “To enable BGP on a routing instance” (p. 916) .
3. Configure global-level BGP on each NE in the VPLS or BGP VPLS. See 28.31 “To configure global-level BGP” (p. 918) .
4. Create a PW template policy. See Chapter 83, “Service PW template policies” for information about PW template policies.
5. Distribute the PW template policy to each NE that is or will be a component of the VPLS or BGP VPLS.
6. Configure a site for BGP VPLS Multi-homing, as required. See 77.111 “To configure a site for BGP VPLS Multi-homing” (p. 2412) .
7. Configure a site for EVPN VXLAN, as required. See 77.57 “To configure EVPN on a VPLS site” (p. 2317)
8. Configure a site for PBB-EVPN, as required. See 77.59 “To configure PBB-EVPN on a VPLS site” (p. 2320) .
9. Configure ethernet segments for a PBB-EVPN, as required. See 77.17 “To configure an Ethernet segment” (p. 2275) .
10. Create SDPs for the BGP VPLS Multi-homing site(s), if manually-provisioned service tunnels are required. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
11. Re-evaluate the PW template policies associated with a BGP AD or BGP VPLS site when a policy is modified. See 77.112 “To re-evaluate the PW Templates associated with a BGP AD or BGP VPLS” (p. 2416) and 83.5 “To reevaluate a PW template policy on a local definition after a configuration change” (p. 2747) .

10

Perform one or more of the following.

a. Create an HVPLS. See 77.6 “To create an HVPLS” (p. 2249) .

b. Create an MVPLS. See 77.7 “To create an MVPLS” (p. 2251) .

c. Create a B-site for VPLS or MVPLS. See 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) .

d. Create an I-VPLS. See 77.8 “To create an I-VPLS” (p. 2253) .

e. Create a VPLS or MVPLS B-L2 access interface. See 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) .

f. Create a VPLS I-L2 access interface. See 77.88 “To create a VPLS I-L2 access interface” (p. 2372) .

g. Configure a GNE site on a VPLS service. See 77.34 “To configure a GNE site on a VPLS service” (p. 2295) .

h. Force a switchover to a redundant spoke SDP binding. See 77.107 “To force a switchover to a redundant spoke SDP binding” (p. 2406) .

i. Run an OAM validation test for a VPLS. See 77.28 “To run a VPLS service OAM validation test” (p. 2289) .

11

As required, modify the FIB entries associated with a VPLS:

a. Add or modify FIB entries. See 77.29 “To add or modify FIB entries associated with a VPLS” (p. 2290) .

b. List FIB entries. See 77.30 “To list FIB entries associated with a VPLS” (p. 2291) .

12

As required, view VPLS information:

a. View the service topology map associated with a VPLS. See 77.12 “To view the service topology associated with a VPLS ” (p. 2269) .

b. View the VPLS operational status. See 77.15 “To view the VPLS operational status” (p. 2274) .

c. View IGMP snooping queriers. See 77.31 “To view IGMP snooping queriers” (p. 2292) .

d. View MLD snooping queriers. See 77.32 “To view MLD snooping queriers” (p. 2292) .

13

As required, modify a VPLS:

a. Using the Show Info form. See 77.10 “To view VPLS contents” (p. 2260) .

b. Using the Manage Services form. See 77.9 “To modify a VPLS” (p. 2259) .

c. Using the topology view. See 77.11 “To modify a VLPS using the topology view” (p. 2262) .

14

As required, delete a VPLS. See 77.13 “To delete a VPLS” (p. 2269) .

77.5 To create a VPLS

77.5.1 Steps

1

Choose Create→Service→VPLS from the NFM-P main menu. The VPLS Service (Create) form opens.

2

Select a customer to associate with the VPLS in the Customer panel.

3

Configure the required general service parameters.

The Service ID and SVC Mgr Service ID parameters are configurable when the Auto-Assign ID parameter is disabled.

The Default Mesh VC ID parameter is configurable when the Inherit Service ID Value parameter is disabled.

4

If you enabled the Automatic Mesh SDP Binding Creation parameter in Step 3 , select a tunnel selection profile in the Automatic Mesh SDP Binding Creation Preferences panel.

5

Save the changes and close the form.

END OF STEPS

77.6 To create an HVPLS

77.6.1 Overview

An HVPLS consists of a VPLS in which one or more sites connect to other sites in the same VPLS, different VPLS, or to VLL services.

Note: One VPLS site in an HVPLS must be configured with an SHG. See 77.2.23 “Split horizon groups” (p. 2234) in this chapter for more information.

77.6.2 Steps

1

CAUTION

Service Disruption

If you are creating an HVPLS that includes two sites in the same VPLS connected by spoke SDPs, do not create mesh SDP bindings between the sites. Mesh SDP binding functionality is available in the spoke SDP bindings between the sites.

Add a VPLS to the HVPLS.

a. Create a new VPLS.

1. Perform 77.5 “To create a VPLS” (p. 2249) .
2. Go to Step 2 .

b. Use an existing VPLS. Go to Step 3 .

2

Create another service for inclusion in the HVPLS.

a. Create a VPLS.

1. Perform 77.5 “To create a VPLS” (p. 2249) .
2. Go to Step 3 .

b. Create a VLL.

1. Perform the appropriate VLL creation procedure in Chapter 76, “VLL service management” .
2. Go to Step 3.

3

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

4

Choose a VPLS that you want to include in the HVPLS and click OK. The VPLS Service (Edit) form opens.

5

On the service tree, expand Sites→site and right-click on Spoke SDP Bindings below the site and choose Create Spoke SDP Bindings. The Spoke SDP Binding (Create) form opens.

6

Perform Step 8 to Step 34 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386). Choose the destination NE for the site that you want to include in the HVPLS as the Tunnel Termination Site.

7

Perform 77.33 “To configure a VPLS site” (p. 2294) to add sites to the HVPLS.

END OF STEPS

77.7 To create an MVPLS

77.7.1 Purpose

An MVPLS runs RSTP or MSTP to manage traffic blocking on the associated VPLS. Perform this procedure to create an MVPLS to run MSTP, or to run RSTP and manage traffic on the associated VPLS SAPs or redundant spoke SDPs. The procedure also applies to the I-Sites and B-Sites used in PBB.

77.7.2 Steps

1

Choose Create→Service→MVPLS from the NFM-P main menu. The MVPLS service (Create) form opens.

2

Select a customer to associate with the MVPLS.

3

Configure the required parameters.

4

To configure service tunnel required bandwidth for the service:

Note: To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

1. Click on the Bandwidth tab.
2. Configure the Bandwidth Method parameter.
3. If the Bandwidth Method parameter is set to Manual or Calculate, configure LSP Path Booking parameter.
4. If the Bandwidth Method parameter is set to Manual, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

5

On the service tree, right-click on Sites under MVPLS and choose one of the following:

- Create MVPLS B-Site
- Create MVPLS I-Site
- Create MVPLS Site

6

Select a site. Depending on your selection in Step 5 , the MVPLS B-Site (Create), MVPLS I-Site (Create), or MVPLS Site (Create) form opens.

7

Perform one of the following:

a. Create an RSTP site for the MVPLS.

1. For regular VPLS sites, perform Step 4 to Step 5 of 77.33 “To configure a VPLS site” (p. 2294) . Specify RSTP as the value for the STP Mode parameter in Step 6 of 77.35 “To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site” (p. 2296) . Go to 4 when completed.
2. For B-Sites, perform Step 5 to Step 20 of 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) . Specify RSTP as the value for the STP Mode parameter in Step 8 4 . Go to 4 when completed.
3. For I-Sites, perform Step 8 to Step 25 of 77.8 “To create an I-VPLS” (p. 2253) . Specify RSTP as the value for the STP Mode parameter in Step 14 4 . Go to 4 when completed.
4. Repeat 1 , 2 , or 3 , as required, for each site that you want to create in the MVPLS.

b. Create an MSTP site for the MVPLS.

1. For regular VPLS sites, perform Step 4 to Step 5 of 77.33 “To configure a VPLS site” (p. 2294) . Specify MSTP as the value for the STP Mode parameter in Step 6 of 77.35 “To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site” (p. 2296) . Go to 4 when completed.
2. For B-Sites, perform Step 5 to Step 20 of 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) . Specify MSTP as the value for the STP Mode parameter in Step 8 4 . Go to 4 when completed.
3. For I-Sites, perform Step 8 to Step 25 of 77.8 “To create an I-VPLS” (p. 2253) . Specify MSTP as the value for the STP Mode parameter in Step 14 4 . Go to 4 when completed.
4. Repeat 1 , 2 , or 3 , as required, for each site that you want to create in the MVPLS.

8

Save the changes and close the forms.

9

If the MVPLS site is to manage traffic on associated VPLS SAPs, create a SAP for the MVPLS with a defined redundant VLAN range.

Note: If an MVPLS site has SAPs that manage traffic on the associated VPLS SAPs, you must define a redundant VLAN range during SAP creation in the MVPLS.

1. For regular VPLS sites, perform Step 4 to Step 29 of 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) . Ensure that you perform 77.76 “To configure a redundant VLAN range on a VPLS L2 access interface” (p. 2353) to specify a redundant VLAN range. Go to 4 when completed.
2. For B-Sites, perform Step 4 to Step 25 of 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) . Ensure that you perform Step 20 to specify a redundant VLAN range. Go to 4 when completed.
3. For I-Sites, perform Step 4 to Step 30 of 77.88 “To create a VPLS I-L2 access interface” (p. 2372) . Ensure that you perform Step 23 to specify a redundant VLAN range. Go to 4 when completed.
4. Save the changes and close the form.The MVPLS service (Edit) form reappears.

10

Save the changes and close the forms.

END OF STEPS

77.8 To create an I-VPLS

77.8.1 Steps

1

Choose Create→Service→VPLS from the NFM-P main menu. The VPLS Service (Create) form opens.

2

Select a customer to associate with the VPLS in the Customer panel.

3

Configure the required parameters.

4

To configure service tunnel required bandwidth for the service:

Note: To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

1. Click on the Bandwidth tab.
2. Configure the Bandwidth Method parameter.
3. If the Bandwidth Method parameter is set to Manual or Calculate, configure LSP Path Booking parameter.

4. If the Bandwidth Method parameter is set to Manual, enter a required bandwidth value for each CoS (see Override Service Configuration).

5

Perform one of the following.

a. Create an I-site for the VPLS. Go to Step 6 .

b. Complete service creation if I-sites, I-L2 access interfaces, and SDP bindings for the VPLS are to be created later. Go to Step 30 .

6

On the service tree, right-click on Sites and choose Create VPLS I-Site.

7

Select a site. The I-Site (Create) form opens.

8

Configure the general parameters.

9

Configure the parameters in the Routed VPLS panel.

10

Configure the parameters in the Load Balancing panel.

11

Click on the Backbone tab.

12

Select the Backbone VPLS site.

Note: The selection of the B-site in this step must be repeated for each I-site you create, since you must select the B-VPLS site that is within the same site as the I-site.

13

Configure the required parameters.

The Force Q Tag Forwarding parameter is only displayed if the NE for this site is in chassis mode D or higher or Sparrow.

14

To configure MFib, STP, FIB, and MAC Protection parameters for the I-site:

1. Click on the Forwarding Control tab.

2. Configure the required parameters.
3. Click on the STP tab to configure STP parameters for the I-site.
4. Configure the bridge-level STP parameters for the I-site.

   Nokia STP in a VPLS interoperates with customer STP implementations as a mechanism for loop detection and prevention. The bridge-level parameters balance the STP resiliency and speed of convergence.

   Modifying the bridge-level parameters must be done within the constraints of the following formulas:

   - 2 (Bridge Forward Delay – 1.0 s) >= Bridge Max Age
   - Bridge Max Age >= 2 (Bridge Hello Time + 1.0 s)

   Note:

   If you are configuring an MVPLS I-site, set the STP Mode parameter to RSTP or MSTP, depending on the MVPLS type. The MSTP option is available only if you are creating an MVPLS. See 77.7 “To create an MVPLS” (p. 2251) for more information about creating an MVPLS.

   MSTP is configurable only on the 7450 ESS, 7750 SR, and 7950 XRS.
5. Click on the FIB tab to configure FIB parameters for the I-site.
6. Configure the required parameters.
7. Configure the required MAC move parameters in the Mac Move panel.
8. Click on the MAC Protection tab to configure the list of protected MAC addresses.
9. Click Create. The MAC Protection (Create) form opens.
10. Configure the Protected Mac Address parameter.
11. Save the changes and close the form.
12. To configure an MVPLS I-site that requires MSTP, click on the MSTP tab.
13. Configure the required parameters.
14. Click on the MST Instances tab.
15. Click Create. The MST Instance (Create) form opens.
16. Configure the required parameters.
17. Click on the VLAN Ranges tab. Click Create. The MST Instance Managed VLAN range (Create) form opens.
18. Configure the required parameters.
19. Save the changes and close the forms.

15

To configure ingress multicast forwarding:

Note: An Operational Channels tab appears when you access the VPLS I-Site form in the Edit mode. It displays data for the operational channels when traffic from a specific multicast source for a specific multicast group passes through the service. You must click Search to refresh the data. See Chapter 49, “Policies overview” for information about listing the operational channel parameters.

1. Click on the Multicast tab, then on the Mcast Path Mgmt tab.
2. Select an ingress multicast info policy.

16

To assign test generation options to the site:

1. Click on the OAM tab, then the Configuration tab.
2. Configure the required Test Generation Options parameters.

17

To configure service tunnel required bandwidth for the site:

Note: To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

1. Click on the Bandwidth tab.
2. Enable the Override Service Configuration parameter.
3. Configure the Bandwidth Method parameter.
4. If the Bandwidth Method parameter is set to Manual, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

18

Configure IGMP snooping for the site.

1. Click on the Multicast tab.
2. Click on the IGMP Snooping tab.
3. Configure the required parameters.

   The Query source address parameter is configurable when the Use query source address parameter is enabled.

19

Configure MLD snooping for the site.

1. Click on the Multicast tab.
2. Click on the MLD Snooping tab.
3. Click on the General tab and configure the required parameters.
4. Click on the MRouters tab to view a list of multicast routers.

20

If you are configuring an MVPLS I-site, go to Step 22 .

21

To configure an SHG on the site:

Note: You must configure an SHG or RSHG if you plan to create a spoke circuit from this VPLS I-site to a VLL or to another VPLS.

1. Click on the Split Horizon Groups tab.
2. Click Create. The Site, New Split Horizon Group (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

22

To create an I-L2 access interface for the I-site, perform Step 3 to Step 30 of 77.88 “To create a VPLS I-L2 access interface” (p. 2372) .

23

To create a virtual MEP on the site, perform Step 3 to Step 12 of 77.52 “To create a Virtual MEP on a VPLS site” (p. 2312) .

24

To create a spoke SDP binding between the I-site and regular VPLS sites, click on the Components tab and perform Step 7 to Step 34 of 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) , as required. I-sites can only use spoke SDP bindings.

Note: You cannot create a spoke SDP binding on an MVPLS site that runs MSTP, or enable MSTP on a site that has a spoke SDP binding.
You cannot enable MSTP on an access interface that has a non-zero encapsulation value.

25

Save the changes and close the form.The I-Site (Create) form closes, and a dialog box appears.

26

Save the changes and close the form.The VPLS Service (Create) form displays the new I-site on the navigation tree under VPLS.

27

Repeat Step 6 to Step 26 to create additional I-sites for the I-VPLS.

28

To add protected MAC addresses at the service level:

1. Click on the Forwarding Control tab.
2. Click on the MAC Protection tab.

3. Click Create. The MAC Protection (Create) form opens.
4. Configure the Protected Mac Address parameter.
5. Save the changes and close the form. Protected MAC addresses that you add on the site level, as performed in Step 14 , are automatically added to the service-level MAC protection list.

29

To configure bandwidth for the service if required, click on the Bandwidth tab:

Note: The Bandwidth tab is only available if service CAC is configured. See the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide* for information about enabling and disabling service CAC.

The ability to configure required bandwidth is only applicable to I-sites.

1. For each CoS, configure the CoS Bandwidth (Mbps) parameter.
2. Click on the General tab to determine the CAC status.

   If the Verify CAC button is enabled, the CAC has not been verified. The Probable Cause and CAC Status fields provide details.

The CAC Status field describes the current CAC status of the service. CAC statuses are as follows.

- CAC Verified indicates that all tunnels have sufficient bandwidth to admit service and that requested bandwidth for the service is booked on the appropriate physical links.
- CAC Failed indicates that the attempt made to admit service into the network was unsuccessful. The most likely cause for this is insufficient bandwidth. See the Probable Cause field for more specific information.
- BW Defined, No CAC Request indicates that the required bandwidth is defined on the service; however, a CAC request has not occurred.
- CAC To be Verified indicates that a tunnel has been configured on the service either manually or through the CLI; however, the required bandwidth has not been verified in the network.

The Probable Cause field describes possible reasons for the current CAC state. Probable causes are as follows.

- No Candidate Tunnels Found indicates that the autobind tunnel function was found, but no suitable tunnels were found.
- Different PBB Tunnels Applied to Service indicates that two or more different PBB tunnels are configured on this service.
- Not Enough Bandwidth on any Candidate Tunnels indicates that one or more candidate tunnels were found, but the available bandwidth was insufficient to admit the service.
- Automatic PBB Tunnel Selection Failed indicates that a suitable PBB tunnel was found, but there were errors when attempting to assign the tunnel to the service. A dialog box will provide more details.
- Site Missing Tunnel indicates that at least one selected site is not configured with a PBB tunnel.

- All PBB Tunnels have not been Verified indicates that all sites within the service have a tunnel configured, but the available bandwidth has not been booked or verified in the network.

3. To manually verify the CAC, click on the Verify CAC button if it is enabled. The NFM-P will attempt to find the most appropriate PBB tunnel for the service based on available bandwidth, and to automatically bind the tunnel to the service if one has not already been assigned.

30

Save the changes and close the forms.

END OF STEPS

77.9 To modify a VPLS

CAUTION

Service Disruption

Modifying parameters can be service-affecting.

77.9.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

The following tabs contain parameter information for the service:

- General tab—displays the general service properties
- MVR tab—displays multicast package policy information for the service
- Bandwidth tab — displays service tunnel bandwidth booking information for the service.
- OAM tab—allows the creation and execution of service-specific diagnostic tests
- Forwarding Control tab—lists the FIB and STP instances
- Sites tab—lists the sites that belong to the service
- L2 Access Interfaces tab—lists the L2 access interfaces that belong to the service
- L2 Management Interfaces tab—lists the L2 management interfaces that belong to the service
- Mesh SDP Bindings tab—displays the mesh SDP bindings that belong to the service
- Spoke SDP Bindings tab—displays the spoke SDP bindings that belong to the service, including active and backup SDP bindings associated with endpoints.

- Endpoints tab—displays the endpoints associated with the service
- Faults tab—displays the faults associated with the service
- Schedulers tab—configures the rate type and displays the egress aggregate rate limit
- QoS tab

3

Navigate to related components for the service using the navigation tree, as required.

Certain service components have related objects that can be easily examined using the navigation tree. For such objects, right-clicking them in the navigation tree offers a “Navigate to...” option in the contextual menu. If you select this option, the properties form for the related object is displayed. This can be very convenient, especially for complex services containing many sites and components.

The VPLS components for which you can navigate to such related objects include:

- I-Sites: you can navigate to the associated B-VPLS or B-(M)VPLS
- L2 Access Interfaces: you can navigate to the opposite SAP in a VLAN Uplink configuration
- Redundant L2 Access Interfaces: you can navigate to the opposite SAP.
- Mesh SDP Bindings: you can navigate to the opposite (return) mesh SDP
- Spoke SDP Bindings: you can navigate to the opposite (return) spoke SDP

4

Modify the components and parameters for the service, as required.

5

Save the changes and close the forms.

END OF STEPS

77.10 To view VPLS contents

77.10.1 Purpose

Use this procedure to view various contents of a VPLS and any modifications you make to it before deploying the changes.

Note: The procedure also provides information on the following policy types associated with the service:

- QoS SAP ingress and QoS SAP egress polices
- ACL IP/IPv6 filters and ACL MAC filter policies

Two information views are available. The Committed info view displays the current contents of a service, and the Committed menu item is always enabled. The Modified info view allows you to review any changes you make to a service before committing them. Modified, created, and deleted attributes and objects are displayed.

77.10.2 Steps

1

Choose Manage Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search and select the service that you want to view or modify.

3

On the service navigation tree, expand the Sites icon and click on the required site. The site properties form is displayed.

4

Perform one of the following:

a. To view the currently committed service contents, go to Step 5 .

b. To view modified service contents before committing any changes, go to Step 6 .

5

Click on the Show Info and choose the Committed menu item. A Committed Values form is displayed and shows various policy attributes that have been previously configured and saved in the local policy. The information displayed in the form is similar to the information retrieved in CLI by running the “config>service>*Service Type#* info” command, where *Service Type#* is the type of service (followed by its Service ID, for example: vpls42) that you want to query.

6

Modify any service parameters or objects, as required. Otherwise go to Step 9 .

7

Click on the Show Info and choose the Modified menu item. A Modified Values form is displayed. The table lists modified, created, and deleted actions, as well as specific attributes and objects, along with their old value, new value, and tab location. The Attribute Title corresponds to the attribute or object name acted upon by your current modifications. For created objects, the values of mandatory attributes are shown in comma-separated format.

8

Select an item in the Modified Values form and then click Show on Form. The service form tab containing the changed item is displayed and the modified attribute is highlighted in blue.

9

Save your changes if required, and close the form.

END OF STEPS

77.11 To modify a VLPS using the topology view

77.11.1 Purpose

The topology view for a service provides a graphical representation of the various components and their interconnections. You can also use this view to add, modify, or just navigate to service components. This provides an alternative approach to performing these functions from the navigation tree view.

Working from the topology view can expedite the creation of the components, since many of the fields you would ordinarily have to set in the configuration forms will be automatically populated using this approach. OAM components can also be configured in this way. The configuration forms can also be accessed directly at any time from this view by right-clicking a component. This allows quick access to conduct more detailed component configuration.

77.11.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click on the Topology View button. The Service Topology map opens.

The remainder of this procedure contains a number of sub-procedures describing the various components that can be created and modified from the topology view. These include:

- Creating a new site. Go to Step 3 .
- Creating site components. Go to Step 9 .
- Creating SDP bindings. Go to Step 22 .

Adding a new site

3

Right-click on any blank space in the service topology map. A contextual menu is displayed. You can choose to create one of the following:

- VPLS Site
- VPLS B-Site
- VPLS I-Site

4

Select the required option. The Select Network Elements form appears.

5

Select one or more sites to add to the service and click OK. The VPLS Site (Create), VPLS I-Site (Create), or VPLS B-Site (Create) form is displayed, depending on your menu item selection. If you selected more than one site, the VPLS Site (Multiple Instances) (Create), VPLS I-Site (Multiple Instances) (Create), or VPLS B-Site (Multiple Instances) (Create) form is displayed, depending on your menu item selection.

6

Save the changes and close the form.The VPLS Site (Create) (or VPLS Site (Multiple Instances) (Create)) form closes and the new site (or sites) is displayed on the map.

7

If you want to perform detailed configuration of site properties for the new site, right-click on the site icon and select Properties. The VPLS Site (Edit) form opens. See 77.5 “To create a VPLS” (p. 2249) for detailed site configuration information.

8

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

Creating site components

9

Right-click on any site icon in the service topology map. A contextual menu is displayed. You can choose to create one of the following:

- VPLS L2 Access Interface (this choice will actually display as either a regular L2 Access Interface, or a B-L2 or I-L2 variant, depending on the site you select). Go to Step 10 .
- VPLS L2 Management Interface. Go to Step 14 .
- VPLS Endpoint. Go to Step 18 .

10

If you choose to create a VPLS L2 Access Interface, then the VPLS L2 Access Interface (Create) form is displayed. If the selected site is a B-Site or I-Site, then the B-L2 or I-L2 Access Interface (Create) form is displayed accordingly.

11

Click on the Port tab and assign a port to the interface.

See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) (or 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) or 77.88 “To create a VPLS I-L2 access interface” (p. 2372) for B-L2 or I-L2 Access Interfaces respectively) for detailed information on further configuring the interface.

12

Save the changes and close the form.The VPLS L2 Access Interface (Create) form closes.

13

Go to Step 21 .

14

If you choose to create a VPLS L2 Management Interface, then the VPLS L2 Management Interface (Create) form is displayed.

15

Click on the Port tab and assign a port to the interface.

See 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) for detailed information on further configuring the interface.

16

Save the changes and close the form.The VPLS L2 Management Interface (Create) form closes.

17

Go to Step 21 .

18

If you choose to create a VPLS Endpoint, then the Endpoint (Create) form is displayed.

19

Configure the Name parameter for the endpoint.

See 77.5 “To create a VPLS” (p. 2249) (or 77.24 “To create a B-site for VPLS or MVPLS” (p. 2283) for B-Sites) for detailed information on further configuring the endpoint.

20

Save the changes and close the form.The Endpoint (Create) form closes and the new endpoint is displayed in the topology view.

21

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

Creating SDP bindings

22

Select the sites you want to connect in the service topology map and right-click on any one of them. A contextual menu is displayed.

Note: If you intend to create either a spoke or a mesh binding between two sites, the order in which you select them is important. The first site you select will become the source site and the second site will become the destination site. Therefore, it is not recommended that you do a marquee-select in the topology view, since you will not be sure of this hierarchy. Instead, select the sites individually, and hold down the Shift or Ctrl key after your first selection.
If you intend to create a full mesh between sites, then the order in which you select the sites is not important.
If you intend to add a service site to an existing mesh of sites, then the first site you select must be the one that is not currently a part of the mesh. The second site you select can be any member of the existing mesh.

23

Select Connect. Depending on the sites you selected, one or more of the following options are available:

- Create Spoke SDP Binding. Go to Step 24 .
- Create Mesh SDP Binding. Go to Step 31 .
- Create Full Mesh. Go to Step 38 .
- Add To Existing Mesh. Go to Step 44 .

24

If you choose the Create Spoke SDP Binding option, then the Spoke SDP Binding (Create) form is displayed.

Note: For this function, it is assumed that you clicked on the source site first and then held down the Shift or Ctrl key while right-clicking on the destination site to display the contextual menu.
You can also create a spoke SDP binding between a site icon and an endpoint icon, or between two endpoint icons in the topology view. Appropriate endpoints must first exist or be created to enable this.

25

Enable the Auto-Select Transport Tunnel parameter.

26

You can manually configure other parameters here if required, or just click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new binding between the two sites is displayed in the topology view. See 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) for more detailed information on creating and configuring spoke SDP bindings.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

27

Assuming that the spoke SDP binding was successfully created in Step 26 , select the same two sites again in the topology view, although this time do so in the opposite order that you originally selected them. This will allow you to create a spoke binding for the return tunnel.

28

Right-click on the second site you selected and choose the Create Spoke SDP Binding ... option. The Spoke SDP Binding (Create) form is displayed.

29

You can manually configure other parameters here if required, or just click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new return binding between the two sites is displayed in the topology view.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

30

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

31

If you choose the Create Mesh SDP Binding option, then the Mesh SDP Binding (Create) form is displayed.

Note: For this function, it is assumed that you clicked on the source site first and then held down the Shift or Ctrl key while right-clicking on the destination site to display the contextual menu.

32

Enable the Auto-Select Transport Tunnel parameter.

33

You can manually configure other parameters here if required, or just click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Mesh SDP Binding (Create) form closes and the new binding between the two sites is displayed in the topology view. See 77.91 “To create a VPLS or MVPLS mesh SDP binding” (p. 2382) for more detailed information on creating and configuring mesh bindings.
- If an available transport tunnel does not exist between the two sites, then an error message

is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

34

Assuming that the mesh SDP binding was successfully created in Step 33 , select the same two sites again in the topology view, although this time do so in the opposite order that you originally selected them. This will allow you to create a mesh binding for the return tunnel.

35

Right-click on the second site you selected and choose the Create Mesh SDP Binding ... option. The Mesh SDP Binding (Create) form is displayed.

36

You can manually configure other parameters here if required, or just click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Mesh SDP Binding (Create) form closes and the new return binding between the two sites is displayed in the topology view.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

37

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

38

If you choose the Create Full Mesh option, the Complete Vpls Mesh for Service Component form is displayed.

39

Click Select to display the Tunnel Selection Profiles form.

40

Click Search to display a list of available Tunnel Selection Profiles.

41

Click on the desired profile entry and click OK. The Tunnel Selection Profiles form closes and your selection is displayed in the Complete Vpls Mesh for Service Component form.

42

Click OK to accept your selection and close the Complete Vpls Mesh for Service Component form. One of the following will result:

- If available transport tunnels exist between the selected sites, then the new mesh bindings between the sites are created and displayed in the topology view. See 77.91 “To create a VPLS or MVPLS mesh SDP binding” (p. 2382) for more detailed information on creating and configuring mesh bindings.
- If available transport tunnels do not exist between all the sites, then the new bindings that can be created are displayed, and an error message is also displayed informing you that the full mesh could not be completed. See Chapter 33, “Service tunnels” for information on how to create the required tunnels. Once the tunnels are created, you can repeat this sub-procedure.

43

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

44

If you choose the Add To Existing Mesh option, the Complete Vpls Mesh for Service Component form is displayed.

45

Click Select to display the Tunnel Selection Profiles form.

46

Click Search to display a list of available Tunnel Selection Profiles.

47

Click on the desired profile entry and click OK. The Tunnel Selection Profiles form closes and your selection is displayed in the Complete Vpls Mesh for Service Component form.

48

Click OK to accept your selection and close the Complete Vpls Mesh for Service Component form. One of the following will result:

- If available transport tunnels exist between the selected sites, then the new mesh bindings between the sites are created and displayed in the topology view. See 77.91 “To create a VPLS or MVPLS mesh SDP binding” (p. 2382) for more detailed information on creating and configuring mesh bindings.
- If available transport tunnels do not exist between the selected sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnels. Once the tunnels are created, you can repeat this sub-procedure.

49

Return to Step 2 for a list of other functions you can perform from the topology view or go to Step 50 to finish.

50

Close the forms.

END OF STEPS

77.12 To view the service topology associated with a VPLS

77.12.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Topology View. The Service Topology map opens.

See Chapter 4, “Topology map management” more information about service topology views.

Note: In the VPLS topology map view, a redundant spoke SDP binding under an endpoint displays differing colors, depending on whether it is in the active or backup state. Backup spoke SDP bindings are shown in purple.

3

Close the Service Topology form.

4

Close the Manage Services form.

END OF STEPS

77.13 To delete a VPLS

CAUTION

Service disruption

Deleting a service may result in a service disruption for customers.

Consider the implication of deleting the service before proceeding.

Note: A VPLS cannot be deleted if the L2 access interface has MSAPs that are in an active state. See Chapter 74, “Residential subscriber management” for more information about MSAPs.

77.13.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service or a range of services.

3

Click on Delete. A warning form appears. This form is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 4 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 4 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 4 .

4

For all services, regardless of the priority configuration, acknowledge the check box that prompts you to confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

5

Click Yes to confirm the action. The service is deleted and removed from the list.

6

Close the Manage Services form.

END OF STEPS

77.14 To copy or move a VPLS

77.14.1 Purpose

Perform this procedure to copy or move a VPLS from one NE to another NE. The VPLS service site and all of the children objects, including SAPs, SDP bindings, MEPs on the SAPs and related CFM tests, are moved to the destination site. MEPs and associated CFM tests are included in VPLS

copy/move operations for B-sites, L2 SAPs and SDP bindings (including return SDP bindings). Return SDP bindings (from remote site to the source site) for SDP bindings from the source site to a remote site are copied or moved. Ingress and egress, scheduler, policer, accounting, and IP/MAC filters (ACLs) policies are copied and distributed to the destination NE.

If there are SAPs on different ports within a VPLS, you must move the service port by port. When the service does not already exist on the destination NE—that is, before an move operation—the service site and children objects are moved to the destination NE. Only SAPs on the specified port are moved to the destination site. After an initial move, moving the remaining SAPs on other ports is a SAP move operation. See 16.32 “To copy or move L2 SAPs between ports” (p. 614) .

Note: Consider the following site limitations when you copy or move a service:

- Copying and moving VPLS services is only supported on 7750 SR, 7450 ESS, and 7950 XRS NEs.
- The NE type and major software release of the source and destination NEs must be the same.
- The chassis mode of the destination NE must be the same or higher than the chassis mode of the source NE.
- The NFM-P copies the source site and children objects and creates them on the destination NE.
- You cannot move sites that are used in a backbone configuration.
- You cannot move sites for which RVPLS is enabled.
- You cannot move sites created by a RADIUS script for dynamic service.

Consider the following SAP limitations when you copy or move a service:

- The source and destination ports must be compatible, for example, have the same encapsulation type, be HSMDA ports or not, or both be access or hybrid.
- You cannot move or copy SAPs that are configured on PW ports and Ethernet tunnel endpoints.
- SAPs with connection profiles (ATM/VLAN) can be copied or moved when there is at least one SAP configured on the source port with the encap value in the specified range on the source service site.
- If the source site only contains SAPs with connection profiles, the site will not be copied or moved.
- If the source site contains SAPs associated with connection profiles on the specified port, and SAPs on a different port, or SAPs on the specified source port but with encap values out of the range, the site will not be copied or moved.
- Connection profiles on source SAPs must not overlap with the encap values of SAPs on the destination port or fall in a range where a SAP on the destination port may exist.

  The encap range specified in the service copy or move procedure does not apply to connection profiles.
- If there is a Ethernet Ring Path Endpoint associated with the source SAP, the association is not be copied to destination SAP.

Take into account the following service tunnel considerations when you move a service:

- If the spoke SDP bindings on the source NE use an SR-ISIS or SR-OSPF service tunnel,

the NFM-P looks for a matching tunnel on the destination NE. If no matching tunnel exist, then the NFM-P creates one.

- If the spoke SDP bindings on the source NE use an MPLS-TP service tunnel, the copy/move operation is allowed, but an MPLS-TP service tunnel must exist on the destination NE.
- You cannot move SDP bindings that use L2TPv3 service tunnels.
- You cannot move spoke SDP bindings if they are not manually created.

77.14.2 Steps

1

Choose Tools→Copy/Move→Service from the NFM-P main menu. The Service Copy/Move form opens.

2

Configure the Action Type parameter.

3

Set the Service Type parameter to VPLS or VPLS+VLL.

Note: Only VPLS is supported for the action type Copy.

4

In the Source panel, choose the NE from which you want to copy or move the service.

5

In the Source Port panel, choose the source port from which you want to copy or move the service.

6

Configure required parameters in the Source panel.

7

In the Destination panel, choose the NE to which you want to copy or move the service.

8

In the Destination Port panel, choose the port to which you want to copy or move the service.

9

Configure the required parameters in the Execution Details panel.

The NFM-P creates SAPs on the destination port with the same encapsulation values as the source SAP if you set the Outer Encap Value Offset and Inner Encap Value Offset parameters

to 0. Otherwise, the NFM-P creates SAPs on the destination port encapsulation values equal to the source encapsulation value plus the Outer Encap Value Offset or Inner Encap Value Offset parameter.

10

Click Result Export Path and specify the file name and location in which to save a text file that contains the results of service copy/move operation.

11

Click Execute.

12

Click on the Results tab to view the results of the copy or move action.

END OF STEPS

Note:

For a VPLS move, if the new site is created successfully with all children copied from the source site, the source site is deleted.

A service move searches for all of the SAPs with specified service type on the source NE, within the encapsulation value range that you specified in Step 6 on the specified port. The NFM-P creates the associated service sites on the destination NE with the corresponding SAPs on the destination port, including all children objects, such as SDP bindings, that existed on the source site.

The source site is not deleted if the new site is created successfully with only some of the SAPs from source site, for instance, there are SAPs remaining on the source site on different ports or with outer or inner encapsulation values out of the specified range. The SAPs that are within the specified encapsulation range that you specified in Step 6 will be deleted from the source port after the move, along with all corresponding SDP bindings.

If the corresponding service site exists on the destination NE, only the SAPs within the specified encapsulation range are copied/moved to the destination site. All other children objects, such as SDP bindings, under the source site are not copied or moved to the destination site. In this case, the source site will be deleted along with all other children objects when the last SAP is moved.

If a SAP exists on the destination site with same encapsulation values as the source SAP (the source SAP encapsulation values plus the Outer Encap Value Offset or Inner Encap Value Offset parameters that you specified in Step 9), the move action fails.

If there is a failure when moving a site or any of the child objects, the move action for this site fails and nothing from this site is moved to the destination NE.

If there are multiple service sites being copied or moved along with the SAPs as specified by source port and encapsulation range, the copy or move action continues to the next site if there is an individual failure.

The NFM-P reports that the copy or move action is finished when the NFM-P database transaction has completed; deployment operations on the node may continue after the database transaction is complete.

77.15 To view the VPLS operational status

77.15.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

View the Aggregated Operational State and State Cause indicators. When the Aggregated Operational State is Down or Partially Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The Object Alarms tab is displayed.

6

Click on the Aggregated Alarms tab to view the aggregated alarms for the object.

7

Click on the Tests tab if the OAM Validation Failed indicator is enabled. You can view the OAM validation test suite results by clicking on the Tested Entity Result tab.

Note: You can run the OAM Validation test suite for this service from this form by clicking on the Validate button. If the Validate button is not visible, click on the More Actions button and choose Validate.
As an alternative, you can also run an OAM validation test on the service by performing a One Time Validation.

8

Save the changes and close the forms.

END OF STEPS

77.16 To configure a VPLS for AA reporting

77.16.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Application Assurance tab. The NSP Analytics Parameters tab is displayed.

4

Configure the required parameters.

5

To configure one or more custom DCP groups, see the custom DCP group configuration procedure in Chapter 87, “Application assurance” .

6

To specify one or more application or application groups for IP detail reporting:

1. Click on the IP Detail Application or IP Detail Application Group tab and click Create. The IP Detail Application or IP Detail Application Group (Create) form opens.
2. Configure the required parameters.

   You can specify up to 10 applications and up to 10 application groups.

7

Save the changes and close the forms.

END OF STEPS

77.17 To configure an Ethernet segment

77.17.1 Purpose

Ethernet segments are used by EVPNs. An Ethernet segment is a group of ports on an NE that are part of the same redundancy group, and are identified by a unique Ethernet Segment Identifier.

77.17.2 Steps

1

Navigate to the NE where you want to configure an Ethernet segment.

2

Right-click on the NE and choose Properties. The Network Element form opens.

3

Click on the Globals tab, and then on the BGP tab. In the Ethernet Segment panel, click Create. The BGP EVPN Ethernet Segment form opens.

4

Configure the required parameters.

5

Save your changes and close the form.

END OF STEPS

77.18 To configure a BGP EVPN

77.18.1 Steps

A 7750 SR, 7450 ESS, 7250 IXR, or 7950 XRS can generate Ethernet auto-discovery routes. When BGP EVPN is configured, a single AD per-ES route with the associated RD and a set of EVI route-targets is advertised.

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, and then on the Service tab.

3

Click on the BGP tab, and then on the EVPN tab.

4

Configure the required BGP EVPN parameters. The following considerations apply:

- The default Route Distinguisher Type is None. In this case, the RD is automatically built using the BGP-AD VPLS ID configured under the VPLS Site BGP tab. If you set the Route Distinguisher Type to Configured, then you need to provide an RD IP address.

- The default AD per-ES routes advertisement type is EVI-RT, which specifies the option to advertise a separate AD per-ES route per service.

  The EVI-RT Set advertisement type specifies the option to advertise a set of AD per-ES routes aggregating the route targets for all the services in the Ethernet segment. If you choose EVI-RT Set, then you need to specify the IP address part of the RD that is being used in the EVI-RT Set option.
- To allow for both EVPN VXLAN and EVPN MPLS to be in an Admin Up state on the same VPLS, different BGP IDs should be assigned to VXLAN and MPLS.
- If two EVPN VXLAN instances are created on the same VPLS, different BGP IDs should be assigned to each VXLAN.
- To allow propagation of attributes to other owners in VPRN, enable the Attribute Uniform Propagation parameter.
- To allow new BGP based selection for RT5 routes, enable BGP Path Selection parameter.
- To ignore the D-PATH domain segment length, enable the D-Path Length Ignore parameter.

5

Save the changes and close the form.

END OF STEPS

77.19 To assign a multicast package policy to a VPLS

77.19.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Multicast tab.

4

Click on the MVR tab.

5

Select a multicast package policy. The Select Multicast Package Policy - VPLS Service - Subscriber form opens.

6

Choose a multicast package policy and click OK. The Select Multicast Package Policy - VPLS Service - Subscriber form closes, and the VPLS Service (Create) form refreshes with the multicast package policy name.

After the multicast package policy is applied to the VPLS, the policy is distributed as the routing policy to all MVR-capable VPLS sites. If you apply another package policy to the site, the new policy is distributed to the site. The previously distributed policy remains on the site.

7

Save the changes and close the forms.

8

Close the Manage Services form.

END OF STEPS

77.20 To configure bandwidth management for a VPLS

Note: In order to configure required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

77.20.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Bandwidth button.

4

Click on the Required Bandwidth tab to configure service tunnel required bandwidth for the service.

5

Configure the Bandwidth Method parameter.

6

If you set the Bandwidth Method parameter to Input Manually or Calculate From SAPs,

configure LSP Path Booking parameter.

7

If you the Bandwidth Method parameter to Input Manually, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

8

Click on the Bandwidth Reserved Tunnel tab to reserve PBB tunnel bandwidths.

9

For each CoS, enter a value for the CoS Reserved Bandwidth (Mbps) to specify how much bandwidth this tunnel will reserve in the network.

10

Click on the Reserve Bandwidth button. This action checks to ensure that all the active links in the B-VPLS have enough bandwidth to admit the tunnel into the network and book the bandwidth. The tunnel status will be updated appropriately, based on the outcome of the action. Once the bandwidth is reserved, the BW Utilization tab on the applicable Physical Link properties form will also show this tunnel and the bandwidth information.

Note: The reserved bandwidth of the tunnel can be changed at any time after the service creation. However, if a change to the reserved bandwidth causes the used bandwidth to be greater than the requested change, or if there is insufficient bandwidth in the network to facilitate this change, then it will be denied and an appropriate message is displayed.
Once bandwidth is reserved on the tunnel, any changes to the topology of the B-VPLS (for example, uplinks added, STP state changes, and so forth) will be updated automatically with the correct bandwidth information on the underlying physical links. If there is insufficient bandwidth available when the changes happen, the bandwidth will still be booked on the physical links and the appropriate alarms will be raised.
To unreserve bandwidth of the tunnel after the service creation, you can set all of the configured Cos Reserved Bandwidth parameters back to 0. However, this can only be done when there are no longer any i-services riding on this PBB Tunnel.
After the service has been created, the Tunnel Usage tab shows all the i-services currently using this tunnel and the specific bandwidth usage per service.

11

Save the changes and close the forms.

END OF STEPS

77.21 To add protected MAC addresses to a VPLS

77.21.1 Purpose

Protected MAC addresses that you add at the VPLS site level are automatically added at the service level. See Step 8 of 77.33 “To configure a VPLS site” (p. 2294) for more information about adding protected MAC addresses at the site level.

77.21.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Forwarding Control tab, then the MAC Protection tab.

4

Click Create. The MAC Protection (Create) form opens.

5

Configure the Protected Mac Address parameter.

6

Save the changes and close the forms.

END OF STEPS

77.22 To connect a G.8032 Ethernet ring to a VPLS

77.22.1 Purpose

Perform the following procedure to connect an existing G.8032 Ethernet ring to a VPLS. See 33.5 “Ethernet G.8032 rings” (p. 1183) for information about using the NFM-P to automatically create G.8032 Ethernet rings. See 33.18 “To create an Ethernet G.8032 ring” (p. 1212) to create G.8032 Ethernet rings using the NFM-P GUI.

Connecting an Ethernet ring adds the sites in the ring to the VPLS, then creates SAPs and paths for the ring.

77.22.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Perform the following bandwidth management configuration as required:

a. To configure service-level bandwidth management parameters for the service, see 77.20 “To configure bandwidth management for a VPLS” (p. 2278) .

b. To configure service tunnel required bandwidth for a VPLS site, see 77.40 “To configure service tunnel required bandwidth for a VPLS site” (p. 2301) .

4

Select one or more sites in the navigation tree and right-click and choose Connect to Ethernet Ring. The sites must be part of the same Ethernet ring. The Connect to Ethernet Ring form opens.

5

Select an Ethernet ring:

a. Click Select to manually select an Ethernet ring.

b. Enable the Automatic Ring Selection parameter to automatically select an Ethernet ring and select a Tunnel Selection Profile.

Note: To use a Tunnel Selection Profile, the Multi-Segment Tunnel Selection system preference must be enabled. For information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

6

Configure the required parameters.

Note: If you are connecting a G.8032 Ethernet ring to a VPLS using a Dot1q SAP, only configure the Data Service Outer Encap parameter. For QinQ SAPs, configure both the Data Service Outer/Inner Encap parameters.

7

As required, repeat Step 3 to perform additional bandwidth management configuration on customer SAPs. If you modify the bandwidth parameters, choose Audit CAC to recalculate service tunnel bandwidth booking.

8

Save the changes and close the forms. All sites in the ring are created in the service, if they were not already present.

END OF STEPS

77.23 To configure custom object attributes for AA reporting

77.23.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

On the General tab, choose an Application Profile for the applicable Group and Partition.

6

Click on the Application Assurance tab, then the NSP Analytics Parameters sub-tab.

7

Click on the Reporting tab and click Create. The AA Reporting (Create) form opens.

8

Configure the required parameters.

9

Save the changes and close the forms.

END OF STEPS

77.24 To create a B-site for VPLS or MVPLS

77.24.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on the Sites icon and choose Create VPLS B-Site.

4

Select a site. The B-Site (Create) form opens.

5

Configure the required parameters.

6

Click on the Backbone tab.

7

Configure the backbone parameters.

The Source MAC Address should not be duplicated for other B-sites within the same B-VPLS.

The Source MAC Address parameter is the only applicable parameter when configuring PBB on a 7210 SAS-M, 7210 SAS-T, or 7210 SAS-X.

The Use SAP Backbone MAC Address parameter is configurable only on the 7750 SR, 7750 SR-c12, 7750 SR-12E, 7450 ESS, or 7950 XRS. The NE must also be configured with an IOM 3 MDA or an XMA.

8

To configure MFIB, STP, FIB, and MRP parameters for the B-site:

1. Click on the Forwarding Control tab.
2. Configure the required parameters.
3. Click on the STP tab to configure STP parameters for the B-site. Otherwise, go to 5 .
4. Configure the bridge-level STP parameters for the B-site.

Nokia STP in a VPLS interoperates with customer STP implementations as a mechanism for loop detection and prevention. The bridge-level parameters balance the STP resiliency and speed of convergence.

Modifying the bridge-level parameters must be done within the constraints of the following formulas:

- 2 (Bridge Forward Delay – 1.0 s) >= Bridge Max Age
- Bridge Max Age >= 2 (Bridge Hello Time + 1.0 s)

Note:

If you are configuring an MVPLS B-site, set the STP Mode parameter to RSTP or MSTP, depending on the MVPLS type. The MSTP option is available only if you are creating an MVPLS. See 77.7 “To create an MVPLS” (p. 2251) for more information about creating an MVPLS.

MSTP is configurable only on the 7450 ESS, 7750 SR, and 7950 XRS.

5. To configure FIB parameters for the B-site, click on the FIB tab.
6. Configure the required parameters.
7. Configure the MAC move parameters in the Mac Move panel.
8. If you are configuring an MVPLS B-site that requires MSTP, click on the MSTP tab.
9. Configure the required parameters.
10. Click on the MST Instances tab and click Create. The MST Instance (Create) form opens.
11. Configure the required parameters.
12. Click on the VLAN Ranges tab. Click Create. The MST Instance Managed VLAN range (Create) form opens.
13. Configure the required parameters.
14. Save the changes and close the forms.
15. Click on the MRP tab and configure the required parameters.

    Note:

    You can view information regarding MMRP Entries for the access interface and/or SDP Binding by clicking on the MMRP Entries tab.

9

To assign test generation options to the site:

1. Click on the OAM tab, then the Configuration tab.
2. Configure the required Test Generation Options parameters.

10

To configure service tunnel required bandwidth for the site:

Note: To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management checkboxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

1. Click on the Bandwidth tab.
2. Enable the Override Service Configuration parameter.
3. Configure the Bandwidth Method parameter.
4. If the Bandwidth Method parameter is set to Manual, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

11

To configure a provider tunnel for the site:

When you configure an inclusive Provider Tunnel, multicast and broadcast Ethernet frames are always forwarded over a P2MP LSP, defined as the Inclusive Provider Multicast Service Interface (I-PMSI). When a Provider Tunnel is not configured, these frames are replicated at the ingress PE and a copy of each frame is forwarded over the P2P PW to each destination PE participating in the VPLS instance. Provider Tunnels can be configured on VPLS and B-VPLS sites.

Note: Provider tunnels are supported on the 7750 SR-7, 7750 SR-12, 7750 SR-c4, 7750 SR-c12, 7750 SR-12E, and 7450 ESS NEs in all chassis modes, as well as on the 7540 ESS, provided all network IP interfaces are on IOM3/IMM ports (chassis mode D).

1. Click on the Provider Tunnel tab.
2. Enable the Enable Provider Tunnel parameter.
3. Configure the required parameters.

   The Administrative State parameter is only configurable when the Enable BGP AD parameter has been enabled for the site.

   The LSP Template parameter is only configurable when the Type parameter is set to RSVP.

12

To create an endpoint for redundancy (dual homing) on the B-site:

1. Click on the Endpoints tab.
2. Click Create. The Endpoint (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the forms.

13

If you are configuring an MVPLS B-site, go to Step 17 .

14

To configure an SHG on the site:

Note: You must configure an SHG or RSHG if you plan to create a spoke circuit from this VPLS B-site to a VLL or to another VPLS.

1. Click on the Split Horizon Groups tab.

2. Click Create. The Site, New Split Horizon Group (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

   Note:

   You can view information about the I-Sites and Epipe sites associated with the B-Site by clicking on the Associated Sites tab.

15

To create a B-L2 access interface for the site, perform Step 4 to Step 30 of 77.87 “To create a VPLS or MVPLS B-L2 access interface” (p. 2366) .

16

To create a mesh SDP binding for the site, perform Step 4 to Step 21 of 77.91 “To create a VPLS or MVPLS mesh SDP binding” (p. 2382) .

17

To create a redundant spoke SDP binding under an endpoint, perform Step 3 to Step 34 of 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

18

To create a spoke SDP binding for the site, perform Step 5 to Step 34 of 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

Note: You cannot create a spoke SDP binding on an MVPLS site that runs MSTP, or enable MSTP on a site that has a spoke SDP binding.
You cannot enable MSTP on a SAP that has a non-zero encapsulation value.

19

To create a virtual MEP on the site, perform Step 3 to Step 12 of 77.52 “To create a Virtual MEP on a VPLS site” (p. 2312) .

20

Save the changes and close the forms.

END OF STEPS

77.25 To view SPB fate-shared objects

77.25.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on the B-site and choose Properties. The VPLS Site (Edit) form opens.

4

Click on the SBP Fate Sharing tab. The SAPs tab is displayed with a list of fate-shared SAPs within the SPB B-VPLS.

5

Click on the Spoke SDPs tab. A list of fate-shared spoke SDP bindings within the SPB B-VPLS is displayed.

6

Click on the SPB User Sites tab. A list of SPB user B-VPLS sites associated with the control B-VPLS is displayed.

7

Close the forms.

END OF STEPS

77.26 To list the SPB instances on an NE

77.26.1 Steps

1

Right-click on an NE in the Equipment navigation tree and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

Click on the SPB Control Instances tab. A list of all SPB instances configured on the NE is displayed.

4

Close the Network Element (Edit) form.

END OF STEPS

77.27 To create a static ISID range on a VPLS B-L2 access interface or spoke SDP binding

77.27.1 Purpose

Perform the following procedure to create a static ISID range to allow I-VPLS services to connect to non-SPB I-VPLS services. You can create one or more static ISID ranges on a B-L2 access interface or spoke SDP binding on a control or user B-VPLS site. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for information about how to enable SPB on a control or user site.

77.27.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the B-VPLS site.

4

Perform one of the following:

a. To add a static ISID range on a B-L2 access interface, expand the L2 Access Interfaces icon, and click on the B-L2 access interface. The Port (Edit) form opens.

b. To add a static ISID range on an SDP binding, expand the Spoke SDP Bindings icon and click on the SDP binding. The Spoke SDP Binding (Edit) form opens.

5

Click on the Static ISID tab.

6

Click Create. The Static ISID Range Entry (Create) form opens.

7

Configure the required parameters.

8

Save the changes and close the form.

9

To view the status information about the configured static ISIDs, click on the Static ISID Ranges tab.

10

Save the changes and close the forms.

END OF STEPS

77.28 To run a VPLS service OAM validation test

77.28.1 Prerequisites

An OAM validator test suite must be created for the tested entity.

Note: As an alternative, you can also run an OAM validation test on the service by performing a One Time Validation.

OAM validation tests are not supported for HVPLS.

77.28.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click Validate. If an OAM validator test suite is not associated to the service, a dialog box appears:

1. Click OK to associate the service with an existing OAM validator test suite. The Choose Validator Test Suite form opens.
2. Select an OAM validator test suite.

4

View the State Cause indicators. When the validation test fails, a check mark appears beside the OAM Validation Failed indicator.

5

Click on the Tests tab, then on the Tested Entity Result tab.

6

Choose an entry and click Properties. The Tested Entity Result (Edit) form opens.

7

Click on the Results tab to display the validation test results.

8

If you need to compare two test results from the same type of test, choose the two test results and click Compare; the Difference form opens. Otherwise go to Step 11 .

9

Compare the test results.

10

Close the Difference form.

11

Close the Tested Entity Result form.

12

Save the changes and close the forms.

END OF STEPS

77.29 To add or modify FIB entries associated with a VPLS

77.29.1 Steps

1

Choose Manage→Service→FIB Entries from the NFM-P main menu.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Forwarding Control tab, then on the FIB Entries tab.

4

Click Search. A list of FIB entries appears.

5

Add or modify FIB entries:

a. To add a FIB entry:
   1. Click Create. The FIB Entry (Create) form opens.
   2. Configure the required parameters.
   3. Select an interface, service circuit, or endpoint from the list on the L2 Interfaces, Service Circuits, or Endpoints tab.
   4. Save the changes and close the form.

b. To modify FIB entries:
   1. Choose a FIB entry from the list and click Properties. The FIB Entry (Edit) form opens.
   2. Configure the parameters and view the information as required.
   3. Save the changes and close the form.

END OF STEPS

77.30 To list FIB entries associated with a VPLS

77.30.1 Steps

1

Choose Manage→Service→FIB Entries from the NFM-P main menu.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Sites tab, then on the VPLS, B-VPLS, or I-VPLS tabs.

4

Select the site can click Properties. The VPLS Site (Edit) form opens.

5

Click on the Forwarding Control tab, then on the FIB Entries tab.

6

Click Resync for the FIB entries on the right side of the form. Resync on the bottom of the form is for re-synchronizing the entire VPLS site.

7

Click Find. A list of FIB entries appears.

END OF STEPS

77.31 To view IGMP snooping queriers

77.31.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, click on a site. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the IGMP Snooping tab.

5

Click on the MRouters tab. The MRouters table displays a list of IGMP snooping queriers and their properties.

6

Click Refresh to view periodic updates to the Mrouters table.

7

Close the forms.

END OF STEPS

77.32 To view MLD snooping queriers

77.32.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expands Sites and click on a site. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then the MLD Snooping tab.

5

Click on the MRouters tab. The MRouters table displays a list of MLD snooping queriers and their properties.

6

Click Refresh to view periodic updates to the Mrouters table.

7

Close the forms.

END OF STEPS

VPLS site management procedures

77.33 To configure a VPLS site

77.33.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on the Site icon and choose Create VPLS Site and select a site, or expand the Site icon and choose Properties. The VPLS Site (Create|Edit) form opens.

Note: The options to create either a B-Site or an I-Site are used when you are configuring a VPLS that utilizes Provider Backbone Bridging. See 77.2.4 “Provider Backbone Bridging in VPLS” (p. 2215) in “VPLS management overview” (p. 2210) for more information.

4

Configure the required general parameters.

The following considerations apply:

- To configure the Enable IP Interface Binding parameter on the 7450 ESS, you must enable mixed mode on the NE.
- For 7210 SAS NEs that support RVPLS, the Enable IP Interface Binding parameter is configurable when the RVPLS parameter is selected. To bind the site to an IES, see 78.30 “To bind an IES L3 access interface to a VPLS site or VPLS I-site” (p. 2475). To bind the site to a VPRN, see 79.98 “To bind a VPRN L3 access interface to a VPLS site or VPLS I-site” (p. 2673).
- Tunnel Fault Notification is configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation.
- If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept, in order to accept the fault propagation from the tunnel facility MEP.
- The Customer VID parameter is configurable only when the SAP Type parameter is set to dot1q-preserve.
- You can configure a VLAN range SAP only when the SAP Type parameter is set to dot1q-range.
- For MVPLS sites, the Allow L2Pt Xstp BPDU parameter is read-only.
- The Enable VXLAN ECMP parameter can only be configured if all cards in the chassis are FP3 or higher.

5

Configure the parameters in the GSMP panel.

6

Configure the parameters in the ETree panel.

You must enable the Etree Enabled parameter before you can configure ETree tags at the SAP or SDP spoke levels.

7

Configure the parameters in the Pim Snooping panel.

You must enable the PIM Snooping Enabled parameter before you can configure PIM snooping parameters on the Multicast tab and at the SAP and SDP spoke levels.

8

Configure the parameters in the Routed VPLS panel.

9

Configure the parameters in the Load Balancing Panel.

10

Save the changes and close the forms.

END OF STEPS

77.34 To configure a GNE site on a VPLS service

77.34.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS service and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on VPLS Service and choose Create GNE Site.

4

Select a site. The GNE Site (Create) form opens.

5

Configure the required parameters.

6

Click Apply.

7

To configure an interface for the GNE site:

1. Click on the GNE Service Interfaces tab and then click on Create. The GNE Service Interface (Create) form opens.
2. Configure the required parameters.
3. Click on the Ports tab and select a Generic NE interface.
4. Configure the required parameters.
5. Save the changes and close the form.

8

Save the changes and close the forms.

You can use the topology maps to view the service. See Chapter 4, “Topology map management” for more information about service topology maps.

END OF STEPS

77.35 To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site

77.35.1 Purpose

Configure MFIB, STP, FIB, and MAC learning protection parameters for a VPLS site.

77.35.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure MFIB, STP, FIB, and MAC learning protection parameters and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Forwarding Control tab. The MFIB tab is displayed.

5

Configure the required MFIB parameters.

6

Click on the STP tab to configure the required STP parameters.

Note: Nokia STP in a VPLS interoperates with customer STP implementations as a mechanism for loop detection and prevention. The bridge-level parameters balance the STP resiliency and speed of convergence. Modifying the bridge-level parameters must be done within the constraints of the following formulas:

- 2 (Bridge Forward Delay – 1.0 s) >= Bridge Max Age
- Bridge Max Age >= 2 (Bridge Hello Time + 1.0 s)
  If you are configuring an MVPLS site, set the STP Mode parameter to RSTP or MSTP, depending on the MVPLS type. The MSTP option is available only if you are creating an MVPLS. See 77.7 “To create an MVPLS” (p. 2251) for more information about creating an MVPLS.
  MSTP is configurable only on the 7450 ESS, 7750 SR, and 7950 XRS.

7

To configure FIB parameters:

1. Click on the FIB tab.
2. Configure the required parameters.

8

To configure the list of protected MAC addresses:

1. Click on the MAC Protection tab and click Create. The MAC Protection (Create) form opens.
2. Configure the Protected Mac Address parameter.
3. Save the changes and close the form.

9

If you are configuring an MVPLS site that requires MSTP:

MSTP is configurable only on the 7450 ESS, 7750 SR, and 7950 XRS.

1. Click on the MSTP tab.
2. Configure the required parameters.
3. Click on the MST Instances tab and click Create. The MST Instance (Create) form opens.
4. Configure the required parameters.

5. Click on the VLAN Ranges tab. Click Create. The MST Instance Managed VLAN range (Create) form opens.
6. Configure the required parameters.
   - Min. VLAN Tag
   - Max. VLAN Tag
7. Save the changes and close the forms.

10

Save the changes and close the forms.

END OF STEPS

77.36 To configure SHCV for a VPLS site

77.36.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure SHCV and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Subscriber Management tab.

5

Select the SHCV Enabled parameter to enable SHCV.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.37 To configure a default gateway for a VPLS site

77.37.1 Purpose

Configure a default gateway for the site.

77.37.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria and click Search. A list of services appears.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand the site on which you want to configure the default gateway and click Properties. The VPLS Site (Edit) form opens.

5

Click on the Default Gateway tab.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.38 To configure ingress multicast forwarding on a VPLS site

77.38.1 Purpose

Configure ingress multicast forwarding.

Note: An Operational Channels tab appears when you access the VPLS Site form in the Edit mode. It displays data for the operational channels when traffic from a specific multicast source for a specific multicast group passes through the service. You must click Search to refresh the data. See Chapter 49, “Policies overview” for information about listing the operational channel parameters.

77.38.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure ingress multicast forwarding and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the Mcast Path Mgmt tab.

5

Select a multicast info policy.

6

Save the changes and close the forms.

END OF STEPS

77.39 To configure a provider tunnel for a VPLS site

77.39.1 Purpose

When you configure an inclusive provider tunnel, multicast and broadcast Ethernet frames are always forwarded over a P2MP LSP, defined as the Inclusive Provider Multicast Service Interface (I-PMSI). When a provider tunnel is not configured, these frames are replicated at the ingress PE and a copy of each frame is forwarded over the P2P PW to each destination PE participating in the VPLS instance. Provider tunnels can be configured on VPLS and B-VPLS sites.

Note: Provider tunnels are supported on the 7750 SR-7, 7750 SR-12, 7750 SR-c4, 7750 SR-c12, 7750 SR-12E, and 7450 ESS NEs in all chassis modes, as well as on the 7540 ESS, provided all network IP interfaces are on IOM3/IMM ports (chassis mode D).

77.39.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure a provider tunnel and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Provider Tunnel tab and enable the Enable Provider Tunnel parameter.

5

Configure the required parameters.

The LSP Template parameter is only configurable when the Type parameter is set to RSVP.

The Administrative State parameter is configurable only when the LSP Template parameter has been configured.

6

Save the changes and close the forms.

END OF STEPS

77.40 To configure service tunnel required bandwidth for a VPLS site

77.40.1 Purpose

To configure service tunnel required bandwidth, you must enable the Multi-Segment Tunnel Selection and Service Bandwidth Management check boxes on the Services tab on the NFM-P System Preferences form. For more information, see the procedure to configure NFM-P system preferences in the *NSP System Administrator Guide*.

77.40.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure service tunnel required bandwidth and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Bandwidth tab.

5

Enable the Override Service Configuration parameter.

6

Configure the required parameters.

If the Bandwidth Method parameter is set to Input Manually, enter a required bandwidth value for each CoS (see CoS 0 Bandwidth - CoS 7 Bandwidth).

7

Save the changes and close the forms.

END OF STEPS

77.41 To configure IGMP snooping on a VPLS site

77.41.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure IGMP snooping and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the IGMP Snooping tab.

5

Configure the required parameters.

If the Administrative State is down, packets will be forwarded along the P2P PWs, otherwise they will be forwarded along the P2MP LSP instance.

6

Save the changes and close the forms.

END OF STEPS

77.42 To configure PIM snooping on a VPLS site

77.42.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure PIM snooping and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the PIM Snooping tab.

5

Configure the required parameters.

If the node has associated Neighbour and (S,G) Entries, then the same is displayed under EVPN MPLS sub-tab.

6

Save the changes and close the forms.

END OF STEPS

77.43 To create an endpoint for redundancy (dual homing) on a VPLS site

77.43.1 Note

You cannot create a VPLS endpoint on a site that has an active or inactive MC ring SAP. See Chapter 45, “MC ring groups” for more information.

77.43.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to create an endpoint for redundancy and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Endpoints tab and click Create. The VPLS Endpoint (Create) form opens.

5

Configure the required parameters.

The EndPoint ID parameter is configurable only if you set the Endpoint Type parameter to Multi Chassis.

The Restrict Protected Source Action parameter is configurable only if you enable the Restrict Protected Source parameter.

6

Save the changes and close the forms.

END OF STEPS

77.44 To configure an SHG on a VPLS site

77.44.1 Purpose

You must configure an SHG or RSHG if you plan to create a spoke circuit from this VPLS site to a VLL or to another VPLS.

77.44.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure SHG and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Split Horizon Groups tab and click Create. The Site, New Split Horizon Group (Create) form opens.

5

Configure the required parameters.

The Restrict Protected Source Action parameter is configurable only if you enable the Restrict Protected Source parameter.

6

Save the changes and close the forms.

END OF STEPS

77.45 To configure an EVPN gateway on a VPLS site

77.45.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure gateway and click Properties. The VPLS Site (Edit) form opens.

4

Click on the EVPN MCast Gateway tab and click Create. The EVPN Gateway (Create) form opens.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

77.46 To configure proxy ARP for a VPLS site

77.46.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure Proxy ARP and click Properties. The VPLS Site (Edit) form opens.

4

Enable the Enable Proxy ARP parameter.

5

Click on the Proxy ARP tab.

6

Configure the required general parameters.

Note: The Administrative State parameter is not configurable on the 7210 SAS-Mxp as it is automatically updated based on the state of global System EVPN Proxy ARP ND parameter. See 12.45 “To configure the global EVPN proxy ARP and node discovery on an NE” (p. 377) for more information.

7

Configure the parameters in the Duplicate Detect panel.

8

To configure a static proxy ARP entry:

1. Click on the Static tab.
2. Click Create. The Static Proxy ARP (Create) form opens.
3. Configure the IP Address and IEEE Address parameters.
4. Save the changes and close the form.

9

Before a dynamic proxy ARP entry can be configured, a MAC List must be configured on the NE; see 12.15 “To configure proxy ARP and proxy node discovery for an NE” (p. 351)

To configure a dynamic proxy ARP entry:

1. Click on the Dynamic tab.
2. Click Create. The Dynamic Proxy ARP (Create) form opens.
3. Configure the IP Address and MAC List parameters.
4. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

77.47 To configure proxy node discovery for a VPLS site

77.47.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure MVR and click Properties. The VPLS Site (Edit) form opens.

4

Enable the Enable Proxy Node Discovery parameter.

5

Click on the Proxy ND tab.

6

Configure the required general parameters.

Note: The Administrative State parameter is not configurable on the 7210 SAS-Mxp as it is automatically updated based on the state of global System EVPN Proxy ARP ND parameter. See 12.45 “To configure the global EVPN proxy ARP and node discovery on an NE” (p. 377) for more information.

7

Configure the parameters in the Duplicate Detect panel.

8

To configure a static proxy node discovery entry:

1. Click on the Static tab.
2. Click Create. The Static Proxy ND (Create) form opens.
3. Configure the IP Address, IEEE Address and ND Type parameters.
4. Save the changes and close the form.

9

Before a dynamic proxy node discovery entry can be configured, a MAC List must be configured on the NE; see 12.15 “To configure proxy ARP and proxy node discovery for an NE” (p. 351)

To configure a dynamic proxy node discovery entry:

1. Click on the Dynamic tab.
2. Click Create. The Dynamic Proxy ND (Create) form opens.
3. Configure the IP Address and MAC List parameters.
4. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

77.48 To configure MVR for a VPLS site

77.48.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure MVR and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the MVR tab.

5

Configure the required parameters.

The Administrative State parameter specifies whether a site is an MVR VPLS site.

After the multicast package policy is applied to the MVPLS, the policy is distributed as the routing policy to all MVR-capable MVPLS sites. If you apply another package policy to the site, the new policy is distributed to the site. The previously distributed policy remains on the site.

6

If you deselect the Use Component Package Policy parameter, you must specify a multicast package policy to associate with the MVR VPLS site. Perform one of the following.

a. Select a multicast package policy.

b. Manually enter a multicast package policy name as the Routing Policy Name parameter value.

7

View the VPLS SAPs that use the MVR VPLS site.

1. Click on the User MVR SAPs tab.
2. Click Search to list the VPLS SAPs.

8

Save the changes and close the forms.

END OF STEPS

77.49 To configure a GSMP group on a VPLS site

77.49.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure a GSMP group and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the GSMP tab.

5

Click Create. The GSMP Group (Create) form opens.

6

Configure the required parameters.

7

Click on the GSMP Group Neighbor tab.

8

Click Create. The GSMP (Create) form opens.

9

Configure the required parameters.

The Priority Precedence parameter is configurable when the Priority Type parameter is set to Precedence.

The Priority Dscp parameter is configurable when the Priority Type parameter is set to Dscp.

10

Save the changes and close the forms.

END OF STEPS

77.50 To configure L2 management interfaces on a VPLS site

77.50.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria and click Search. A list of services appears.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand the site on which you want to configure L2 management interfaces and click Properties. The VPLS Site (Edit) form opens.

5

Click on the L2 Management Interfaces tab.

6

Click Create. The VPLS L2 Management Interface, Service - Subscriber (Create) form opens.

7

Configure the required parameters.

Note: The NFM-P does not support the configuration of static ARP on a VPLS (MVPLS) L2 management interface. If required, this configuration must be done using CLI. Static ARP configuration must be removed from the interface using CLI before this interface can be removed using the NFM-P.

8

Click on the Addresses tab and click Create. The IP Address, Service - Subscriber (Create) Form opens.

9

Configure the required parameters.

10

Save the changes and close the forms.

END OF STEPS

77.51 To configure MLD snooping on a VPLS site

77.51.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure MLD snooping and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Multicast tab, then on the MLD Snooping tab.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

77.52 To create a Virtual MEP on a VPLS site

77.52.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to create a virtual MEP and click Properties. The VPLS Site (Edit) form opens.

4

Click on the OAM tab, then the ETH-CFM tab.

5

Configure the Tunnel Fault Notification parameter in the Facility MEPs panel.

6

Click Create in the Virtual MEP panel. The MEP (Create) form opens.

7

Select a MEG in the Maintenance Entity Group panel.

8

Configure the parameters in the MEP panel.

9

Configure the parameters in the CCM panel.

The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

10

Select a MEG sub-group in the MEG Sub-Grouping panel.

11

Configure the required Test Generation Options parameters.

Note: The Test Generation Options parameters are only displayed when you set the Direction parameter to Up.

12

Save the changes and close the forms.

Note: After the site is created and is operational, you can open the MEP (Edit) form and view an information field on the General tab that displays the number of Continuously Running Tests currently on the MEP.

END OF STEPS

77.53 To configure MVR for MLD on a VPLS site

77.53.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria and click Search. A list of services appears.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand the site on which you want to configure MVR for MLD and click Properties. The VPLS Site (Edit) form opens.

5

Click on the Multicast tab.

6

Click on the MVR (MLD) tab.

7

Configure the required parameters.

The Administrative State parameter specifies whether the site is an MVR VPLS site.

The Routing Policy Name parameter is configurable when the Use Component Package Policy is disabled.

After a multicast package policy is applied to an MVPLS, the policy is distributed as the routing policy to all MVR-capable MVPLS sites. If you apply another package policy to the site, the new policy is distributed to the site and the previously distributed policy remains on the site.

8

If you deselect the Use Component Package Policy parameter, you must specify a multicast package policy to associate with the MVR VPLS site. Perform one of the following.

a. Select a multicast package policy.

b. Manually enter a multicast package policy name as the Routing Policy Name parameter value.

9

Click on the User MVR SAPs tab to view the VPLS SAPs that use the MVR VPLS site and click Search to list the VPLS SAPs.

10

Save the changes and close the forms.

END OF STEPS

77.54 To configure IGMP host tracking on a VPLS site

77.54.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure IGMP host tracking and click Properties. The VPLS Site (Edit) form opens.

4

Click on the IGMP Host Tracking tab.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

77.55 To configure WLAN GW L2 wholesale forwarding on a VPLS site

77.55.1 Purpose

Perform this procedure to enable L2 wholesale forwarding on a VPLS service that is bound to a WLAN GW as part of a L2 wholesale-retail configuration.

77.55.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure a VPLS SAP template policy. Click Properties. The VPLS Site (Edit) form opens.

4

Click on the WLAN GW tab.

5

Select a SAP template policy (if required) and set the Administrative State parameter to Enabled.

6

Save the changes and close the forms.

END OF STEPS

77.56 To configure a non-system IP address VXLAN termination

77.56.1 Purpose

Use this procedure to configure a non-system IP address for VXLAN termination. The address terminating the VXLAN tunnels can be an IPv4 or IPv6 address, and the service where VXLAN is terminated can be:

- a VPLS service
- an R-VPLS service connected to a VPRN service

VXLAN is supported on 7450 ESS, 7750 SR, and 7950 XRS nodes in chassis mode D.

77.56.2 Steps

1

Create an FPE on the NE, with an SPD ID range configured and VXLAN termination enabled. See 12.41 “To create an FPE” (p. 374)

2

Create a VXLAN tunnel termination:

1. On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.
2. Click on the Globals tab, then the Service sub-tab, then the VXLAN sub-tab, then the Tunnel Termination sub-tab.
3. Enter the tunnel termination IP address. The address must be a loopback interface address and a non-system IP address.
4. In the FPE panel, click Select and choose the FPE ID created in the previous step.
5. Click Apply and close the form.

3

Create a VPLS Site. See 77.33 “To configure a VPLS site” (p. 2294).

In the VXLAN Tunnel End Point panel, click Select and choose the VXLAN tunnel termination created in the previous step.

4

Create a system interface with an IPv4 address at the NE routing instance. See 27.17 “To create an L3 network interface on a routing instance” (p. 856).

5

Configure the assisted replication IP address with the address configured in the previous step.

1. On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.
2. Click on the Globals tab, then the Service sub-tab, then the General sub-tab.
3. Click Select and choose the interface created in the previous step, or enter the IP address of the interface.

6

Create a VXLAN identifier at the VPLS service:

1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
2. Choose a VPLS service and click Properties. The VPLS Service (Edit) form opens.
3. Click on the BGP sub-tab, then the EVPN sub-tab.
4. Configure the parameters in the VXLAN panel.

   To configure the Assisted Replication type parameter as Replicator, step Step 5 must be completed.

END OF STEPS

77.57 To configure EVPN on a VPLS site

77.57.1 Purpose

Use this procedure to configure EVPN parameters on a VPLS site for EVPN -VXLAN or EVPN-MPLS. VXLAN is supported on 7450 ESS, 7750 SR, and 7950 XRS nodes in chassis mode D.

77.57.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites and click on the site on which you want to configure EVPN.

4

Click on the BGP tab, and then on the EVPN sub-tab.

5

Configure BGP EVPN.

1. Click Create in the BGP EVPN panel. The BGP EVPN (Create) form opens.
2. Configure the parameters, as required.
3. Click Create in the VXLAN panel and configure the VXLAN parameters on the form that appears.

   If two EVPN VXLAN instances are created on the same VPLS, different BGP IDs should be assigned to each VXLAN.
4. Click Create in the MPLS panel and configure the MPLS parameters on the form that appears.

   To allow for both EVPN VXLAN and EVPN MPLS to be in an Admin Up state on the same VPLS, different BGP IDs should be assigned to VXLAN and MPLS.
5. Click Create in the Segment Routing V6 panel and configure the parameters on the form that appears.
6. Click OK and confirm to close the form.

6

Configure a network identifier for the VXLAN.

1. Click Create in the VXLAN panel. The VXLAN (Create) form opens.
2. Configure the network identifier.
3. Select a policy in the Network Ingress QoS Policy panel.
4. Select a queue group in the Forwarding Plane Queue Group panel.
5. Configure the Forwarding Plane Queue Group Instance ID parameter.
6. Configure the Assisted Replication Type parameter as needed. An assisted replication IP address must be previously configured in the NE properties.
7. Click OK and confirm to close the form.

7

Create one or more EVPN static MACs.

1. Click on the EVPN Static MAC tab, and click Create. The Conditional Static Mac (Create) form opens.
2. Configure a MAC address for the static MAC.
3. Select a SAP in the SAP panel.
4. Select a spoke SDP binding in the Spoke DSP Binding panel.
5. Click OK and confirm to close the form.

8

Save and close the form.

END OF STEPS

77.58 To configure segment routing v6 on a VPLS site

77.58.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to configure segment routing with IPv6 and click Properties. The VPLS Site (Edit) form opens.

4

Click on the Segment Routing V6 tab.

5

Create or select an entry and click Properties. The Segment Routing V6 (Create/Edit) form opens.

6

Configure the required parameters in the General tab.

7

Create or select an entry in the Locator panel and click Properties. The SRv6 Function (Create/Edit) form opens.

1. Associate a locator under Locator panel.
2. Configure the parameters under End Function panels.
3. Save your changes and close the form.

8

Save your changes and close the forms.

END OF STEPS

77.59 To configure PBB-EVPN on a VPLS site

77.59.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites, click on the B-VPLS site on which you want to configure PBB-EVPN, and click Properties.

4

Click on the Backbone tab. Configure a source MAC address and enable the Use Ethernet Segment Backbone MAC parameter.

5

Click on the BGP tab, and then on the EVPN sub-tab.

6

Click the Create button in the BGP EVPN panel and configure the parameters on the form that appears.

7

Add one or more ISID route targets for the PBB-EVPN by performing the following.

1. Click on the ISIS Route Target sub-tab and click Create. The ISIS Route Target (Create) form opens.
2. Configure the required parameters.
3. Click OK to close the form.

8

Configure static MACs for the PBB-EVPN by performing the following.

1. Click on the Static MACs tab and click Create. The Conditional Static Mac form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

9

Configure ISID policies for the PBB-EVPN by performing the following.

1. Click on the ISID Policy tab and click Create. The ISID Policy Range Entry form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

10

Save your changes and close the forms.

END OF STEPS

77.60 To configure a black hole MAC address on a VPLS site

77.60.1 Purpose

Use this procedure to configure a conditional static black hole MAC address FDB entry on a VPLS service site to mitigate potential DOS, DDOS, and worm attacks and to quarantine hostile traffic. This feature is similar to a black hole static-route for VPRNs but is not associated with any particular SAP or SDP binding. If there is a hit on the black hole FDB entry, all frame packets are immediately discarded in the data path to a null route.

You can also use this procedure to configure a black hole MAC address for IP duplicate detection and anti-MAC address spoofing by enabling the VPLS services with a proxy ARP or proxy ND. The AS-MAC address provides a method to push traffic to a given IP address when a duplicate IP address is detected. However, the NFM-P drops the traffic addressed to the AS-MAC address.

The feature is supported on 7450 ESS, 7750 SR, and 7950 XRS devices.

77.60.2 Steps

Create a VPLS and enable BGP-EVPN

1

Create a VPLS for a supported device type if required; see 77.5 “To create a VPLS” (p. 2249) .

2

Create a VPLS site for the VPLS; see 77.33 “To configure a VPLS site” (p. 2294) .

3

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

4

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

5

On the VPLS Service tree, expand the Sites icon and click on the site on which you want to configure a black hole MAC address and click Properties. The VPLS Site form opens.

6

Click on the BGP tab, and then on the EVPN sub-tab.

7

Click the Create button in the BGP EVPN panel and configure the parameters on the form that appears.

Configure the static MAC address and black hole option

8

Click on the EVPN Static MAC sub-tab and click Create. The Conditional Static Mac form opens.

9

Configure a MAC address for the static MAC.

10

For Locale parameter, select the Blackhole option and click OK to close the form.

11

Save and close the form.

Enable a static black hole for a Proxy ARP or Proxy ND

12

Configure a proxy ARP (address Resolution Protocol) for a VPLS site; see 77.46 “To configure proxy ARP for a VPLS site” (p. 2306) .

13

Configure a proxy ND (node discovery) for a VPLS site; see 77.47 “To configure proxy node discovery for a VPLS site” (p. 2307) .

14

On the VPLS Service tree, expand the Sites icon and click on the site on which you want to configure a black hole MAC address and click Properties. The VPLS Site form opens.

15

Perform one of the following:

a. For a Proxy ARP, click on the Proxy ARP tab and configure the Anti-Spoofing MAC Address parameter and Static Black-Hole parameter on the Proxy ARP sub-tab.

b. For a Proxy ND, click on the Proxy ND tab and configure the Anti-Spoofing MAC Address parameter and Static Black-Hole parameter on the Proxy HD sub-tab.

16

Save and close the form.

END OF STEPS

77.61 To enable SPB on a control B-VPLS site

77.61.1 Purpose

Perform the following procedure to enable SPB on the B-VPLS and configure the B-site as a control B-VPLS.

Control and user B-VPLS SAPs and SDP bindings are fate-shared. You must configure these SAPs and SDP bindings to use the same links and resources. SAPs and spoke SDP bindings that do not use the same links and resources (non-fate-shared) can exist in the user B-VPLS configuration but do not carry traffic.

Consider the following when you enable SPB on a B-VPLS:

- B-VPLS mesh SDP bindings are not supported in control or user SPB B-VPLSs.
- STP can be configured but not enabled.
- two Dot1Q SAPs on the same physical port cannot exist in a user or control SPB B-VPLS.
- Two QinQ SAPs with same top bit on the same physical port cannot exist in a user or control SPB B-VPLS.
- Control or user SPB B-VPLS SAPs cannot exist on physical ports that are managed by MSTP.
- MVPLSs and SPB B-VPLSs cannot share the same spoke SDP tunnel.
- .0 and 0.* SAPs cannot exist on a user or control SPB B-VPLS.
- SPB B-VPLS services that support multiple I-VPLSs do not share forwarding IDs. A B-VPLS that has I-VPLS services must have a different forwarding ID.
- A control SPB B-VPLS can be associated with none or many user SPB B-VPLSs.
- A user SPB B-VPLS must be associated with a control SPB B-VPLS.
- MVPLS and SPB B-VPLS cannot share the same spoke SDP tunnel.
- SHG is not supported.
- G803.1 tunnels are not supported.
- G.8032 access is not supported.

Note: All of the B-sites in the VPLS must be configured as control B-VPLS sites.

77.61.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the VPLS on which you want to enable SPB and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on the B-site and choose Properties. The VPLS Site (Edit) form opens.

4

Click on the Backbone tab, then on the SPB tab.

5

Set the SPB Mode parameter to Control B-VPLS.

6

Click Create in the SPB Instance for Control Service panel. The SPB (Create) form opens.

Note: The SPB instance configuration, except the ISIS instance and the Bridge Priority parameter on the Level 1 tab, must be identical on all the sites in a control B-VPLS.

7

Configure the required general SPB control instance parameters.

8

Configure required the SPB ISIS parameters in the ISIS panel.

The Overload Timeout (seconds) parameter is configurable only if the Overload parameter is set to Enabled.

The Overload On Boot Timeout (seconds) parameter is configurable only if the Overload On Boot parameter is set to Enabled.

9

Click on the Level 1 tab and configure the level 1 area parameters.

10

Create the IS-IS interface, or SPB link, of an access interface.

If you enabled the Automatic Enable Interface in Step 7 , all of the L2 access interfaces on the site and all of the L2 access interfaces that are subsequently created on the B-site are automatically enabled for SPB.

1. Click on the Access Interfaces tab.
2. Click Create. The Access Interface, ISIS Instance (Create) form opens.
3. Select an L2 access interface in the Interface panel.
4. Configure the required general parameters.
5. Click on the Behavior tab and configure the required parameters.
6. Click on the Level 1 tab and configure the required parameters.
7. Save your changes and close the form.

11

Configure SPB on the B-VPLS spoke SDP binding.

If you enabled the Automatic Enable Interface in Step 7 , all of the spoke SDP bindings on the site and all of the spoke SDP bindings that are subsequently created on the B-site are automatically enabled for SPB.

1. Click on the Spoke SDP Bindings tab.
2. Click Create. The Spoke SDP Binding Interface, ISIS Instance (Create) form opens.
3. Select a spoke SDP binding in the Circuit panel.
4. Configure the required general parameters.
5. Click on the Behavior tab and configure the required parameters.
6. Click on the Level 1 tab and configure the required parameters.
7. Save your changes and close the form.

12

To create one or more ECT-to-FID range mappings:

1. Click on the ECT Algorithm to FID Range Mappings tab.
2. Click Create. The ECT Algorithm to FID Range Mapping (Create) form opens.
3. Configure the required FID range parameters.
4. Save your changes and close the form.

   Note:

   The ECT-to FID range mapping parameter values must be identical on all the sites in the control SPB B-VPLS. In addition, each site must have the same number of ECT-to-FID range mapping objects.

13

Save your changes and close the forms.

END OF STEPS

77.62 To enable SPB on a user B-VPLS site

77.62.1 Purpose

Perform the following procedure to enable SPB on the B-VPLS and configure the B-site as a user B-VPLS. A control SPB B-VPLS service must already exist. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) for more information about how to create a control SPB B-VPLS.

Control and user B-VPLS SAPs and SDP bindings are fate-shared. You must configure these SAPs and SDP bindings to use the same links and resources. SAPs and spoke SDP bindings that do not use the same links and resources (non-fate-shared) can exist in the user B-VPLS configuration but do not carry traffic.

Consider the following when you enable SPB on a B-VPLS:

- B-VPLS mesh SDP bindings are not supported in control or user SPB B-VPLSs.
- STP can be configured but not enabled.
- two Dot1Q SAPs on the same physical port cannot exist in a user or control SPB B-VPLS.
- Two QinQ SAPs with same top bit on the same physical port cannot exist in a user or control SPB B-VPLS.
- Control or user SPB B-VPLS SAPs cannot exist on physical ports that are managed by MSTP.
- MVPLSs and SPB B-VPLSs cannot share the same spoke SDP tunnel.
- .0 and 0.* SAPs cannot exist on a user or control SPB B-VPLS.
- SPB B-VPLS services that support multiple I-VPLSs do not share forwarding IDs. A B-VPLS that has I-VPLS services must have a different forwarding ID.
- A control SPB B-VPLS can be associated with none or many user SPB B-VPLSs.
- A user SPB B-VPLS must be associated with a control SPB B-VPLS.
- MVPLS and SPB B-VPLS cannot share the same spoke SDP tunnel.
- SHG is not supported.
- G803.1 tunnels are not supported.
- G.8032 access is not supported.

77.62.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select the VPLS on which you want to enable SPB and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, right-click on the B-site and choose Properties. The VPLS Site (Edit) form opens.

4

Click on the Backbone tab, then the SPB tab.

5

Set the SPB Mode parameter to User B-VPLS.

6

Select a user in the User Service panel to associate the user B-VPLS with a control B-VPLS.

7

Configure the Forwarding Identifier (FID) parameter.

Note: The Forwarding Identifier (FID) parameter must be the same for all the user B-VPLS sites within the SPB B-VPLS.

8

Save your changes and close the forms.

Note: You can perform an RCA audit on the SPB user B-VPLS service to ensure a correct configuration.

END OF STEPS

77.63 To view the last cleared BFD statistics and sessions on a VPLS site

77.63.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPLS service and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, click on the site on which you want to view the last cleared BFD statistics and sessions.

4

Click on the Clear Status tab and view the five last cleared BFD statistics and the five last cleared BFD sessions.

5

Close the forms.

END OF STEPS

77.64 To enable the automatic selection of an RD on a VPLS site

77.64.1 Purpose

Since an RD must be unique on each PE in the network, you can allocate either a route distinguisher that you manually select or an NE-selected route distinguisher for each service. When you configure an auto-RD on a VPLS site, B-site, or I-site, a Type-1 RD is automatically allocated by the NE based on the community range that you configure.

77.64.2 Steps

1

Before you configure a site for auto-RD selection, you must:

1. Enable BGP on the routing instance of the NE. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.
2. Enable the automatic selection of an RD on the NE and specify the community range. See 12.11 “To enable the automatic selection of an RD on an NE” (p. 349) .

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the navigation tree, right-click on VPLS Service and choose Create VPLS Site, or right-click on a site and choose Properties. The VPLS Site (Create|Edit) form opens.

5

Click on the BGP tab, then on the General tab, and enable the Enable BGP parameter.

6

Configure the Auto Route Distinguisher parameter in the RD panel.

The Operational RD is displayed after you apply the changes to the site and to the service.

7

Save the changes and close the forms.

END OF STEPS

77.65 To create a static B-MAC on a B-VPLS site

77.65.1 Purpose

Perform the following procedure to create a static B-MAC to allow remote PBB Epipe VLL services to connect to an SPB network. You can create a static B-MAC on a B-L2 access interface or spoke SDP binding on a control or user B-VPLS site. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for information about how to enable SPB on a control or user site.

77.65.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, click on the B-VPLS site. The Site (Edit) form opens.

4

Click on the Static MACs tab.

5

Click Create. The Conditional Static MAC (Create) form opens.

6

Configure the MAC Address parameter.

7

Perform one of the following:

a. Select a SAP.

b. Select a spoke SDP binding.

8

Save the changes and close the forms.

END OF STEPS

77.66 To create an ISID policy on a control or user B-VPLS site

77.66.1 Purpose

An ISID policy defines a group of actions for individual ISIDs or ISID ranges. You can define whether the default multicast tree is used and whether the advertisement of ISIDs in SPB is suppressed when I-VPLS or static ISIDs are used for unicast services. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for information about how to enable SPB on a control or user site.

77.66.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, click on the B-VPLS site. The Site (Edit) form opens.

4

Click on the ISID Policy tab.

5

Click Create. The ISID Policy Range Entry (Create) form opens.

6

Configure the required parameters.

You cannot set the Use Default Multicast Tree and the Advertise Local parameters to False at the same time.

7

Save the changes and close the forms.

END OF STEPS

VPLS access interface management procedures

77.67 To create a VPLS or MVPLS L2 access interface

CAUTION

Service Disruption

The creation of a SAP that uses the same port and encapsulation values as an existing inactive MSAP fails under the following conditions:

- If you try to use the NFM-P to create a SAP, the configuration fails and the NFM-P displays an error message.
- If you use a CLI to create a SAP in a service other than the service that contains the MSAP, the configuration succeeds but the MSAP is inactive until the regular SAP is deleted. Although the NFM-P displays the SAP and MSAP, the MSAP remains inactive and consumes resources.
- If you use a CLI to create a SAP in the service that contains the MSAP, the SAP creation fails.

Nokia recommends that you delete an inactive MSAP from the NFM-P if you need to create a regular SAP on the same port using the same encapsulation values.

Note: If 7210 SAS-D, 7210 SAS-E, or 7210 SAS-K sites are connected in a ring network, you must configure an uplink SAP between each pair of sites.

77.67.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site to which you want to add the access interface.

For information on adding a GNE service interface to a VPLS service, see 77.34 “To configure a GNE site on a VPLS service” (p. 2295) .

4

Right-click on L2 Access Interfaces and choose Create VPLS L2 Access Interface. The VPLS L2 Access Interface (Create) form opens.

5

Configure the required parameters.

If you set the SAP Sub Type parameter to Regular, go to Step 10 .

If you set the SAP Sub Type parameter to Capture, the VPLS L2 Access Interface (Create) form refreshes to allow the configuration of the Capture SAP. To create a Capture SAP, see 74.26 “To configure a capture SAP” (p. 2049) .

6

Select an SHG for the interface in the Split Horizon Group panel.

Note: You must configure an SHG or residential SHG for a VPLS if you plan to create a spoke circuit from this VPLS site to a VLL or another VPLS.

7

Select an application profile for the L2 access interface.

8

To associate an AA transit prefix policy with the L2 access interface, select a transit prefix policy in the Transit Prefix Policy panel.

Note: To bind a transit policy to an access interface, a port must already exist on the interface.
You can bind a transit policy to only one access interface or spoke SDP binding per NE.
The transit policy and the application profile must belong to the same application assurance group or partition.

9

To associate a host lockout policy with the L2 access interface, select a host lockout policy in the Host Lockout panel.

10

Select an operational group for the L2 access interface to join as a member. To select an operational group for the L2 access interface to monitor, go to Step 11 .

11

Select an operational group for the L2 access interface to monitor.

12

Click on the Port tab and select a port for the L2 access interface.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.
If you select an Ethernet Tunnel Endpoint, the Port form is refreshed and an Ethernet Tunnel tab is added.
If you are configuring a VLAN range SAP, you must choose a port with dot1q encapsulation type.

If you select the private port of an ISA tunnel, an L2TPV3 tab is added. On the L2TPV3 tab, select VPRN or Base Router in the Router ID parameter and enter the Tunnel Group name.

13

Configure the required parameters.

When the selected port uses dot1q encapsulation, you can enable the Auto-Assign ID parameter to have the Outer Encapsulation Value parameter automatically assigned. If you choose this, the system assigns the lowest unused encapsulation value.

You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter in the User Preferences form. Choose Application→ User Preferences from the main menu.

If the port you have chosen is an Ethernet Tunnel Endpoint, you will be able to set the Outer Encapsulation Value to 8191. This automatically enables the Ethernet Tunnel Endpoint Control SAP parameter.

SAP types X.0 and X.* as well as SAP types *.null and *.* can be configured on the same Q in Q port. You must enable the Enable Q in Q Untagged Sap parameter on the NE which allows the creation of the following default SAP types:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- The SAP type *.* functions as a default SAP for double-tagged frames on a Q in Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

See 12.23 “To enable a Q in Q untagged SAP on an NE” (p. 360) for more information.

If you are configuring the port for an L2 Access Interface in a Control VPLS for an Ethernet ring, the Outer Encapsulation Value and Inner Encapsulation Value parameters for the port in each L2 Access Interface must have the same value as the R-APS Tag (Outer Encapsulation Value) and R-APS Tag (Inner Encapsulation Value) respectively, that you use for the particular path endpoint. This defines the interface as a Control SAP for the Ethernet ring. See 33.18 “To create an Ethernet G.8032 ring” (p. 1212) for additional information.

The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with Q in Q encapsulation.

The Outer Encapsulation Value (VPI) and Inner Encapsulation Value (VCI) parameters are configurable only when the port is an ATM port.

14

If a VLAN Connection Profile policy was created and you want to use it for a VPLS service, configure the parameters in the VLAN Connection Profile panel on the Port tab. See Chapter 62, “VLAN Connection Profile policies” for details.

15

To configure VLAN ranges on a 7210 SAS:

1. Enable the VLAN Range SAP parameter.

You can configure the VLAN Range SAP parameter only when the SAP Type parameter is set to dot1q-range. See 77.33 “To configure a VPLS site” (p. 2294) .

2. Select a VLAN range.

   For more information about VLAN ranges for the 7210 SAS, including restrictions, see Chapter 63, “Connection profile policies” . See 63.4 “To configure a VLAN range for a 7210 SAS VPLS or VLL Epipe service” (p. 1825) for information about how to create a connection profile.

16

Configure the required parameters.

If the port you have chosen is an Ethernet port and uses ATM, dot1q or QinQ encapsulation, you can enable ingress VLAN translation.

17

Configure the Ethernet Tunnel Endpoint Control SAP parameter.

Enabling the Ethernet Tunnel Endpoint Control SAP parameter creates the control L2 Access Interface (also known as a Control SAP). It also automatically sets the value of the Outer Encapsulation Value parameter to 8191.

If you are currently creating a same-fate SAP or an L2 Access Interface for an Ethernet ring, the Ethernet Tunnel Endpoint Control SAP parameter must not be enabled.

18

If you are creating this L2 Access Interface for the Control VPLS of an Ethernet ring or for a data service using the ring, configure the ID parameter in the Ethernet Ring Element section to select the required Ethernet Ring Element. See 33.16 “To configure an Ethernet Ring Element” (p. 1208) for additional information.

19

If you are creating an VPLS and use dot1q or QinQ encapsulation, you can enable ingress VLAN translation. Configure the required parameters in the QinQ VLAN Translation panel.

20

Depending on the port that you have chosen, the Egress Multicast Group tab is configurable.

1. Click on the Egress Multicast Group tab.
2. Select an Egress Multicast Group.

   Only EMGs that have the same egress filter and encapsulation type as the interface are listed.

21

Configure the Enable Split Horizon parameter.

Note: If you are configuring the control SAP for a 7210 SAS-M (in Access Uplink mode) to be used as an interconnected NE in an Ethernet sub-ring, you must enable the Enable Split Horizon parameter.

22

For QTag Manipulation on the 7250 IXR, configure the parameters in the QTag Manipulation panel.

For the Ingress Action parameter, when you select the Push Outer or Replace Outer option, the Outer Tag parameter is available. You must configure the Outer Tag parameter.

If required, click Reset to restore the default values.

23

To assign ingress and egress ACL filters to the interface:

1. Click on the ACL tab.

   Note:

   When you assign ACL filters on a 7210 SAS, you must configure the system resource profile appropriately. See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) for more information.
2. Select the required ACL filter policies.

24

To assign an accounting policy to the interface:

1. Click on the Accounting tab.
2. Select an accounting policy.
3. Configure the required parameters.

   Collect Egress Queue Statistics parameter can be configured only during service creation.

25

To assign a virtual port to the interface:

1. Click on the Virtual Port Name tab.
2. Configure the required parameters.

26

To assign a time of day suite to the interface.

1. Click on the TOD Suite tab.
2. Select a Time Of Day Suite.

   Note:

You cannot assign a ToD suite to a L2 access interface if accounting statistics collection is enabled on the L2 access interface. You must first disable the Collect Accounting Statistics parameter in Step 24 .

SapEgrQosPlcyStats and SapIngQosPlcyStats statistics will only be collected if a Time Of Day Suite is applied on the SAP.

27

The ATM tab is configurable when the interface port is an ATM port. To specify OAM functionality and assign ingress and egress ATM policies to the interface:

1. Click on the ATM tab.
2. Configure the required parameters.
3. Select an ingress ATM policy in the Ingress ATM Policy panel.
4. Select an egress ATM policy in the Egress ATM Policy panel.

28

To configure policy overrides.

1. Click on the Override Policy Items tab.

   Note:

   The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.

   - If you configured a non-HSMDA port, the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.
   - If you configured an HSMDA port, the Access Ingress Queues, Access Egress HSMDA Queues, and Ingress Policer tabs are active.
2. Set the policy overrides, as described in 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) . To configure meter overrides on a 7210 SAS, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) . To configure queue overrides on a 7210 SAS, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

29

Save the changes and close the forms.

30

To configure LAG per-link hashing on the L2 access interface, perform 77.68 “To configure LAG per-link hashing on a VPLS L2 access interface” (p. 2339) .

31

To assign QoS policies or to enable a MAC override address to an L2 access interface, perform 77.69 “To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface” (p. 2340) .

32

To configure scheduling on the L2 access interface, perform 77.71 “To configure scheduling on an L2 access interface” (p. 2345) .

33

To configure BPDU Termination, STP, and FIB parameters for the L2 access interface, perform 77.72 “To configure BPDU Termination, STP, and FIB parameters for the VPLS L2 access interface” (p. 2348) .

34

To assign a DoS protection policy or DDoS protection policy to the L2 access interface, perform 77.73 “To assign a DoS protection policy or DDoS protection policy to the VPLS L2 access interface” (p. 2350) .

35

To configure residential subscriber management on the L2 access interface, perform 77.74 “To configure residential subscriber management for the VPLS L2 access interface” (p. 2351) .

36

To configure an Ethernet tunnel on the L2 access interface, perform 77.75 “To configure an Ethernet tunnel on a VPLS L2 access interface” (p. 2352) .

37

To configure a redundant VLAN range on the L2 access interface, perform 77.76 “To configure a redundant VLAN range on a VPLS L2 access interface” (p. 2353) .

38

To configure IGMP snooping on the L2 access interface, perform 77.77 “To configure IGMP snooping for a VPLS L2 access interface” (p. 2354) .

39

To configure the ARP host for the L2 access interface, perform 77.78 “To configure the ARP host for the VPLS L2 access interface” (p. 2355) .

40

To configure DHCP on the L2 access interface, perform 77.79 “To configure DHCP for the VPLS L2 access interface” (p. 2356) .

41

To configure MVR for the L2 access interface, perform 77.80 “To configure MVR for a VPLS L2 access interface” (p. 2357) .

42

To configure anti-spoofing on the L2 access interface, perform 77.81 “To configure anti-spoofing filters for a VPLS L2 access interface” (p. 2358) .

43

To configure DHCPv6 snooping on the L2 access interface, perform 77.90 “To configure DHCPv6 snooping for a VPLS or MVPLS L2 access interface” (p. 2381) .

44

To associate a MEP to the L2 access interface, perform 77.82 “To create MIPs and MEPs on a VPLS L2 access interface” (p. 2360) .

45

To assign an ANCP policy to the L2 access interface, perform 77.83 “To assign an ANCP policy to a VPLS L2 access interface” (p. 2362) .

46

To configure PIM snooping on the L2 access interface, perform 77.42 “To configure PIM snooping on a VPLS site” (p. 2303) .

47

To configure MLD snooping on the L2 access interface, perform 77.85 “To configure MLD snooping for a VPLS L2 access interface” (p. 2364) .

48

To configure MVR (MLD) on the L2 access interface, perform 77.86 “To configure MVR (MLD) for a VPLS L2 access interface” (p. 2365) .

49

To configure custom object attributes for AA reporting on the L2 access interface, perform 77.23 “To configure custom object attributes for AA reporting” (p. 2282) .

END OF STEPS

77.68 To configure LAG per-link hashing on a VPLS L2 access interface

77.68.1 Purpose

You can configure weighted per-link hashing on a VPLS L2 access interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member. This procedure applies to B-VPLS, I-VPLS and M-VPLS interfaces.

77.68.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Site→L2 Access Interfaces.

4

Right-click on the L2 access interface you want to modify and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the LAG Per Link Hash tab.

6

Configure the Class and Weight parameters.

7

Save your changes and close the forms.

END OF STEPS

77.69 To assign QoS policies or to enable a MAC override address to a VPLS or MVPLS L2 access interface

77.69.1 Before you begin

The available panels and parameters vary depending on the NE, chassis type, and release.

77.69.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interface.

4

Right-click on the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the QoS tab.

6

Configure the Ingress Match QinQ Dot1P parameter.

The Ingress Match QinQ Dot1P parameter is configurable only when the encapsulation type of the port is dot1q, BCP dot1q, or QinQ.

7

Select an ingress policy in the Ingress Policy panel.

Note: If you choose an access ingress policy which has a forwarding class mapped to an ingress queue group, you must ensure that port that you are configuring for the VPLS L2 access interface has the access ingress queue group with the same name created on it. See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

8

Select an ingress queue group template policy in the Forwarding Plane Redirect panel.

9

Configure the Instance ID parameter.

10

Configure the Egress Mark QinQ Top Bits Only parameter.

The Egress Mark QinQ Top Bits Only parameter is configurable only when the encapsulation type of the port is dot1q, BCP dot1q, or QinQ.

11

Configure the required parameters in the Aggregate Rate Limit panel.

12

Select an egress policy in the Egress Policy panel.

Note: If you choose an access egress policy which has a forwarding class mapped to an egress queue group, you must ensure that port that you are configuring for the VPLS L2 access interface has the access egress queue group with the same name created on it.

See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

13

Select an egress queue group template policy in the Port Redirect panel.

Note: Selecting an Egress Queue Group Template Policy permits the redirection of Ethernet traffic packets to a queue ID specified in the egress port queue group of the SAP. The following properties and restrictions apply:

- If an Egress Queue Group Template Policy is specified here, the policy must have port redirection enabled.
- You cannot use policy-based redirection with the queue group when the SAP has SAP-based redirection enabled.
- Port access egress redirection is only supported on Ethernet/LAG ports. It is not supported on SAPs bound on non-Ethernet, Eth-tunnel, or CCAG ports.
- Supported ports include access, hybrid, and HSMDA.
- Queue groups can be applied to SAPs that incorporate LAGs. The LAGs can include port members from just a single card or from multiple cards.
- If you edit a LAG incorporated by the SAP, you cannot remove the last LAG member if a queue group reference exists to the containing SAP.
- You cannot add a secondary LAG member that has a queue group mismatch with primary LAG member.

14

If you are configuring an L2 access interface for a 7705 SAR, or if the port you selected in 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) is not an HSMDA port, go to Step 6 of 77.71 “To configure scheduling on an L2 access interface” (p. 2345) .

15

Configure the Packet Byte Offset (bytes) parameter. Before you can configure the parameter, you must enable the associated Override check box.

16

Select a WRR policy in the Egress HSMDA Override panel.

17

Select an HSMDA egress secondary shaper policy in the Egress HSMDA Override panel.

18

In the Shaper Group panel, select a Shaper Group for the access ingress port or access egress port.

19

In the IXR Specific panel, select an Egress Remark policy, Egress VLAN QoS policy, and shared policer, and configure all other parameters, as required. See 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634) and 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594).

20

To enable a destination MAC override address to the L2 access interface, select the Enable Mac Overwrite check box and enter a MAC address in the IEEE Address parameter.

21

Select an HS secondary shaper in the HS Overrides panel, if required.

22

Save the changes and close the forms.

END OF STEPS

77.70 To assign ingress and egress QoS policies to a VPLS L2 access interface on a 7210 SAS site

77.70.1 Before you begin

You can only assign policies and configure parameters for 7210 SAS NEs when the system resource profile is appropriately configured for the device. See 6.5.13 “System resource profile” (p. 220) in Chapter 6, “Device support” .

The available parameters and policies vary depending on the device type and chassis variant. The configurations that are supported on the site NE are shown on the form.

77.70.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the navigation tree, expand Sites→site→L2 Access Interfaces, right-click on the L2 access interface, and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

If you are configuring an L2 Uplink SAP on a 7210 SAS-K site to preserve the dot1p values of ingress packets, expand to the Uplink SAP, open the properties form, and perform Step 4 and Step 10.

4

Click on the QoS tab and expand the 7210/7250 Specific panel.

5

Select a SAP Access Ingress policy in the Ingress Policy panel.

When you assign an ingress policy, the Color Mode parameter setting for the meters in the policy must coincide with the Enable DEI parameter setting on the physical port. When Enable DEI is selected, the Color Mode for meters must be set to Color Aware. When Enable DEI is not selected, the Color Mode for meters must be set to Color Blind. See 16.24 “To configure Ethernet ports” (p. 599) and 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .

To support H-metering, you must choose an ingress policy with all meter rate modes set to trTCM (RFC 4115).

For 7250 IXR sites, the selected SAP Access Ingress policy must contain an assigned 7250 Ingress CoS policy.

6

Select an egress policy in the Egress Policy panel.

See 50.32 “To configure a 7210 SAP access egress policy” (p. 1558)for more information about 7210 SAP access egress policies.

SAP-based remarking is defined in a 7210 SAP access egress policy. To enable SAP-based remarking on the 7210 SAS-X, you must enable the SAP QoS Marking parameter on the port; see 16.24 “To configure Ethernet ports” (p. 599).

7

To enable table-based color-aware ingress classification, select the Enable Table Classification parameter. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

8

Select an Egress Remarking policy in the Egress Remark Policy panel.

9

Configure the parameters in the Aggregate Rate Limit panel.

You can configure the Ingress Meter parameter only during SAP creation. The parameter must be set to true to support H-metering.

You can configure the Ingress Meter Rate (kbps) and Ingress Meter Burst parameters only after SAP creation.

You can configure the Egress Meter Rate and Egress Meter Burst parameters only when resources are allocated to the SAP Egress Aggregate Meter parameter in the system resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380). To allocate resources on the 7210 SAS-R, configure the Egress SAP Aggregate Meter

parameter in the system resource profile policy assigned to the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).

You must also enable port-based scheduling on 7210 SAS-Mxp and 7210 SAS-R NEs; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384).

The Enable Egress Meter Stats parameter is available when a value is configured for the Egress Meter Rate parameter.

10

For interfaces on 7210 SAS-K sites, configure the required parameters in the Egress Dot1p Remarking panel.

For access ports, configured parameters take effect only when the Remarking parameter is set to true in the 7210 SAP Access Egress policy assigned to the interface; see Step 6.

For L2 Uplink ports, configured parameters take effect only when the Remarking parameter is set to true in the 7210 and 1830 Network policy assigned to the port; see 16.45 “To assign QoS policies to a 7210 SAS Ethernet port” (p. 636).

11

Save the changes and close the forms.

END OF STEPS

77.71 To configure scheduling on an L2 access interface

77.71.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria and click Search. A list of services appears.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand Sites→site→L2 Access Interfaces.

5

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

6

Click on the Schedulers tab and configure the required parameters.

Note: The Schedulers tab is configurable only if a port is assigned to the interface.

7

To configure scheduling on a 7210 SAS site, enable and configure the required parameters in the Egress Aggregate Rate Limit panel and go to Step 30 .

8

To configure scheduling on a 7705 SAR site:

Note: For the 7705 SAR, scheduler behavior is determined by the scheduler mode, which is 4-Priority by default. You can only configure the Egress and Ingress Aggregate Rate Limit parameters when the Scheduler Mode parameter is set to 16-Priority. You can set the Scheduler Mode to 16-Priority only when the port is on an MDA that supports 16-Priority. See the 7705 SAR documentation for more information.

If you change the Scheduler Mode parameter from 16-Priority to 4-Priority, the NFM-P automatically restores the default settings for the Egress Aggregate Rate Limit and Ingress Aggregate Rate Limit panels when you Apply your changes.

1. In the Egress Scheduler panel, configure the Scheduler Mode parameter.
2. In the Ingress Scheduler panel, configure the Scheduler Mode parameter.
3. If you set the Scheduler Mode parameter to 16-Priority in the Egress Scheduler panel or Ingress Scheduler panel, configure the required parameters in the Egress Aggregate Rate Limit panel and Ingress Aggregate Rate Limit panel.
4. Go to Step 30 .

9

To apply an aggregation scheduler policy to the interface:

Note: You cannot specify an access scheduler policy if the port you selected 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) is an HSMDA port. Go to Step 30 .

1. Set the Aggregation parameter to On.
2. Select an aggregation scheduler in the Aggregation Scheduler panel.
3. Go to Step 30 .

10

To specify that an aggregation scheduler policy is not being applied to the interface, set the Aggregation parameter to Off.

Note: The Aggregation parameter is not configurable if the port you selected in 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) is an HSMDA port.

11

Select an ingress scheduler in the Ingress Scheduler panel to choose an ingress scheduler.

12

Configure the required parameters.

The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

The Frame-Based Accounting parameter is not configurable if the port you selected in 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) is an HSMDA port.

You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.

13

Select an ingress policer control policy in the Ingress Policer Control Policy panel.

14

Click Create in the Ingress Policer Control Override panel. The Ingress Policer Control Override (Create) form opens.

15

Configure the required parameters.

16

Click the Level Override Policy Items tab.

17

Select an item from the list and click Properties. The Ingress Policy Policer Level Override form opens.

18

Click the Override tab.

19

Configure the override for the Maximum Cumulative Buffer Space (bytes) parameter.

20

Close the forms.

21

If the port you selected in 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332) is an HSMDA port, go to Step 30 .

22

Select an egress scheduler in the Egress Scheduler panel.

23

Select an egress policer control policy in the Egress Policer Control Policy panel.

24

Click Create in the Egress Policer Control Override panel. The Egress Policer Control Override (Create) form opens.

25

Configure the required parameters.

26

Click the Level Override Policy Items tab.

27

Select an item from the list and click Properties. The Egress Policy Policer Level Override form opens.

28

Click the Override tab.

29

Configure the override for the Maximum Cumulative Buffer Space (bytes) parameter.

30

Save the changes and close the forms.

END OF STEPS

77.72 To configure BPDU Termination, STP, and FIB parameters for the VPLS L2 access interface

77.72.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Forwarding Control tab. Depending on the device being configured, the BPDU Termination tab is displayed. Otherwise, go to Step 7 .

Note: For L2 access interfaces on the 7210 SAS-R, the BPDU Termination tab is displayed only when the Allow L2Pt Xstp BPDU parameter is enabled for the site. See 77.33 “To configure a VPLS site” (p. 2294) .

6

Configure the required parameters.

You can enable the L2 Protocol Termination parameter only when STP is disabled on the site. The Administrative State parameter for STP must be set to Down. See 77.35 “To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site” (p. 2296) .

When the L2 Protocol Termination parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required. You must select at least one protocol.

The Force L2PT on Managed L2 Access Interface parameter is only available for MVPLS L2 access interfaces. When the Force L2PT on Managed L2 Access Interface parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required.

7

Click on the STP tab and configure the required parameters.

8

Click on the FIB tab and configure the required parameters.

The Restrict Protected Source Action parameter is configurable only when the Restrict Protected Source is set to true.

9

If you are creating an MVPLS to run MSTP:

1. Click on the MST Instances tab to edit a SAP MST instance.
2. Select an MST instance and click Properties. The L2 Access Interface MST Instance (Edit) form opens. Configure the required parameters.
3. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

77.73 To assign a DoS protection policy or DDoS protection policy to the VPLS L2 access interface

77.73.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Security tab.

6

Select a DoS protection policy or DDos protection policy in the NE DoS Protection panel.

7

Configure the MAC Monitoring parameter.

8

Configure the Ethernet CFM Monitor Flags parameter.

9

Save the changes and close the forms.

END OF STEPS

77.74 To configure residential subscriber management for the VPLS L2 access interface

77.74.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Subscriber Management tab.

6

Enable the SHCV Enabled parameter to enable SHCV. Otherwise, go to Step 8 .

7

Configure the required parameters.

8

Click on the IGMP Host Tracking tab.

9

Select the import policy used to filter IGMP packets.

10

Configure the required parameters.

11

Click on the Profiles tab and configure the required parameters.

12

Select a default subscriber profile for the interface in the Default Subscriber Profile panel.

13

Select a default SLA profile for the SAP in the Default SLA Profile panel.

14

Select a subscriber identification policy for the SAP in the Subscriber Identification Policy panel.

15

Select a default application profile for the SAP in the Default Application Profile panel.

16

Select a non-subscriber subscriber profile for the SAP in the Non-Subscriber Traffic Subscriber Profile panel.

17

Select a non-subscriber traffic SLA profile for the SAP in the Non-Subscriber Traffic SLA Profile panel.

18

Select a non-subscriber traffic application profile for the SAP in the Non-Subscriber Traffic Application Profile panel.

19

Click on the Host Tracking Info tab to view a list of hosts that are being tracked on this L2 access interface.

20

Save the changes and close the forms.

END OF STEPS

77.75 To configure an Ethernet tunnel on a VPLS L2 access interface

77.75.1 Purpose

If you are creating a fate-sharing Ethernet Tunnel Endpoint SAP—or same-fate SAP—perform the following procedure to configure an Ethernet tunnel.

77.75.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Ethernet Tunnel tab and click Create. The Ethernet Tunnel (Create) form opens.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.76 To configure a redundant VLAN range on a VPLS L2 access interface

77.76.1 Purpose

Configure a redundant VLAN range.

Note: If an MVPLS site has SAPs that manage traffic on the associated VPLS SAPs, you must define a redundant VLAN range during SAP creation. The redundant VLAN range defines the range of VC IDs for VPLS SAPs that the MVPLS manages.

77.76.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Redundancy tab and click Create. The RedundantVlanRange (Create) form opens.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.77 To configure IGMP snooping for a VPLS L2 access interface

77.77.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the IGMP Snooping tab.

6

Configure the required parameters.

The General Query Interval (seconds), Maximum Response Interval (seconds), Robust count, and IGMP Version parameters are configurable when the Send Queries parameter is enabled.

7

To configure a multicast CAC policy, select a policy in the Multicast CAC Policy panel.

8

Configure the required parameters.

9

Click on the Static Mcast Group tab to configure a static multicast group.

10

Click Create. The Access Interface Igmp Snooping Mcast Group Display (Create) form opens.

11

Configure the required parameters.

12

Close the form.

13

Save the changes and close the form.

14

To configure the LAG port down parameters:

1. Click on the LAG Port Down tab.
2. Click Create. The LAG Port Down (Create) form opens.
3. Configure the required parameters.
4. Click on the Levels tab and click Create. The Multicast CAC Level (Create) form opens.
5. Configure the required parameters.
6. Save the changes and close the form.

15

Save the changes and close the forms.

END OF STEPS

77.78 To configure the ARP host for the VPLS L2 access interface

77.78.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the ARP Host Configuration tab.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.79 To configure DHCP for the VPLS L2 access interface

77.79.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click on the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the DHCP tab.

6

Configure the required parameters.

When the Circuit ID or Remote ID parameter is set to HEX, you must configure the accompanying Circuit or Remote HEX String parameter.

When the Remote ID parameter is set to String, you must configure the accompanying Remote ID String parameter.

The Enable Lease Populate parameter is configurable when the Enable parameter is enabled.

Note: For access interfaces on 7210 SAS-E sites, you can only enable the Snooping parameter when the system resource profile for the 7210 SAS-E device is appropriately configured. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

7

To configure the VPLS L2 access interface proxy server:

Note: You can configure a VPLS L2 access interface proxy server on a 7450 ESS or 7750 SR.

1. Click on the Server tab
2. Configure the required parameters.

   The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time RADIUS Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.

8

Save the changes and close the forms.

END OF STEPS

77.80 To configure MVR for a VPLS L2 access interface

77.80.1 Purpose

Depending on the type of device that you are configuring, the MVR tab is configurable. Configure MVR for the SAP.

77.80.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the MVR tab.

6

Select a source MVR VPLS in the Source MVR VPLS panel to associate an MVR VPLS with the SAP.

7

Select a proxy MVR SAP to which the multicast traffic is sent in the Proxy MVR SAP panel.

Note: If the SAP already has an MVR proxy SAP or is the MVR proxy SAP of another SAP, the SAP cannot be an MVR proxy SAP.

8

Save the changes and close the forms.

END OF STEPS

77.81 To configure anti-spoofing filters for a VPLS L2 access interface

77.81.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the Anti-Spoofing tab.

6

Configure the required parameters.

The ARP Reply Agent parameter is configurable only when an IP address is specified for the static hosts on the SAP.

7

Click on the Static Hosts tab to configure a static subscriber host entry for each subscriber host that is not managed by DHCP.

8

Click Create. The Access Interface Anti-Spoofing Static Host Display (Create) form opens.

9

Configure the required parameters.

Note: You must specify at least one IP address or MAC address for each static host. The values that are specified for the Anti-Spoofing and ARP Reply Agent parameters determine the type of address entry that is required for the static host. For example, if you set the Anti-Spoofing parameter to Source Ip Addr, you must specify at least the IP address for the static host.

10

To configure residential subscriber management for the static host:

1. Select a subscriber profile for the static host in the Subscriber Profile panel.
2. Select a subscriber profile for the static host in the Subscriber Profile panel.
3. Select an SLA profile for the static host in the SLA Profile panel.
4. Select an application profile for the static host in the Application Profile panel.

11

Save the changes and close the forms.

END OF STEPS

77.82 To create MIPs and MEPs on a VPLS L2 access interface

77.82.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the navigation tree, expand the Sites icon and click the required site. The site's properties form is displayed.

4

Expand the required site's icon on the navigation tree to show the L2 Access Interfaces.

5

Right-click the required L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

6

Click on the OAM tab, then on the Configuration sub-tab.

7

Configure the required Test Generation Options parameters.

Note: You can propagate test generation role settings to all MEPs on a SAP or service site by clicking Propagate to MEPs.

8

Click on the ETH-CFM sub-tab.

9

To create MIPS:

1. Click Create in the MIP Configurations panel. The MIP Configuration (Create) form opens.
2. Configure the parameters, then click OK. The MIP Configuration (Create) form closes.

   The Primary VLAN Enable, VLAN ID, and MAC Address parameters are configurable only when a port is assigned to the interface.

   Note: To enable Primary VLAN configuration for egress MIPs on supporting 7210 SAS-R NEs, the Bi Directional MIP Egress parameter in the system resource profile policy

assigned to the card must be set appropriately. See 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382) .

Note: You cannot configure MIPs on an SPB-enabled SAP. See 77.2.6 “SPB in VPLS” (p. 2222) in “VPLS management overview” (p. 2210) for more information about SPB.

3. Repeat the previous two steps to create additional MIPs, as required.
4. Click Apply in the VPLS L2 Access Interface (Edit) form. The default MD objects are then automatically created on the service site.
5. Go to the service site's properties form and click the OAM tab and then on the ETH-CFM sub-tab.
6. Scroll down to the Default Domain panel and multi-select all required entries, then click Properties. The CFM Vlan Default Domain (Multiple Instances) form opens.
7. Configure the parameters for the selected entries and click OK. The CFM Vlan Default Domain (Multiple Instances) form closes.
8. Click Apply and then return to the VPLS L2 Access Interface (Edit) form.
9. Scroll down to the MIPs panel and click Resync MIPs. The default Up and Down MIPs that were automatically created by the NE are displayed in the list.

10

To enable a tunnel facility MEP on the L2 access interface, set the Tunnel Fault Notification parameter in the Facility MEPs panel to Accept.

Tunnel Fault Notification is only configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation.

11

If you are configuring a B-L2 Access Interface, go to Step 13 .

12

Configure the Enable Virtual MEP Filter parameter in the Virtual MEPs panel.

13

Configure the parameters in the Squelch Ingress Level panel.

14

Configure the parameters on the LMM Session Stats Collection panel as required.

15

To create a MEP:

1. Click Create in the MEPs panel. The MEP (Create) form opens.
2. Select a MEG in the MEP panel.
3. Configure the general parameters in the MEP panel.

For L2 access interfaces on 7210 SAS-R sites, you can create Up MEPs only when the system resource profile is set appropriately. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

The Type and Interface Type parameters are automatically populated based on whether the MEP is created on a SAP, SDP binding, or Ethernet Tunnel Path Endpoint.

4. Configure the parameters in the CCM panel.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

5. Select a MEG sub-group in the MEG Sub-Grouping panel.

6. Configure the required Test Generation Options parameters.

   The Test Generation Options parameters are only displayed when you set the Direction parameter to Up.

7. If the MD for the MEP has a Maintenance Domain Name Type of none and the associated MEG has a Maintenance Entity group Name Type of icc-based, then the Y.1731 Tests and AIS tabs are configurable. Click on the Y.1731 Tests tab and configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

8. Click on the AIS tab and configure the required parameters.

   The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.

9. Save the changes and close the form.

   Note: If you have configured a SAP, the SAP information is filled in when you configure the MEP.

   Once the site has been created and is operational, you can open the MEP (Edit) form and view an information field that displays the number of Continuously Running Tests currently on the MEP. This appears on the General tab of that form.

16

Save the changes and close the forms.

END OF STEPS

77.83 To assign an ANCP policy to a VPLS L2 access interface

77.83.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the ANCP Static Map tab and click Create. The ANCP Static Map (Create) form opens.

6

Configure the ANCP String parameter.

7

Select an ANCP Policy.

8

Save the changes and close the form.The ANCP Static Map form closes.

9

Save the changes and close the forms.

END OF STEPS

77.84 To configure PIM snooping on a VPLS L2 access interface

77.84.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the PIM Snooping tab and configure the Maximum Number of Groups parameter.

6

Save the changes and close the forms.

END OF STEPS

77.85 To configure MLD snooping for a VPLS L2 access interface

77.85.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the MLD Snooping tab and configure the required parameters.

The General Query Interval (seconds), Maximum Response Interval (seconds), Robust count, and MLD version parameters are configurable when the Send Queries parameter is enabled.

6

To configure one or more static multicast groups:

1. Click on the Static Mcast Group tab.
2. Click Create. The Access Interface Mld Snooping Mcast Group Display (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

7

Save the changes and close the forms.

END OF STEPS

77.86 To configure MVR (MLD) for a VPLS L2 access interface

77.86.1 Purpose

Depending on the type of device that you are configuring, the MVR (MLD) tab is configurable. Use the MVR (MLD) tab to use MLD snooping on the SAP. Configure MVR for the L2 access interface.

77.86.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the MVR (MLD) tab.

6

Select an MVR VPLS with the SAP in the Source MVR VPLS panel.

7

Select a proxy MVR SAP to which the multicast traffic is sent in the Proxy MVR SAP panel.

If the SAP already has an MVR proxy SAP or is the MVR proxy SAP of another SAP, the SAP cannot be an MVR proxy SAP.

8

Save the changes and close the forms.

END OF STEPS

77.87 To create a VPLS or MVPLS B-L2 access interface

77.87.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Site→L2 Access Interfaces.

4

Right-click on L2 Access Interfaces and choose Create B-L2 Access Interface. The B-L2 Access Interface (Create) form opens.

5

Configure the required parameters.

6

Select a transit prefix policy for the interface in the Transit Prefix Policy panel.

7

Select an SHG for the interface in the Split Horizon Group panel.

Note: You must configure an SHG or residential SHG for a VPLS if you plan to create a spoke circuit from this VPLS site to a VLL or another VPLS.

8

Click on the Port tab.

9

Select a port for the B-L2 access interface.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.

10

CAUTION

Service Disruption

The creation of a SAP that uses the same port and encapsulation values as an existing inactive MSAP fails under the following conditions:

- If you try to use the NFM-P to create a SAP, the configuration fails and the NFM-P displays an error message.
- If you use a CLI to create a SAP in a service other than the service that contains the MSAP, the configuration succeeds but the MSAP is inactive until the regular SAP is deleted. Although the NFM-P displays the SAP and MSAP, the MSAP remains inactive and consumes resources.
- If you use a CLI to create a SAP in the service that contains the MSAP, the SAP creation fails.

Nokia recommends that you delete an inactive MSAP from the NFM-P if you need to create a regular SAP on the same port using the same encapsulation values.

Configure the required parameters.

For a B-L2 access interface, only Null, dot1q, and QinQ encapsulation are supported.

When the selected port uses dot1q encapsulation, you can enable the Auto-Assign ID parameter to have the Outer Encapsulation Value parameter automatically assigned. If you choose this, the system assigns the lowest unused encapsulation value.

Note: You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter in the User Preferences form. Choose Application→ User Preferences from the main menu.

11

To configure the EMG:

1. Click on the Egress Multicast Group tab.
2. Select an EMG.

   Note:

   The Egress Multicast Group-L2 Interface form lists only EMGs that have the same egress filter and encapsulation type as the interface.

12

If the selected port uses FR encapsulation, configure Frame Relay for the interface:

1. Click on the Frame Relay tab.
2. Set the FRF-12 Mode parameter to Enabled.
3. Configure the required parameters.

13

To assign ingress and egress QoS policies to the interface:

1. Click on the QoS tab.

   Note:

   Items such as policies, schedulers, and filters can be applied later to multiple service components at once by selecting and right-clicking the components in the service navigation tree, choosing Properties, and configuring the parameters on the appropriate tab. This action opens a properties form in a new window for the component that was right-clicked. The navigation tree is not displayed in this new window.

2. Configure the Ingress Match QinQ Dot1P parameter.

   The Ingress Match QinQ Dot1P parameter is configurable only when the encapsulation type of the port is dot1q or QinQ.

3. Select an ingress in the Ingress Policy panel.

   If you select an access ingress policy that has a forwarding class mapped to an ingress queue group, the port that you choose for the VPLS L2 access interface must use the same access ingress queue group.

   See 16.24 “To configure Ethernet ports” (p. 599) for information about configuring Ethernet ports. See Chapter 49, “Policies overview” for information about queue group template policies.

4. Select an ingress queue group template policy in the Forwarding Plane Redirect panel.
5. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
6. Configure the Egress Mark QinQ Top Bits Only parameter.

   The Egress Mark QinQ Top Bits Only parameters are configurable only when the encapsulation type of the port is dot1q or QinQ.

7. Select an egress policy in the Egress Policy panel.

   Note:

   If you select an access egress policy that has a forwarding class mapped to an egress queue group, the port that you choose for the VPLS L2 access interface must use the same access egress queue group.

   See 16.24 “To configure Ethernet ports” (p. 599) for information about configuring Ethernet ports. See Chapter 49, “Policies overview” for information about queue group template policies.

8. If the port you selected in Step 9 is not an HSMDA port, go to Step 14 .
9. Configure the Packet Byte Offset (bytes) parameter. Before you can configure the parameter, you must enable the associated Override parameter.
10. Select a WRR policy in the Egress HSMDA Override panel.
11. Select an HSMDA egress secondary shaper policy in the Egress HSMDA Override panel.
12. Select an HS secondary shaper in the HS Overrides panel, if required.
13. Configure any required parameters.

14

To configure scheduling:

1. Click on the Schedulers tab.

   Note:

   The Schedulers tab is configurable only when a port is assigned to the interface.

2. To specify that an aggregation scheduler policy is not applied to the interface, set the Aggregation parameter to off and configure the required parameters.

   The Aggregation parameter is not configurable if the port you selected in Step 9 is an HSMDA port.

   The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

   The Frame-Based Accounting parameter is not configurable if the port you selected in Step 9 is an HSMDA port.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.

3. Select and ingress and egress scheduler, and ingress and egress control policy, and an ingress and egress policer control policy.If the port you selected in Step 9 is an HSMDA port, go to Step 15 .

4. To specify that an access scheduler policy is applied to the interface, set the Aggregation parameter to on and configure the Frame-Based Accounting parameter.

   Note:

   You cannot specify an access scheduler policy if the port you selected in Step 9 is an HSMDA port. Go to Step 15 .

5. Select an aggregation scheduler in the Aggregation Scheduler panel.

15

To assign ingress and egress ACL filters to the interface:

Note: Only MAC filters are allowed for B-L2 Access Interfaces.

1. Click on the ACL tab.
2. Select an ingress filter in the Ingress Filter panel.
3. Select an egress filter in the Egress Filter panel.

16

To assign an accounting policy to the interface:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics and Ingress Counter Type parameters.
3. Select an accounting policy.

17

To configure BPDU Termination, STP, FIB, and MRP parameters for the interface:

1. Click on the Forwarding Control tab.
2. Configure the required parameters.

   When the L2 Protocol Termination parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required.

   The Force L2PT on Managed L2 Access Interface parameter is only available for MVPLS B-L2 access interfaces. When the Force L2PT on Managed L2 Access Interface parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required.
3. Click on the STP tab and configure the required parameters.
4. Click on the FIB tab and configure the required parameters.
5. If you are creating an MVPLS to run MSTP, the MST Instances tab is configurable. Click on the MST Instances tab to edit a SAP MST instance.
6. Select an MST instance and click Properties.
7. The B-L2 Access Interface MST Instance (Edit) form opens. Configure the required parameters.
8. Save the changes and close the form.
9. If you are configuring an MVPLS B-L2 access interface, go to Step 18 .
10. Click on the MRP tab and configure the required parameters.

    Note:

    You can view information regarding MMRP Entries by clicking on the MMRP tab.
11. Select an PBB MRP Policy.

18

To assign a DoS protection policy to the interface:

Note: A default DoS protection policy is automatically assigned to the interface.

1. Click on the Security tab.
2. Select an NE DoS protection policy.
3. Configure the MAC Monitoring parameter.

19

To assign test generation options to the interface:

1. Click on the OAM tab, then the Configuration tab.
2. Configure the required Test Generation Options parameters.

20

To configure a redundant VLAN range.:

Note: If an MVPLS site has SAPs that manage traffic on the associated VPLS SAPs, you must define a redundant VLAN range during SAP creation. The redundant VLAN range defines the range of VC IDs for VPLS SAPs that the MVPLS manages.

1. Click on the Redundancy tab.
2. Click Create. The RedundantVlanRange (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

21

To configure anti-spoofing filters for the interface:

1. Click on the Anti-Spoofing tab.
2. Configure the MAC Pinning parameter.

22

To specify QoS policy overrides:

1. Click on the Override Policy Items tab.

   Note: The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.

   - If you configured a non-HSMDA port, then the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.
   - If you configured an HSMDA port, then the Access Ingress Queues, Access Egress HSMDA Queues and Ingress Policer tabs are active.
2. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) for information about setting policy overrides.

23

To associate a MEP with the B-L2 Access interface:

1. Click on the OAM tab, then on the ETH-CFM tab.
2. Perform Step 9 to Step 16 in 77.82 “To create MIPs and MEPs on a VPLS L2 access interface” (p. 2360) .

24

To assign an ANCP policy to the interface:

1. Click on the ANCP Static Map tab. The ANCP Static Map (Create) form opens.
2. Configure the ANCP String parameter.
3. Select an ANCP Policy.
4. Save the changes and close the form.

25

Save the changes and close the forms.

26

Close the Manage Services form.

END OF STEPS

77.88 To create a VPLS I-L2 access interface

77.88.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click on L2 Access Interfaces and choose Create I-L2 Access Interface. The I-L2 Access Interface (Create) form opens.

5

Configure the required parameters.

Tunnel Fault Notification is configurable on interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept.

6

Select a transit prefix policy for the interface in the Transit Prefix Policy panel.

7

Select an SHG for the interface in the Split Horizon Group panel.

Note: You must configure an SHG or residential SHG for a VPLS if you plan to create a spoke circuit from this VPLS site to a VLL or another VPLS.

8

Click on the Port tab.

9

Select a port for the I-L2 access interface.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.
If you select an Ethernet Tunnel Endpoint, the Port form is refreshed and an Ethernet Tunnel tab is added.

10

CAUTION

Service Disruption

The creation of a SAP that uses the same port and encapsulation values as an existing inactive MSAP fails under the following conditions:

- If you try to use the NFM-P to create a SAP, the configuration fails and the NFM-P displays an error message.
- If you use a CLI to create a SAP in a service other than the service that contains the MSAP, the configuration succeeds but the MSAP is inactive until the regular SAP is deleted. Although the NFM-P displays the SAP and MSAP, the MSAP remains inactive and consumes resources.
- If you use a CLI to create a SAP in the service that contains the MSAP, the SAP creation fails.

Nokia recommends that you delete an inactive MSAP from the NFM-P if you need to create a regular SAP on the same port using the same encapsulation values.

Configure the required parameters.

For an I-L2 access interface, only Null, dot1q, and QinQ encapsulations are supported. In the case of I-MPVLS L2-access interfaces, Null encapsulation is not supported. I-MPVLS L2-access interfaces with MSTP configured support QinQ. These ports must also be of Ethernet type PBB.

When the selected port uses dot1q encapsulation, you can enable the Auto-Assign ID parameter to have the Outer Encapsulation Value parameter automatically assigned. If you choose this, the system assigns the lowest unused encapsulation value.

Note: You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter in the User Preferences form. Choose Application→ User Preferences from the main menu.

The Inner Encapsulation Value is configurable only when the port is an Ethernet with QinQ encapsulation.

SAP types X.0 and X.* as well as SAP types *.null and *.* can be configured on the same Q in Q port. You must enable the Enable Q in Q Untagged Sap parameter on the NE which allows the creation of the following default SAP types:

- The SAP type *.null functions as a default SAP for single-tagged frames on a Q in Q port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- The SAP type *.* functions as a default SAP for double-tagged frames on a Q in Q port. This SAP accepts untagged, single-tagged, and double-tagged frames with tags in the range 0 to 4095.

See 12.23 “To enable a Q in Q untagged SAP on an NE” (p. 360) for more information.

If the port you have chosen is an Ethernet Tunnel Endpoint, you will be able to set the Outer Encapsulation Value to 8191. This automatically enables the Ethernet Tunnel Endpoint Control SAP parameter.

11

Configure the Ethernet Tunnel Endpoint Control SAP parameter.

Note: Enabling the Ethernet Tunnel Endpoint Control SAP parameter creates the control L2 Access Interface (also known as a Control SAP). It also automatically sets the value of the Outer Encapsulation Value parameter to 8191.
If you are currently creating a same-fate SAP, the Ethernet Tunnel Endpoint Control SAP parameter must not be enabled.

12

Depending on the port that you have chosen, the Egress Multicast Group tab is configurable. To configure the EMG:

1. Click on the Egress Multicast Group tab.
2. Select an EMG.

   Note:

   The Egress Multicast Group-L2 Interface form lists only EMGs that have the same egress filter and encapsulation type as the interface.

13

If the selected port uses FR encapsulation, configure Frame Relay for the interface:

1. Click on the Frame Relay tab.
2. Set the FRF-12 Mode parameter to Enabled.
3. Configure the required parameters.

14

To assign ingress and egress QoS policies to the interface:

1. Click on the QoS tab.

   Note:

   Items such as policies, schedulers, and filters can be applied later to multiple service components at once by selecting and right-clicking the components in the service navigation tree, choosing Properties, and configuring the parameters on the appropriate tab. This action opens a properties form in a new window for the component that was right-clicked. The navigation tree is not displayed in this new window.

2. Configure the Ingress Match QinQ Dot1P parameter.

   The Ingress Match QinQ Dot1P parameter is configurable only when the encapsulation type of the port is dot1q or QinQ.

3. Select an ingress Qos policy in the Ingress Policy panel.

   Note:

   If you select an access ingress policy which has a forwarding class mapped to an ingress queue group, you must ensure that port that you are configuring for the VPLS L2 access interface has the access ingress queue group with the same name created on it.

   See 16.24 “To configure Ethernet ports” (p. 599) in Chapter 16, “Port and channel object configuration” for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

4. Select an ingress queue group template policy in the Forwarding Plane Redirect panel to choose a policy.
5. Configure the Instance ID parameter.
6. Configure the Egress Mark QinQ Top Bits Only parameter.

   The Egress Mark QinQ Top Bits Only parameter is configurable only when the encapsulation type of the port is dot1q or QinQ.

7. Select an egress QoS policy in the Egress Policy panel.

   Note:

   If you select an access egress policy which has a forwarding class mapped to an egress queue group, you must ensure that port that you are configuring for the VPLS L2 access interface has the access egress queue group with the same name created on it.

   See 16.24 “To configure Ethernet ports” (p. 599) in Chapter 16, “Port and channel object configuration” for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

8. If the port you selected is not an HSMDA port then go to Step 15 .
9. Configure the Packet Byte Offset (bytes) parameter. Before you can configure the parameter, you must enable the associated Override parameter.
10. Select a WRR policy in the Egress HSMDA Override panel.
11. Select an HSMDA egress secondary shaper policy in the Egress HSMDA Override panel.
12. Select an HS secondary shaper in the HS Overrides panel, if required.

13. Configure the required parameters.

15

Click on the Schedulers tab to configure scheduling:

Note: The Schedulers tab is configurable only when a port is assigned to the interface.

16

Perform one of the following.

a. To specify that an aggregation scheduler policy is not applied to the interface.
   1. Set the Aggregation parameter to off.

      Note:

      The Aggregation parameter is not configurable if the port you selected is an HSMDA port.
   2. Configure the required parameters.

      Note:

      The Aggregate Rate Limit (kbps) and Frame-Based Accounting parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

      The Frame-Based Accounting parameter is not configurable if the port you selected is an HSMDA port.

      You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.
   3. Select and ingress scheduler in the Ingress Scheduler panel.
   4. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
   5. If the port you selected is an HSMDA port, go to Step 17 .
   6. Select an egress scheduler in the Egress Scheduler panel.
   7. Select a policer control policy in the Egress Policer Control Policy panel.
   8. Go to Step 17 .

b. To specify that an access scheduler policy is applied to the interface:

   Note: You cannot specify an access scheduler policy if the port you selected is an HSMDA port. Go to Step 17 .
   1. Set the Aggregation parameter to on.
   2. Select an aggregation scheduler in the Aggregation Scheduler panel.

17

To assign ingress and egress ACL filters to the interface:

1. Click on the ACL tab.
2. Select an ingress ACL filter in the Ingress Filter panel.
3. Select an egress ACL filter in the Egress Filter panel.

4. Select an IPv6 ingress ACL filter in the IPv6 Ingress Filter panel.
5. Select an IPv6 egress ACL filter in the IPv6 Egress Filter panel.

18

To assign an accounting policy to the interface:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics and Ingress Counter Type parameters.
3. Select an accounting policy.

19

To configure BPDU Termination, STP, and FIB parameters for the interface:

1. Click on the Forwarding Control tab. Depending on the device being configured, the BPDU Termination tab is displayed. Otherwise, go to 3 .
2. Configure the required parameters.

   When the L2 Protocol Termination parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required.

   The Force L2PT on Managed L2 Access Interface parameter is only available for MVPLS I-L2 access interfaces. When the Force L2PT on Managed L2 Access Interface parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required.
3. Click on the STP tab and configure the required parameters.
4. Click on the FIB tab and configure the required parameters.
5. If you are creating an MVPLS to run MSTP, Click on the MST Instances tab to edit a SAP MST instance. Otherwise, go to Step 20 .
6. Select an MST instance and click Properties.
7. The I-L2 Access Interface MST Instance (Edit) form opens. Configure the required parameters.
8. Save the changes and close the form.

20

To assign a DoS protection policy or DDoS protection policy to the interface:

Note: A default DoS protection policy is automatically assigned to the interface.

1. Click on the Security tab.
2. Select a DoS protection policy in the NE DoS Protection panel.
3. Configure the MAC Monitoring parameter.

21

To assign test generation options to the interface:

1. Click on the OAM tab, then the Configuration tab.

2. Configure the required Test Generation Options parameters.

22

To configure an Ethernet tunnel:

Note: You can only configure Ethernet tunnel SAP path parameters if you are creating a same-fate SAP.

1. Click on the Ethernet Tunnel tab.
2. If you are configuring a fate-sharing Ethernet Tunnel Endpoint SAP (also referred to as same-fate SAP) then go to 3 . Otherwise, go to Step 23 .
3. Click Create. The Ethernet Tunnel (Create) form opens.
4. Configure the required parameters.
5. Save the changes and close the form.

23

To configure a redundant VLAN range:

Note: If an MVPLS site has SAPs that manage traffic on the associated VPLS SAPs, you must define a redundant VLAN range during SAP creation. The redundant VLAN range defines the range of VC IDs for VPLS SAPs that the MVPLS manages.

1. Click on the Redundancy tab.
2. Click Create. The RedundantVlanRange (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

24

To configure IGMP snooping for the interface:

1. Click on the IGMP Snooping tab.
2. Configure the required parameters.

   The General Query Interval (seconds), Maximum Response Interval (seconds), Robust count, and IGMP Version parameters are configurable when the Send Queries parameter is enabled.
3. Click on the Static Mcast Group tab to configure one or more static multicast groups. Otherwise, go to Step 26 .
4. Click Create. The Access Interface Igmp Snooping Mcast Group Display (Create) form opens.
5. Configure the required parameters.
6. Save the changes and close the form.

25

To configure MLD snooping for the interface:

1. Click on the MLD Snooping tab.
2. Configure the parameters in the General tab.
3. Click on the Static Mcast Group tab to configure one or more static multicast groups. Otherwise, go to Step 26 .
4. Click Create. The Access Interface Mld Snooping MCast Group Display (Create) form opens.
5. Configure the required parameters.
6. Save the changes and close the form.

26

To configure anti-spoofing filters for the interface:

1. Click on the Anti-Spoofing tab.
2. Configure the MAC Pinning parameter.

27

To specify QoS policy overrides:

1. Click on the Override Policy Items tab.

   Note: The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.
   - If you configured a non-HSMDA port, then the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.
   - If you configured an HSMDA port, then the Access Ingress Queues, Access Egress HSMDA Queues and Ingress Policer tabs are active.
2. See 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) for information about setting policy overrides.

28

To associate a MEP with the I-L2 Access interface:

1. Click on the OAM tab, then on the ETH-CFM tab.
2. Perform Step 9 to Step 16 in 77.82 “To create MIPs and MEPs on a VPLS L2 access interface” (p. 2360) .

29

To assign an ANCP policy to the interface:

1. Click on the ANCP Static Map tab.
2. Click Create. The ANCP Static Map (Create) form opens.

3. Configure the ANCP String parameter.
4. Select an ANCP Policy.
5. Save the changes and close the form.

30

Save the changes and close the forms.

END OF STEPS

77.89 To configure ETree on a VPLS L2 access interface

77.89.1 Purpose

Use this procedure to configure a SAP as an ETree leaf. If the interface is not configured as a leaf, it will configured as a root by default.

The following prerequisites apply:

- the ETree Enabled procedure must be configured in the VPLS site
- The encapsulation type of the terminating port must be Dot1Q or QinQ

77.89.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Configure the parameters in the Etree panel.

The leaf tag cannot match the dot1q tag on the port. Null and the asterisk (*) are not supported as leaf tags.

5

Save the changes and close the forms.

END OF STEPS

77.90 To configure DHCPv6 snooping for a VPLS or MVPLS L2 access interface

77.90.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→L2 Access Interfaces.

4

Right-click the L2 access interface and choose Properties. The VPLS L2 Access Interface (Edit) form opens.

5

Click on the DHCPv6 tab.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

VPLS SDP binding procedures

77.91 To create a VPLS or MVPLS mesh SDP binding

77.91.1 Prerequisites

The value of the Automatic Mesh SDP Binding Creation parameter determines the way that mesh SDP binding creation occurs in the VPLS.

Note: The maximum combined number of mesh and spoke SDP bindings that you can configure on a service site depends on the NE type and release on which the site is configured. A deployment error occurs when you attempt to create an SDP binding if this limit is reached on the NE.

The maximum number of VPLS sites that can be automatically connected using mesh SDP bindings is 100. You can perform automatic creation of a full mesh in a VPLS in one or more passes, as required.

Nokia recommends that you perform automatic mesh SDP binding creation on only 50 sites at one time. Performing automatic mesh SDP binding creation on 100 sites at one time may result in a substantial delay in building the network.

77.91.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh SDP Bindings.

4

Right-click on Mesh SDP Bindings and choose Create Mesh SDP Binding. The Mesh SDP Binding (Create) form opens.

5

If automatic mesh SDP binding creation is enabled on the VPLS, a dialog box appears. The message in the dialog box discourages manual mesh SDP binding creation when automatic mesh SDP binding creation is specified for a VPLS. Perform one of the following.

a. Choose not to override automatic SDP binding creation. Nokia recommends this action.

   1. Click No.
   2. Click cancel to close the VPLS management form with no action.

b. Choose to override automatic SDP binding creation. Nokia does not recommend this action.
   1. Click Yes. The Mesh SDP Binding (Create) form opens.
   2. Consult an Nokia technical representative before proceeding.

6

Specify a destination NE for the mesh SDP binding:

a. If the destination NE is a managed NE, select an NE from a list of managed NEs.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

7

Configure the required general parameters.

The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

8

To specify a transport tunnel for the mesh SDP binding:

a. Let the NFM-P configure the transport tunnel automatically.
   1. Enable the Auto-Select Transport Tunnel parameter.
   2. Configure either the Profile Name or the Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually by selecting a service tunnel for the mesh SDP binding in the Tunnel panel.

9

Configure the required general parameters.

10

Configure the parameters in the Hash Label panel, if required.

11

Configure the Force VLAN VC Forwarding parameter in the VLAN panel.

The Force VLAN VC Forwarding parameter does not appear if you are creating a Mesh SDP binding for a B-site.

You cannot enable the Force VLAN VC Forwarding parameter if you enable the Force Q-in-Q VC Forwarding parameter in Step 12 .

12

Configure the Force Q-in-Q VC Forwarding parameter in the Q in Q panel.

The Force Q-in-Q VC Forwarding parameter does not appear if you are creating a Mesh SDP binding for a B-site.

You cannot enable the Force Q-in-Q VC Forwarding parameter if you enable the Force VLAN VC Forwarding parameter in Step 11 .

You cannot enable the Force Q-in-Q VC Forwarding parameter if the Enable IP Interface Binding parameter is enabled on the VPLS site in 77.33 “To configure a VPLS site” (p. 2294) .

13

Click on the States tab to specify whether the mesh SDP binding is in service and configure the Administrative State parameter.

14

To specify a transport tunnel for the Return SDP binding:

Note: You can create a return tunnel only between sites that are in the same service. If the sites are not in the same service, the Return tab does not appear.

a. Let the NFM-P configure the transport tunnel automatically.
   1. Click on the Return tab.
   2. Enable the Auto Select Return Transport Tunnel parameter.
   3. Configure either the Profile Name or the Return Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually.
   1. Click on the Return tab.
   2. Select a return tunnel in the Return Tunnel panel.

15

If you are configuring a 7705 SAR, go to Step 17 .

16

To configure QoS:

1. Click on the QoS tab.
2. Select a network policy in the Forwarding Plane Redirect panel.
3. Select an ingress queue group template policy in the Forwarding Plane Redirect panel.
4. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
5. Select a network policy in the Port Redirect panel.
6. Select an egress queue group template policy in the Port Redirect panel.
7. Configure the Instance ID parameter in the Port Redirect panel.

17

Click on the Pseudowire OAM tab and configure the Control Word parameter.

18

To assign ingress and egress ACL filters to the mesh SDP binding:

1. Click on the ACL tab.
2. Select an ingress filter in the Ingress Filter panel.
3. Select an egress filter in the Egress Filter panel.
4. Select an IPv6 ingress ACL filter in the IPv6 Ingress Filter panel.
5. Select an IPv6 egress ACL filter in the IPv6 Egress Filter panel.

19

To configure anti-spoofing for the mesh SDP binding:

1. Click on the Anti-Spoofing tab.
2. Configure the MAC Pinning parameter.

20

To assign an accounting policy to the mesh SDP binding:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy.

21

Save the changes and close the forms.

22

To assign a DoS protection policy to the mesh SDP binding, perform 77.94 “To assign a DoS protection policy to a VPLS SDP binding” (p. 2393) .

23

To add or delete MFIB allowed daughter cards or configure MRP on the mesh SDP binding, perform 77.105 “To configure learning protection parameters on a VPLS SDP binding” (p. 2403) .

24

To configure DHCP on the mesh SDP binding, perform 77.95 “To configure DHCP for the VPLS SDP binding” (p. 2394) .

25

To configure IGMP snooping on the mesh SDP binding, perform 77.96 “To configure IGMP snooping for the VPLS SDP binding” (p. 2395) .

26

To configure DHCPv6 snooping on the mesh SDP binding, perform 77.108 “To configure DHCPv6 snooping for the VPLS or MVPLS SDP binding” (p. 2407) .

27

To associate a MEP to the mesh SDP binding, perform 77.98 “To create a MIP on a VPLS SDP binding” (p. 2396) .

28

To configure MLD snooping on the mesh SDP binding, perform 77.100 “To configure MLD Snooping for the VPLS SDP binding” (p. 2399) .

END OF STEPS

77.92 To create a VPLS or MVPLS spoke SDP binding

Note: The maximum combined number of mesh and spoke SDP bindings that you can configure on a service site depends on the NE type and release on which the site is configured. A deployment error occurs when you attempt to create an SDP binding if this limit is reached on the NE.

You cannot create a spoke SDP binding on an MVPLS site that runs MSTP. Likewise, you cannot enable MSTP for a site that has a spoke SDP binding, or on a SAP with a non-zero encapsulation value.

For services employing BGP AD and BGP VPLS, you should create SDP bindings manually at non-BGP AD or BGP VPLS enabled sites, or to other BGP AD or BGP VPLS sites where auto-created pseudowires are not expected to be created.

77.92.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

If you are creating a redundant spoke SDP binding under an endpoint, go to Step 4 . Otherwise go to Step 6 .

Note: Redundant spoke SDP bindings under an endpoint are only valid for VPLS regular sites and B-Sites. They do not apply to I-Sites.

4

On the service tree, expand Sites→site→Spoke SDP Bindings.

5

Right-click on Spoke SDP Bindings below the endpoint and choose Create Spoke SDP Bindings. The Spoke SDP Binding (Create) form opens. Go to Step 8 .

6

On the service tree, expand the site on which you want to create a VPLS spoke SDP binding.

7

Right-click on Spoke SDP Bindings and choose Create Spoke SDP Bindings. The Spoke SDP Binding (Create) form opens.

8

Specify a destination NE for the spoke SDP binding:

a. If the destination NE is a managed NE, select an NE from a list of managed NEs.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

9

Configure the required parameters.

The Block on Mesh SDP Failure parameter does not appear if you are creating a Spoke SDP binding for a B-site or an I-site.

The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

10

Perform one of the following to specify a transport tunnel for the spoke SDP binding.

a. Let the NFM-P configure the transport tunnel automatically.

1. Enable the Auto-Select Transport Tunnel parameter.
2. Configure either the Profile Name or the Tunnel Auto-Selection Transport Preference parameter.

Note: If you are creating a Spoke SDP binding for a B-site, the Tunnel Auto-Selection Preference must be either MPLS:LDP or MPLS:RSVP. For I-sites, you can use MPLS:LDP, MPLS:RSVP, GRE, or Any.

b. Configure the transport tunnel manually by selecting a tunnel in the Tunnel panel.

c. Configure an MPLS-TP transport tunnel manually by selecting a tunnel in the Tunnel panel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for more information about how to create an MPLS-TP service tunnel.

11

If you are creating the spoke SDP binding for a B-Site or an I-Site, go to Step 19 .

12

If you are creating the spoke SDP binding under an endpoint, go to Step 14 .

13

Select an endpoint in the Redundancy panel.

14

Configure the required general parameters.

15

Configure the parameters in the Hash Label panel, if required.

The Enable Hash Label parameter is not configurable if you selected an MPLS-TP tunnel in Step 10 .

16

Configure the Force VLAN VC Forwarding parameter in the VLAN panel.

The Force VLAN VC Forwarding parameter does not appear if you are creating a Spoke SDP binding for a B-Site or an I-Site.

You cannot enable the Force VLAN VC Forwarding parameter if you enable the Force Q-in-Q VC Forwarding parameter in Step 17 .

17

Configure the Force Q-in-Q VC Forwarding parameter in the Q in Q panel.

The Force Q-in-Q VC Forwarding parameter does not appear if you are creating a Spoke SDP binding for a B-site or an I-Site.

You cannot enable the Force Q-in-Q VC Forwarding parameter if you enable the Force VLAN VC Forwarding parameter in Step 16 .

You cannot enable the Force Q-in-Q VC Forwarding parameter if the Enable IP Interface Binding parameter is enabled on the VPLS site in 77.33 “To configure a VPLS site” (p. 2294) .

18

Configure the Enable PW Status Signaling parameter in the Pseudowire Signaling panel.

The Enable PW Status Signaling parameter is not configurable if you selected an MPLS-TP tunnel in Step 10 .

19

Select an SHG for the spoke SDP binding in the Split Horizon Group panel.

Note: You must configure an SHG or residential SHG on a spoke SDP binding for an HVPLS that includes another VPLS or a VLL service.

20

Select an application profile for the spoke SDP binding.

Only local profiles on the NE can be selected.

21

Select a transit prefix policy to associate with the spoke SDP binding in the Transit Prefix Policy panel.

Note: You can bind a transit policy to only one access interface or spoke SDP binding per NE.
The transit policy and the application profile must belong to the same application assurance group or partition.

22

Select an operational group for the spoke SDP Binding to join as a member.

23

Select an operational group for the spoke SDP Binding to monitor.

24

Perform one of the following to specify a transport tunnel for the Return SDP binding.

Note: You can create a return tunnel only between sites that are in the same service. If the sites are not in the same service, the Return tab does not appear.
If you are creating a Return SDP binding for a B-site, the Return Tunnel Auto-Selection Preference must be either MPLS:LDP or MPLS:RSVP. For I-sites, you can use MPLS:LDP, MPLS:RSVP, GRE, or Any.

a. Let the NFM-P configure the transport tunnel automatically.
   1. Click on the Return tab.
   2. Enable the Auto Select Return Transport Tunnel parameter.
   3. Configure either the Profile Name or the Return Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually.
   1. Click on the Return tab.
   2. Select a return tunnel in the Return Tunnel panel.

25

To configure QoS:

1. Click on the QoS tab.
2. Select a network policy in the Forwarding Plane Redirect panel.
3. Select an ingress queue group template policy in the Forwarding Plane Redirect panel.

4. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
5. Select a network policy in the Port Redirect panel.
6. Select an egress queue group template policy in the Port Redirect panel.
7. Configure the Instance ID parameter in the Port Redirect panel.

26

Click on the States tab to specify whether the spoke SDP binding is in or out of service and configure the Administrative State parameter.

27

Click on the Pseudowire OAM tab and configure the Control Word parameter.

Note: If you are creating a spoke SDP binding using an MPLS-TP service tunnel for pseudowire static configuration, you must set the Control Word parameter to Preferred.

28

If you are configuring a 7210 SAS-M, 7210 SAS-Mxp, 7210 SAS-R, 7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T, or 7210 SAS-X, go to Step 31 .

29

To assign ingress and egress ACL filters to the spoke SDP binding:

1. Click on the ACL tab.
2. Select an ingress filter in the Ingress Filter panel.
3. Select an egress filter in the Egress Filter panel to choose an ingress ACL filter.
4. Select an IPv6 ingress filter in the IPv6 Ingress Filter panel.
5. Select an IPv6 egress filter in the IPv6 Egress Filter panel.

30

If you are configuring a 7705 SAR, go to Step 32 .

31

To assign an accounting policy to the interface:

1. Click on the Accounting tab.
2. Select an accounting policy.
3. Configure the required parameters.

32

To configure anti-spoofing for the spoke SDP binding.

1. Click on the Anti-Spoofing tab.

2. Configure the MAC Pinning parameter.

33

If you are configuring a 7210 SAS-M, 7210 SAS-R, 7210 SAS-S, 7210 SAS-Sx, 7210 SAS-T, or 7210 SAS-X, go to Step 34 .

34

Save the changes and close the forms.

35

To assign a DoS protection policy to the spoke SDP binding, perform 77.94 “To assign a DoS protection policy to a VPLS SDP binding” (p. 2393) .

36

To configure MFIB, STP, FIB, MRP, and MAC learning parameters on the spoke SDP binding, perform 77.105 “To configure learning protection parameters on a VPLS SDP binding” (p. 2403) .

37

To configure DHCP on the spoke SDP binding, perform 77.95 “To configure DHCP for the VPLS SDP binding” (p. 2394) .

38

To configure IGMP snooping on the spoke SDP binding, perform 77.96 “To configure IGMP snooping for the VPLS SDP binding” (p. 2395) .

39

To configure DHCPv6 snooping on the spoke SDP binding, perform 77.108 “To configure DHCPv6 snooping for the VPLS or MVPLS SDP binding” (p. 2407) .

40

To associate a MEP to the spoke SDP binding, perform 77.98 “To create a MIP on a VPLS SDP binding” (p. 2396) .

41

To configure MLD snooping on the spoke SDP binding, perform 77.100 “To configure MLD Snooping for the VPLS SDP binding” (p. 2399) .

42

To configure PIM snooping for the spoke SDP binding, perform 77.104 “To configure PIM snooping for a VPLS spoke SDP binding” (p. 2403) .

43

To configure custom object attributes for AA reporting, perform 77.106 “To configure custom object attributes for AA reporting for a spoke SDP binding” (p. 2405) .

END OF STEPS

77.93 To configure an MPLS-TP static pseudowire on a VPLS spoke SDP binding

77.93.1 Purpose

Create an MPLS-TP static pseudowire on the spoke SDP binding. An MPLS-TP service tunnel must be used in the SDP binding, and the Control Word parameter for pseudowire OAM must be set to Preferred. See 77.92 “To create a VPLS or MVPLS spoke SDP binding” (p. 2386) .

77.93.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Spoke SDP Bindings.

4

Click on the SDP binding on which you want to configure the MPLS-TP static pseudowire. The Spoke SDP Binding (Edit) form opens.

5

Click on the Control Channel tab and configure the required parameters.

6

Click on the Static PW tab and click Create. The PW Path ID (Create) form opens.

7

Configure the Path AGI parameter.

8

Configure the parameters in the Source Attachment Individual Identifier panel.

9

Configure the parameters in the Target Attachment Individual Identifier panel.

10

Save the changes and close the forms.

END OF STEPS

77.94 To assign a DoS protection policy to a VPLS SDP binding

Note: A default DoS protection policy is automatically assigned.

77.94.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the SDP binding you want to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the Security tab.

6

Select and NE DoS protection policy.

7

Configure the MAC Monitoring parameter.

8

Configure the Ethernet CFM Monitor Flags parameter.

9

Save the changes and close the forms.

END OF STEPS

77.95 To configure DHCP for the VPLS SDP binding

77.95.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the DHCP tab.

Note: The DHCP tab does not appear if you are configuring a mesh or spoke SDP binding for a B-site.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

77.96 To configure IGMP snooping for the VPLS SDP binding

77.96.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria and click Search. A list of services appears.

3

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

5

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

6

Click on the IGMP Snooping tab.

7

Configure the required parameters.

The General Query Interval (seconds), Maximum Response Interval (seconds), Robust count, and IGMP Version parameters are configurable when the Send Queries parameter is enabled.

8

To configure a multicast CAC policy, select a policy in the Multicast CAC Policy panel.

9

Configure the required parameters.

10

Click on the Static Mcast Group tab to configure a static multicast group.

11

Click Create. The Access Interface Igmp Snooping Mcast Group Display (Create) form opens.

12

Configure the required parameters.

13

Save the changes and close the forms.

END OF STEPS

77.97 To configure ETree on a VPLS SDP binding

77.97.1 Purpose

Use this procedure to configure a spoke SDP as an ETree leaf. If the interface is not configured as a leaf, it will configured as a root by default.

The VC Type parameter must be set to VLAN.

77.97.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Spoke SDP Bindings.

4

Configure the parameters in the Etree panel.

5

Save the changes and close the forms.

END OF STEPS

77.98 To create a MIP on a VPLS SDP binding

77.98.1 Purpose

This procedure is applicable to both VPLS spoke and mesh SDP bindings. You can also create an SDP binding MIP from the service topology view.

Note: You cannot configure MIPs on SPB-enabled SDP bindings. See 77.2.6 “SPB in VPLS” (p. 2222) in “VPLS management overview” (p. 2210) for more information about SPB.

77.98.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the OAM tab, then the Tests sub-tab, then the CFM Global MEG sub-tab.

4

Select a Global MEG or refer to 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) to create one.

Note: A Global MEG suitable for MIP creation on an SDP binding has a couple of specific requirements. When creating a Global MEG in 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) , on the General tab you must set the Initial MHF-Creation parameter to default. On the Service you create there, the MIP(s) Creation on SDP Bindings parameter in the Auto-Creation of MIP(s) panel must also be enabled.

5

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

6

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

7

Click on the OAM tab, then the ETH-CFM sub-tab.

8

To configure a MIP, enable the MIP parameter and configure the MIP MAC Address parameter in the MIP Configuration panel.

Note: If the MIP is later disabled and subsequently re-enabled, the MIP MAC Address will return to a default value on the NE.

9

To view MIPs on the SDP binding, click Resync MIPs and then Search in the MIPs panel.

10

Save the changes and close the forms.

END OF STEPS

77.99 To create a MEP on a VPLS SDP binding

77.99.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the OAM tab, then the ETH-CFM tab.

6

Configure the Enable Virtual MEP Filter parameter in the Virtual MEPs panel.

7

Configure the parameters in the Squelch Ingress Level panel.

Levels must be assigned contiguously from Level 0. If you select a level greater than 0, then all levels lower than the one you chose will automatically be selected.

The Squelch Ingress Levels configured for all the mesh or spoke bindings on a site can be compared in one view when you click on the Mesh SDP Bindings or Spoke SDP Bindings item below a site on the navigation tree. Use the scroll bar to locate the Squelch Ingress Level columns.

8

Configure the parameters on the LMM Session Stats Collection panel as required.

9

To create a MEP:

1. Click Create in the MEPs panel. The MEP (Create) form opens.
2. Select a MEG in the MEP panel.
3. Configure the general parameters in the MEP panel.

   For L2 access interfaces on 7210 SAS-R sites, you can create Up MEPs only when the system resource profile is set appropriately. See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) .

   The Type and Interface Type parameters are automatically populated based on whether the MEP is created on a SAP, SDP binding, or Ethernet Tunnel Path Endpoint.
4. Configure the parameters in the CCM panel.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.
5. Select a MEG sub-group in the MEG Sub-Grouping panel.
6. If the MD for the MEP has a Maintenance Domain Name Type of none and the associated MEG has a Maintenance Entity group Name Type of icc-based, then the Y.1731 Tests and AIS tabs are configurable. Click on the Y.1731 Tests tab and configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.
7. Click on the AIS tab and configure the required parameters.

   The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.

10

To add a MIP to the SDP binding, perform 77.98 “To create a MIP on a VPLS SDP binding” (p. 2396) .

11

Save the changes and close the forms.

END OF STEPS

77.100 To configure MLD Snooping for the VPLS SDP binding

77.100.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the SDP binding you want to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the MLD Snooping tab.

Note: The MLD Snooping tab does not appear if you are creating a Mesh SDP binding for a B-site or I-site.

6

Configure the required parameters.

The General Query Interval (seconds), Maximum Response Interval (seconds), Robust count, and MLD version parameters are configurable when the Send Queries parameter is enabled.

7

To configure one or more static multicast groups:

1. Click on the Static Mcast Group tab.
2. Click Create. The Circuit Mld Snooping Mcast Group Display (Create) form opens.
3. Configure the required parameters.

8

Save the changes and close the forms.

END OF STEPS

77.101 To configure BFD on a VPLS SDP binding

77.101.1 Purpose

BFD is used over the VCCV control channel for PW fault detection. BFD carried over a PW associated channel enables the monitoring of the PW between the terminating PEs, regardless of whether the service spans multiple hops. This allows faults that are local to individual PWs to be detected.

77.101.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPLS service and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings or Mesh SDP Bindings and click on the SDP binding on which you want to configure BFD.

4

Click on the BFD tab and configure BFD on the spoke SDP binding:

1. On the General tab, enable the Enable BFD parameter.
2. Choose a BFD template. To create a BFD template, see 28.25 “To configure a BFD template policy” (p. 911) .

   Note:

   You must assign a BFD template to the SDP binding if you enable BFD.
3. In the Failure Action dropdown, select the action to be performed on BFD failure.
4. Configure the Up-Timer parameter.

   Note:

   The Up-Timer parameter is applicable only when the value of Failure Action is set to down.

5

Save and close the forms.

END OF STEPS

77.102 To clear BFD sessions and statistics on a VPLS SDP binding

77.102.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPLS service and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to view the BFD session information.

4

Clear BFD sessions or statistics on the spoke SDP binding:

1. Click on the BFD tab, then on the BFD Session tab.
2. Click Clear All to clear all BFD sessions.
3. Click Clear All Statistics to clear all BFD statistics.

5

Close the forms.

END OF STEPS

77.103 To view the BFD session status on a VPLS SDP binding

77.103.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPLS service and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and view the status of the BFD session on the spoke SDP binding:

1. Click on the BFD Session tab.
2. Choose a BFD session and click Properties. The BFD Session (View) form opens.

5

Close the forms.

END OF STEPS

77.104 To configure PIM snooping for a VPLS spoke SDP binding

77.104.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure PIM snooping.

4

Click on the PIM Snooping tab.

5

Configure the Max Number of Groups parameter.

6

Save the changes and close the forms.

END OF STEPS

77.105 To configure learning protection parameters on a VPLS SDP binding

77.105.1 Purpose

This procedure applies to MFIB, STP, FIB, MRP, and MAC learning protection parameters on a VPLS SDP binding.

Note: The single-slot models of the 7450 ESS and 7750 SR support the addition or deletion of MFIB allowed daughter cards.

The MFIB Allowed Daughter Card tab does not appear if you are creating a Mesh SDP binding for a B-site.

77.105.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the Forwarding Control tab.

6

Perform the following if you are configuring a spoke SDP binding. Otherwise, go to Step 7 .

1. Depending on the device being configured, the BPDU Termination tab is displayed.

   Note:

   For spoke SDPs on the 7210 SAS-R, the BPDU Termination tab is displayed only when the Allow L2Pt Xstp BPDU parameter is enabled for the site. See 77.33 “To configure a VPLS site” (p. 2294) .
2. Configure the required parameters.

   You can enable the L2 Protocol Termination parameter only when STP is disabled for the site. The Administrative State parameter for STP must be set to Down; see 77.35 “To configure MFIB, STP, FIB, and MAC learning protection for a VPLS site” (p. 2296) .

   When the L2 Protocol Termination parameter is set to Enabled, a list of L2PT protocols is displayed. Select all that are required. You must select at least one protocol.
3. Click on the STP tab and configure the required parameters.
4. Click on the FIB tab and configure the required parameters.

   The Maximum Entries parameter does not appear if you are creating a Spoke SDP binding for a B-site.

   The Restrict Protected Source Action parameter is configurable when the Restrict Protected Source parameter is set to true.

7

If you are configuring a mesh SDP binding, click on the FIB tab and configure the parameters. Otherwise, go to Step 8 .

8

Click on the MFIB Allowed Daughter Card tab and click Create. The SDP Binding MFIB Allowed Daughter Card (Create) form opens.

9

Select a daughter card, save the changes and close the form.

10

Click on the MRP tab if you are configuring a mesh or spoke SDP binding for a B-site and configure the parameters. Otherwise, go to Step 12 .

Note: The MRP tab does not appear if you are creating an SDP binding for an I-site.

11

Select an PBB MRP Policy.

12

Save the changes and close the forms.

END OF STEPS

77.106 To configure custom object attributes for AA reporting for a spoke SDP binding

77.106.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure custom object attributes for AA reporting.

4

On the General tab, choose an Application Profile for the applicable Group and Partition.

5

Click on the Application Assurance tab, then the NSP Analytics Parameters sub-tab.

6

Click on the Reporting tab and click Create. The AA Reporting (Create) form opens.

7

Configure the required parameters.

8

Save the changes and close the forms.

END OF STEPS

77.107 To force a switchover to a redundant spoke SDP binding

77.107.1 Prerequisites

This procedure can only be performed on a VPLS that has been configured with endpoints that are associated redundant spoke SDP bindings.

77.107.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand the site on which you want to perform the switchover.

4

Right-click on the redundant spoke SDP binding under an endpoint for the site for which you want to perform the switchover and choose Force Switchover.

5

You can clear the switchover at a later time by performing Step 1 to Step 4 again and selecting the Clear Forced Switchover.

Note: You must clear a manually forced switchover by using the Clear Forced Switchover button when the active spoke SDP binding is restored. The NFM-P cannot automatically switch to another active spoke SDP binding if this is not done.

6

Close the forms.

END OF STEPS

77.108 To configure DHCPv6 snooping for the VPLS or MVPLS SDP binding

77.108.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites→site→Mesh|Spoke SDP Bindings.

4

Right-click on the mesh or spoke SDP binding you need to configure and choose Properties. The [Mesh|Spoke] SDP Binding (Edit) form opens.

5

Click on the DHCPv6 tab.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

BGP AD and BGP VPLS procedures

77.109 To configure the VPLS for BGP auto-discovery

77.109.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the BGP AD tab.

4

Set the BGP AD Administrative Status parameter to Up. The BGP AD Service Identification group appears.

5

Configure the VPLS ID parameter. This parameter must be a unique network-wide ID.

6

Save the changes and close the forms.

END OF STEPS

77.110 To configure a site for BGP AD or BGP VPLS

77.110.1 Purpose

This procedure provides the steps required to enable BGP Auto Discovery or configure BGP VPLS on a VPLS site.

- BGP AD enables a VPLS PE router to discover other PE routers that are part of the same VPLS domain. T-LDP based label signaling is used for the pseudowire.
- BGP VPLS provides the mechanism for service member auto-discovery based on Route Target. MP-BGP based label signaling is used for the pseudowire.

The following considerations apply:

- BGP AD and BGP VPLS implementations only apply to regular VPLS sites and B-Sites, but not to I-Sites. For BGP VPLS, the B-Site cannot be used as a backbone for an I-Site or Epipe.

- For 7210 SAS NEs that support BGP AD, only VPLS service sites are supported (not MVPLS or B-sites).
- Up to two BGP instances can be created per VPLS site.

77.110.2 Steps

1

Prior to configuring a site for BGP AD or BGP VPLS, you must complete the following actions:

a. Create a routing policy statement to define the required community members. See 54.5 “To configure a routing policy statement” (p. 1745) . This defines the VSI Import/Export Routing Targets.

b. Enable BGP on the routing instance of each NE in the VPLS or BGP VPLS. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.

c. Configure global-level BGP on each NE in the VPLS or BGP VPLS. See 28.31 “To configure global-level BGP” (p. 918) for more information.

The following items are required for BGP AD implementation:

- You must enable the L2 VPN parameter in the Family panel on one or more of the following:
  - the VPN tab of the BGP site form
  - the VPN tab of the BGP Peer Group form, which contains the peers that are involved in the BGP multi-homing
  - the VPN tab of the BGP peer form, which participates in BGP signaling
- Create a peer group under BGP. This peer group is used to collectively define the peers involved in the VPLS.
- Create the required peers under the peer group. These peers are the NEs involved in the VPLS.

d. Create a PW template policy. See Chapter 83, “Service PW template policies” for information about PW template policies.

e. Distribute the PW template policy to each NE that is or will be a component of the VPLS or BGP VPLS.

f. Ensure that the BGP AD Administrative Status parameter is set to Up and that the VPLS ID parameter is configured.

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand Sites and click on the site on which you want to configure BGP AD or BGP VPLS.

5

Click on the BGP tab.

6

Enable the Enable BGP parameter. The BGP panel is displayed.

7

Click Create. The BGP Configuration form opens.

8

Configure the Route Distinguisher parameter.

9

Click on the VSI Import Policies tab.

10

Select up to five import policies. Alternatively, you can enter the policy names manually.

11

Perform one of the following:

a. Enter the Import Route Target name manually in the provided field. The format be “target:x:y”.

b. Select an import route target. The community members you defined when creating a routing policy in Step 1 can be selected.

12

Click on the VSI Export Policies tab and select up to five export policies. Alternatively, you can enter the policy names manually.

13

Perform one of the following:

a. Enter the Export Route Target name manually in the provided field. The format must be “target:x:y”.

b. Select an export route target. The community members you defined when creating a routing policy in Step 1 can be selected.

14

Click on the PW Template Binding tab and perform one of the following:

a. Select an existing PW Template Binding.

b. Add a new PW Template Binding. Go to Step 15 .

15

Create one or more PW Template Bindings:

1. Click Create. The PW Template Binding (Create) form opens.
2. Enter a Split Horizon Group name.
3. Select the PW Template you created in Step 1 .
4. Select an operational group.
5. Click on the BFD tab and configure the required BFD parameters.
6. Click on the PW Templates Binding Route Target tab and click Create. The PW Template Binding Route Target (Create) form opens.

   Note:

   This Import Route Target is used by the NE to decide which PW Template to use to create SDP bindings. If a far-end neighbor has a matching export target (that is, to the PW Template Import Target being defined here), then this PW Template is selected by the NE to create the pseudowire that is used to link both sites of the VPLS. If nothing is entered, and multiple PW Templates are defined, the first one found by the NE is used (most likely the one with the lowest PW Template Policy ID).

   Enter the required Import Route Target in the field and click OK. The PW Template Binding Route Target (Create) form closes and the entered Import Route Target is displayed in the table on the PW Template Binding Route Target tab.
7. Click OK to accept the selection. The PW Template Binding (Create) form closes and the new PW Template Binding is displayed in the table on the PW Template Binding tab.
8. Select the required entry or entries from the list.

16

Save the changes and close the form.

17

If you are configuring a BGP VPLS, go to Step 18 . Otherwise go to Step 21 .

18

Enable the Enable BGP VPLS parameter.

19

Configure the required BGP parameters.

20

Select an operational group for the BGP site to monitor.

21

To configure BGP AD, enable the Enable BGP AD parameter and configure the required BGP AD parameters.

Note: The Global Service VPLS ID is set to the value defined at the service level.
If VPLS ID is defined at the service level, then the NFM-P ensures that each site has the same VPLS ID. If a site has a different VPLS ID, an alarm is raised and the ID mismatch is indicated in the Status panel of the VPLS site properties form.
The same VPLS ID value is propagated to each site in a VPLS. If you change the VPLS ID of a site without using the NFM-P, the NFM-P displays a warning form.

22

Save and close the forms.

END OF STEPS

77.111 To configure a site for BGP VPLS Multi-homing

77.111.1 Purpose

This procedure provides the steps required to configure a site for BGP VPLS Multi-homing. BGP VPLS Multi-homing provides redundancy support through the configuration of a number of multi-homed sites, rather than through the use of MC-LAG or MC-Ring as access mechanisms. Dual-homing between a CE device and a pair of VPLS PE devices (potentially in different autonomous systems) is an example of such a configuration.

Note: VPLS sites, I-sites, and B-Sites can be configured for BGP VPLS Multi-homing. However, the B-Sites cannot be used as a backbone for an Epipe.

MVPLS services cannot be configured for this application.

An RD or RT configured under the BGP of a VPLS site cannot be removed as long as there is a multi-homing site ID configured whose administration state is up.

You can see a list of all current BGP VPLS Multi-homing sites in a multi-homing VPLS service by viewing the BGP Multi-homing Sites tab on the service configuration form.

77.111.2 Steps

1

Prior to configuring a site for BGP VPLS Multi-homing, you must complete the following actions:

a. Create a routing policy statement to define the required community members. See 54.5 “To configure a routing policy statement” (p. 1745) for more information. This defines the VSI Import/Export Routing Targets.

b. Enable BGP on the routing instance of each NE in the VPLS or BGP VPLS. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.

c. Configure global-level BGP on each NE in the VPLS or BGP VPLS. See 28.31 “To configure global-level BGP” (p. 918) for more information.

   The following items are required:

   - Enable the L2 VPN parameter on one or more of the following:
     - the VPN tab of the BGP site form
     - the VPN tab of the BGP Peer Group form, which contains the peers that are involved in the BGP multi-homing
     - the VPN tab of the BGP peer form, which participates in BGP signaling
   - Create a peer group under BGP. This peer group is used to collectively define the peers involved in the VPLS.
   - Create the required peers under the peer group. These peers are the NEs involved in the VPLS.**Note:** For optimal processing while a BGP multi-homing site is activated or de-activated, or the system is rebooted, you should also:
   - Enable the L2 VPN parameter in the Rapid Update Address Family panel on the BGP site VPN tab.
   - Enable the Enable Rapid Withdrawal parameter on the BGP site Behavior tab.

d. Create a PW template. See Chapter 83, “Service PW template policies” for information about PW templates.

e. Distribute the PW Template to each NE that is or will be a component of the VPLS or BGP VPLS.

f. Create SDPs for the BGP VPLS Multi-homing site(s), if manually-provisioned service tunnels are required. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190)

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose the required VPLS or HVPLS and click Properties. The VPLS Service (Edit) form opens.

4

On the service tree, expand Sites and click on the site on which you want to configure BGP VPLS multi-homing.

5

Click on the BGP tab.

6

Enable the Enable BGP parameter.

7

Configure the Route Distinguisher parameter.

8

Click Configuration to configure the Route Targets and the PW Templates. The BGP Configuration form opens.

9

Select up to five import policies. Alternatively, you can enter the policy names manually.

10

Perform one of the following:

a. Enter the Import Route Target name manually. The format must be “target:x:y”.

b. Select an import route target. The community members you defined when creating a routing policy in Step 1 can be selected.

11

Click on the VSI Export Policies tab and select up to five export policies. Alternatively, you can enter the policy names manually.

12

Perform one of the following:

a. Enter the Export Route Target name manually. The format must be “target:x:y”.

b. Select an export route target. The community members you defined when creating a routing policy in Step 1 can be selected.

13

Click on the PW Template Binding tab and perform one of the following:

a. Select an existing PW Template Binding.

b. Create a new PW Template Binding.

1. Click Create. The PW Template Binding (Create) form opens.
2. Enter a Split Horizon Group name.
3. Select the PW Template you created in Step 1 .
4. Click on the PW Templates Binding Route Target tab and click Create. The PW Template Binding Route Target (Create) form is displayed.

   This Import Route Target is used by the NE to decide which PW Template to use to create SDP bindings. If a far-end neighbor has a matching export target (that is, to the PW Template Import Target being defined here), then this PW Template is selected by the NE to create the pseudowire that is used to link both sites of the VPLS. If nothing is entered, and multiple PW Templates are defined, the first one found by the NE is used (most likely the one with the lowest PW Template Policy ID).

Enter the required Import Route Target in the field and click OK. The PW Template Binding Route Target (Create) form closes and the entered Import Route Target is displayed in the table on the PW Template Binding Route Target tab.

5. Click OK to accept the selection.
6. Select the required entry or entries from the list.

14

Save the changes and close the form.

15

Click on the Multi-homing tab and click Create. The BGP Multi-homing Site (Create) form opens.

16

Configure the required parameters in the Service Panel.

You cannot turn up a BGP multi-homing site without specifying a Multi-homing ID. In addition, the Multi-homing ID of the site can only be removed at a later time when the site is shutdown.

17

Select a SAP, Spoke SDP, Split Horizon Group, or Mesh SDP Binding, based on the option you selected for the Enable Multi-homing to parameter.

You cannot turn up a BGP multi-homing site without specifying a SAP. The SAP configured under this site can be part of an endpoint, but the only other object allowed in the endpoint is an ICB spoke-SDP. In addition, an assigned SAP can only be removed at a later time when the BGP multi-homing site is shutdown.

18

Configure the required parameters in the Timer panel.

If you enable either the Use Node Level Boot Timer and/or Use Node Level Site Activation Timer parameters, then the associated Boot Timer (seconds) and/or Activation Timer (seconds) parameters are not configurable. These parameter values are inherited from the network element configuration.

The Boot Timer (seconds) and Activation Timer (seconds) parameters can be configured for an NE on the BGP Multi-homing sub-tab under the Redundancy tab in the Network Element (Edit) form. See 12.5 “To modify NE properties” (p. 343) for more information about how to change device properties.

If you enable the Use Node Level Site-Down Minimum Timer parameter, the Site-Down Minimum Timer (seconds) parameter and Site-Down Minimum Timer Remaining (seconds) parameters are ignored and the values configured on the NE are used.

19

Select a Monitored Group Name.

20

Save the changes and close the form.

21

Check the following indicators:

- Operational State: indicates the operational status of the multi-homing site.
- Designated Forwarder: indicates whether this site has been declared as designated forwarder, depending on the result of the BGP election.

22

Save the changes and close the forms.

END OF STEPS

77.112 To re-evaluate the PW Templates associated with a BGP AD or BGP VPLS

77.112.1 Purpose

Use this procedure to re-evaluate changes made to the Route Targets associated with the PW Template bindings of an existing site with BGP configuration. This procedure can only be performed on an existing VPLS that has been configured with BGP AD or BGP VPLS. The procedure allows you to make configuration changes and propagate them to the service without having to shutdown and then turn up a site.

77.112.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu.

2

Choose the required VPLS and click Properties. The VPLS Service (Edit) form opens.

3

On the service tree, expand Sites and click on the site for which you want to re-evaluate the PW templates.

4

Click on the BGP tab.

5

Click Configuration to open the BGP Configuration form.

6

Make any required changes to the Route Targets or PW Template bindings. See 77.110 “To configure a site for BGP AD or BGP VPLS” (p. 2408) or 77.111 “To configure a site for BGP VPLS Multi-homing” (p. 2412) as required, for detailed instructions on configuring these items.

7

Click on the PW Template Binding tab and click Apply to apply your changes.

8

Perform one of the following:

a. If you need to update the PW Templates under other sites for this service, repeat Step 4 to Step 7 as required, then go to Step 11 .

b. If you do not need to update the PW Templates under other sites for this service go to Step 9 .

9

Click on the Re-evaluate PW Template button to run an evaluation of the PW Template bindings. A pop-up window appears indicating if the re-evaluation was successful. If it was not, the reason for the failure is displayed.

If you make any subsequent modifications, you can re-evaluate the template again.

10

Save the changes and close the forms.

11

Return to the VPLS Service (Edit) form.

12

Click Re-evaluate PW Template. The Add form opens to allow you to select one of the PW Templates you modified.

13

Select the required PW template.

14

Click OK to run an evaluation of the selected PW Template. A pop-up window appears indicating if the re-evaluation was successful. If it was not, the reason for the failure is displayed.

15

Repeat Step 12 to Step 14 for any other sites that you want to re-evaluate a PW Template for.

16

Close the forms.

END OF STEPS

77.113 To assign tunnel administrative groups to a BGP or BGP AD VPLS

77.113.1 Purpose

Perform this procedure to assign tunnel administrative groups to a BGP or BGP AD VPLS. The tunnel administrative groups are applicable only when the VPLS is configured as a PBB tunnel. The VPLS must be configured to include B-sites and Multi-Segment Tunnel Selection must be enabled in the system preferences. See Chapter 88, “Tunnel administrative groups” for more information about tunnel administrative groups.

77.113.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu.

2

Select the required BGP or BGP AD VPLS and click Properties. The VPLS Service (Edit) form opens.

3

Click on the Tunnel Admin Groups tab.

4

Select the tunnel administrative groups in the Unassigned list that you want to assign to the tunnel and click the right arrow to move them into the Assigned list.

5

Close the forms.

END OF STEPS

78 IES management

IES management

78.1 Overview

78.1.1 General information

An IES is a routed connectivity service in which the customer traffic passes through an L3 IP router interface to the Internet.

IES allows customer-facing IP interfaces in the same routing instance to be used for service network core-routing connectivity. IES requires that the IP addressing scheme that is used by the customer be unique among other provider addressing schemes and potentially the entire Internet.

Packets that arrive at the edge device are associated with an IES based on the access interface on which they arrive. An access interface is uniquely identified using:

- port
- service ID
- IP address

78.1.2 IES configuration overview

The NFM-P supports end-to-end IES configuration using tabbed configuration forms with an embedded navigation tree.

The NFM-P supports the configuration in IES of an L3 aggregation mechanism called routed CO. Routed CO uses DHCP relay to manage dynamic subscriber hosts; the network resources for static subscriber hosts are explicitly provisioned. Routed CO supports all residential subscriber management functions of the NFM-P. See Chapter 74, “Residential subscriber management” for more information about residential subscriber management and routed CO.

Routed CO uses a subscriber interface that defines up to 256 subnets. A subscriber interface has child objects called group interfaces. A group interface supports the configuration of multiple SAPs as child objects. A SAP in a group interface supports all residential subscriber management functions. A group interface does not allow the specification of IP subnets or addresses, but inherits the addressing scheme of the parent subscriber interface. The NFM-P service topology map displays IES subscriber interfaces, group interfaces, and the associated SAPs.

You can configure NAT for dynamic subscriber hosts in a routed CO deployment. NAT implementation in an IES requires a NAT configuration on the NE base routing instance and a NAT policy that is associated with a subscriber profile. See Chapter 27, “NE routing and forwarding” for information about configuring and deploying NAT on a base routing instance. See Chapter 66, “NAT policies” for information about configuring a NAT policy. See Chapter 74, “Residential subscriber management” for information about associating a NAT policy with a subscriber profile.

When you use the NFM-P to create or discover a service, the NFM-P assigns a default Service Tier value to the service. The Service Tier parameter value is relevant only in the context of composite

service topology views. See Chapter 85, “Composite service management” for more information about the hierarchical organization of composite services.

Common to all device services, such as IES, are policies that are assigned to the service. Policies are defined at a global level and can then be applied to components of the service, such as interfaces or circuits, when the service is configured or modified. The following policies are common to all device services:

- QoS policies define ingress classification, policing, shaping, and marking on the ingress side of the interface. QoS policies are configured using the SAP Access Ingress Policy form, the SAP Access Egress Policy form, and the ATM QoS Policy form.
- Policer control policies to control access ingress policers and access egress policers under a common hierarchy. Policer control policies are configured using the Policer Control Policy Manager.
- Scheduling policies define hierarchical rate limiting and scheduling to govern the scheduling of queues. Scheduler policies are configured using the Scheduler Policy form.
- Port scheduler policies define hierarchical bandwidth allocation and scheduling at the egress port level. Port scheduler policies are configured using the Port Scheduler Policy form.
- Filter policies control network traffic into or out of an interface or circuit based on DHCP, IP, or MAC matching criteria. Filter policies are configured using the ACL IP Filter form, the ACL MAC Filter form, and the DHCP Filter form.
- Accounting policies measure the traffic on a service to ensure proper billing and enforcement of SLAs. Accounting policies are configured using the Accounting Policy form.
- ANCP policies provide status and control information based on port-up and port-down messages and current line rate changes between the edge device and the access NE. ANCP policies are configured using the Manage Subscriber Policies form.
- Routing policies control the size and content of the routing tables, the routes that are advertised, and the best route to take to reach a destination. Routing policies are configured using the Routing Policy Manager.
- Time of day suites specify time and day restriction policies that are assigned to QoS policies and schedulers, ACL filters, and aggregation schedulers. Time of day suites and time range policies are configured using the Time of Day Suite form and Time Range form, respectively.

See Chapter 49, “Policies overview” for more information about policies.

Although IES is part of the routing domain, the usable IP address space may be limited. IES allows a portion of the service provider address space to be reserved for service IP provisioning and to be administered by a separate, but subordinate, address authority.

Multiple IESs can be created to separate customer-owned IP interfaces. More than one IES can be created for one customer. More than one IP interface can be created in one IES. All IP interfaces created in an IES belong to the same customer.

The IES IP interfaces are restricted to the routing protocols that can be defined on the interface based on the fact that the customer has a different routing domain for this service. The IP interfaces support the following routing protocols:

- RIP
- OSPF

- BGP
- IS-IS
- PIM
- IGMP

Customer routes can be advertised to the network core using static routes, RIP, or BGP. BGP and static routes are the most commonly used routing methods.

An IES can be connected to a VLL service or to a VPLS by an internal cross-connect through a CCAG adapter. This configuration eliminates the need for the physical port, cable, and other MDA-specific components and results in a less costly and more reliable interconnection. See Chapter 85, "Composite service management" for information about joining services to form composite services.

78.1.3 ATM SAP terminations for IES

CE routers that have access to an ATM network can connect with an IES service using ATM SAP terminations on a 7750 SR. The interconnection between ATM point-to-point and L3 services uses RFC 2684-encapsulated IPv4 traffic over an ATM PVC that terminates on a specially configured SAP. All RFC 2684-encapsulated traffic can be routed over ATM networks, frame relay, and directly through ATM connections.

The figure below shows how CPE router A in an existing ATM network can access L3 IP services, such as an IES, using a statically configured ATM PVC on a 7750 SR (SR #1). CPE router B is connected to a frame relay, which connects to ATM switch 2 through IWF (service interworking). The RFC 2684-encapsulated traffic moves from both CPE routers through the ATM access network to a SAP configured on a 7750 SR #1 to serve a specific IES. At the same time, SDPs on the router are configured to a service to forward traffic over the IP/MPLS core. Destination CPE router C can receive RFC 2684-encapsulated traffic over an IP network over an ATM switch connected directly using 7750 SR #2.

Figure 78-1 ATM SAP network connection to an IES

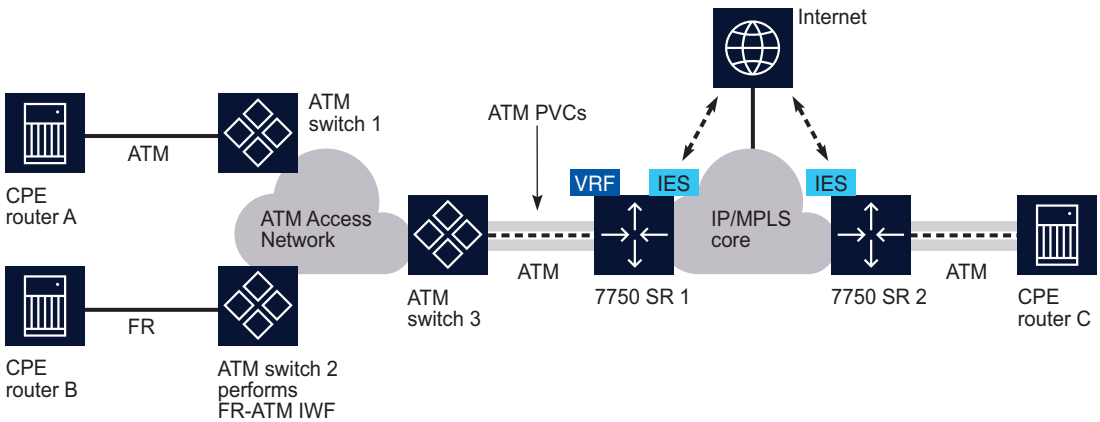

The two connection methods used between an ATM network and the IES router: LLC/SNAP encapsulation and VC multiplexing.

A VLL Epipe service can terminate directly on an IES service using an SDP spoke on the 7750 SR, 7450 ESS, or 7950 XRS. Traffic that terminates on an IES service is identified by the interface ID of the SDP on the L2 access router and the VC ID label in the service packet. All routing protocols supported by IES are also supported for spoke SDP termination.

The figure below shows a spoke SDP terminating directly on an IES. The spoke SDP could be tied to VLL Epipe or VPLS. No configuration is required for the CE-to-PE connection on the SAP.

Figure 78-2 SDP spoke termination on an L2 service

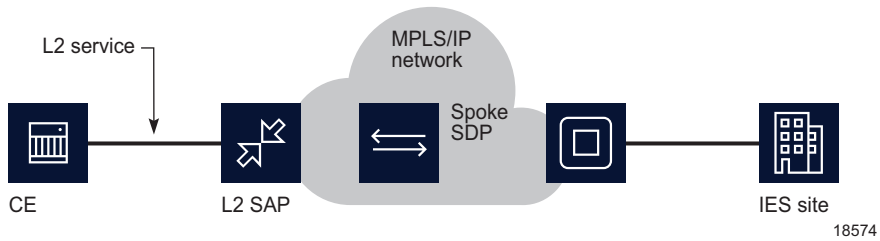

78.1.4 Routed CO dual homing using SRRP

Subscriber Router Redundancy Protocol (SRRP) allows two separate connections to an access NE such as DSLAM to operate in an active/standby configuration similar to the way in which VRRP interfaces operate. SRRP is a collection of functions and messaging protocols that allows a system to create a set of redundant gateway IP addresses that are shared by a local and remote NE.

Each SRRP instance is created within the context of a subscriber group IP interface and is identified by a unique SRRP instance ID, which must be unique within the NE. This SRRP instance controls the redundant routing for all subscriber subnets configured or associated with the group interface. One SRRP instance is supported for each group interface and the SRRP ID must be the same as the SRRP instance ID on the group IP interface on the redundant NE.

A subscriber subnet redundant gateway IP host address is assigned at the subscriber IP interface level and is used for each SRRP instance associated with the subscriber subnet. The redundant IP host address must be configured for a subscriber subnet before it can be associated with an SRRP instance.

When SRRP is active on a group interface, the SRRP instance advertises to a remote NE using in-band messaging on the group-interface SAPs and out-of-band messaging on the group-interface redundant interface. If the remote NE uses the same SRRP instance ID, one NE enters a master state, while the other NE enters a backup state. Since the NEs share a common SRRP gateway MAC address (used for the SRRP gateway IP address and for proxy ARP functions), either NE can act as the default gateway for the attached subscriber hosts. This functionality helps to preserve subscriber QoS enforcement. The master state allows routing to and from the subscriber hosts associated with the group IP interface. The backup state stops ingress forwarding for packets

destined to the SRRP gateway MAC and causes all packets destined to subscriber hosts on the group IP interface to be forwarded to a redundant IP interface associated with the group IP interface.

Normally, when anti-spoofing is enabled on a group-interface SAP, the SAP drops SRRP packets because they do not contain a subscriber MAC or IP address. However, you can use a configuration option to enable anti-spoofing for subscriber hosts on a group-interface SAP that participates in SRRP advertisements.

The underlying mechanism that controls state transitions is based on a dynamic priority level that an SRRP instance maintains. The SRRP instance with the highest priority level assumes the master operating state. An SRRP instance with a higher current priority level always preempts an SRRP instance with a lower priority level. If the priority levels are equal, the SRRP instance with the lowest source SRRP host IP address assumes the master state. The local SRRP instance priority may also be controlled by associating the instance with an existing VRRP policy.

To prevent a flood of AccessInterfaceDown alarms that an SRRP fault or link failure may generate for LAG-based MSAPs, the NFM-P performs alarm suppression. See Chapter 74, “Residential subscriber management” for more information.

The redundant IP interface is a special interface that connects two systems with one or more common SRRP instances. The interface is configured with a /31 address and a spoke SDP binding, creating an Ethernet pseudowire shortcut between the redundant NEs. When the SRRP instance is in backup state, the group interface associated with this instance is not allowed to forward or route traffic downstream towards the subscriber. As a result of this, the packets are shunted across the redundant interface so that the active group interface does the forwarding or routing.

If the redundant IP interface goes down, the system allows the group IP interfaces associated with the down interface to forward locally downstream, when they are in the backup SRRP state. While forwarding downstream in the backup state, the system uses the MAC address associated with the group IP interface, not the SRRP redundant gateway MAC address.

SRRP is supported on the 7450 ESS in mixed mode and 7750 SR.

78.1.5 DoS protection

To protect an IES from a high incoming packet rate that characterizes a DoS attack, you can use the NFM-P to create DoS protection policies for the IES L3 access interfaces. A DoS protection policy limits the number of control-plane packets that an interface receives each second, and optionally logs a violation notification if a policy limit is exceeded. You can use the NE System Security form to view the violations for a specific NE.

You can configure a DoS protection policy to control the following on an IES L3 access interface:

- the control-plane packet arrival rate per subscriber host on the interface
- the overall control-plane packet arrival rate for the interface
- whether an NE sends a notification trap if a policy limit is exceeded

Each IES L3 access interface on an NE that supports DoS protection is automatically assigned a default DoS protection policy. This default policy limits only the overall packet arrival rate for the interface, and cannot be deleted or modified. See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for information about creating a DoS protection policy.

78.1.6 DDoS protection

To protect an IES from a high incoming packet rate that characterizes a DDoS attack, you can use the NFM-P to configure TMS interfaces to route malicious traffic through the ISA-TMS MDA where the malicious traffic is cleaned before being released to the network.

The TMS interface on an IES consists of one IES and two VPRNs as follows:

- The management VPRN communicates with the TMS server to determine whether incoming packets are malicious. Configuration of this VPRN is optional.
- The Off-Ramp VPRN receives the malicious traffic and routes it through the ISA-TMS MDA where the traffic is scrubbed. Configuration of this VPRN is mandatory.
- The On-Ramp IES returns the cleaned traffic to the network. Configuration of this IES is mandatory.

See 78.25 “To add a TMS interface to an IES” (p. 2467) for information about creating a TMS interface on an IES.

You can configure a DDoS protection policy on an IES group interface SAP or L3 access interface. See the procedure to configure an NE DDoS protection policy in the *NSP System Administrator Guide* for information about configuring a DDoS protection policy.

78.1.7 Local DHCP servers

A local DHCP server can be associated with a network interface or L3 access interface on an IES. See 74.2 “Residential subscriber components” (p. 1992) .

78.1.8 Local user database

A local user database can be associated with a local DHCP server and PPPoE configurations on group interfaces. See 74.2 “Residential subscriber components” (p. 1992) .

78.1.9 PPPoE protocol on IES

An IES can be configured to support PPPoE. PPPoE is used in subscriber networks to encapsulate PPP frames inside Ethernet frames. PPPoE combines the point-to-point protocol used by DSL sessions with Ethernet framing to support multiple subscribers in a LAN. Using the group interface configuration form, you can assign a PPPoE policy and a local user database to authenticate PPPoE subscribers.

78.1.10 L2TP configuration for IES

An IES group interface can be configured to terminate LNS PPP sessions. L2TP is a session-layer protocol that extends the PPP model by allowing L2 and PPP endpoints to reside on different devices that are interconnected by a PSN. L2TP extends the PPP sessions between the CPE and PPP/L2TP termination points on the L2TP network server (LNS), via an intermediate L2TP access concentrator (LAC). The LAC is the initiator of session-generated L2TP tunnels; the LNS is the server that waits for new tunnels. Manually configured and initiated L2TP tunnels can be initiated or stopped from either the LNS or LAC.

After a tunnel is established, the network traffic between the peers is bidirectional. If a tunnel carrying a session fails, another tunnel from the same tunnel group re-establishes the session. Within each L2TP tunnel, one or more L2TP sessions can exist. Each L2TP session transports PPP packets.

At least one ISA-LNS group must be configured for the LNS NE.

On an LNS NE, L2TP destinations configured for L2TP tunnel profiles can include the following:

- loopback L3 access interfaces for a VPRN or IES service
- loopback interfaces configured for a base routing instance

See Chapter 13, “Logical group object configuration” for more information about ISA-LNS groups. See 13.12 “To configure an ISA-LNS group” (p. 426) for information about creating and configuring an ISA-LNS group. See Chapter 28, “Routing protocol configuration” for more information about L2TP. See 78.19 “To configure a group interface on an IES” (p. 2449) for information about configuring an IES group interface to terminate LNS PPP sessions.

78.1.11 PIM on IES group interfaces

You can configure Protocol-Independent Multicast (PIM) functionality on an IES group interface. PIM on a subscriber group interface allows SAP-level replication over ESM group interfaces by establishing PIM adjacency to a downstream L3 router. On each group interface, a single Ethernet SAP is configured (LAG or physical port). Multiple default hosts can be configured, and a static host is configured for connectivity to a downstream L3 aggregation device. When PIM is enabled on the static host interface, a SAP configured under the group interface is added to the OIF list for SAP-level replication towards a downstream L3 router.

To configure PIM on an IES group interface, see 78.24 “To configure a PIM interface on an IES group interface” (p. 2466) .

78.2 IES configuration

78.2.1 Sample IES configuration

Figure 78-3 Sample IES configuration

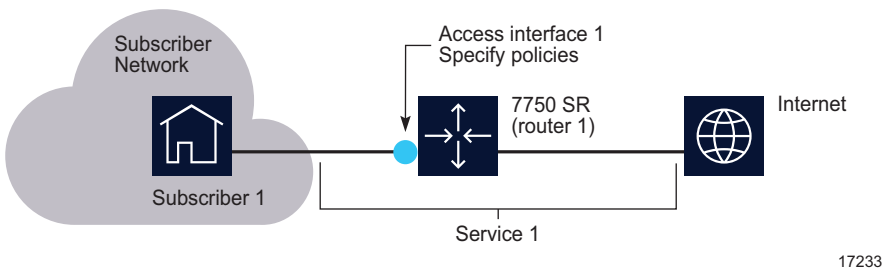

78.2.2 Sample configuration steps

The following high-level tasks are necessary to configure this sample IES.

1

Configure policies as required, for example, QoS access ingress and egress interface, Scheduler, ACL IP, Accounting, and ANCP.

2

Create and configure Subscriber 1.

3

Create and configure Service 1.

78.2.3 Workflow to configure an IES

The following workflow lists the high-level steps required to create an IES. As a prerequisite for creating a IES, this workflow assumes the following:

- a group or customer with the required user access privileges has been configured.
- the IP or IP/MPLS core network exists.
- any required service tunnels are created including the static or dynamic LSP required to create the service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
- the access ports for the service are created; see Chapter 16, “Port and channel object configuration” for more information.
- any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite policies are created; see Chapter 49, “Policies overview” for more information. You do not have to create pre-defined policies if policies are created on a per-service basis.
- any required MP-BGP for PE-to-PE routing is configured; see Chapter 28, “Routing protocol configuration” for more information about protocol configuration.

78.2.4 Stages

1

Create the IES. See 78.3 “To create an IES” (p. 2430) .

1. Choose a customer to associate with the IES.
2. Define the general properties for the IES.
3. Turn up the IES service.

2

As required, configure AA reporting. See 78.4 “To configure an IES for AA reporting” (p. 2430) .

3

As required, create an IES site or GNE IES site. See 78.5 “To configure an IES site” (p. 2431) and 78.6 “To configure a GNE site on an IES service” (p. 2432) .

4

As required, perform one or more of the following on an IES site.

a. Apply routing protocols, such as OSPF, RIP, or IS-IS, to an IES site, if required. See 78.7 “To apply OSPF, RIP, or IS-IS to an IES site” (p. 2433) .

b. Add an IGMP interface to an IES. See 78.8 “To add an IGMP interface to an IES” (p. 2433) .

c. Add a PIM interface to an IES. See 78.9 “To add a PIM interface to an IES” (p. 2435) .

d. Create an L2 SDP spoke termination on an IES service. See 78.10 “To create an L2 SDP spoke termination on an IES service” (p. 2437) .

e. Add a subscriber interface to an IES. See 78.16 “To configure a subscriber interface on an IES” (p. 2444) .

f. Add an AA interface to an IES. See 78.17 “To add an AA interface to an IES or a VPRN site” (p. 2447) .

g. Add an AARP interface to an IES. See 78.18 “To add an AARP interface to an IES or a VPRN site” (p. 2448) .

h. Add a group interface to an IES. See 78.19 “To configure a group interface on an IES” (p. 2449) .

i. Configure WLAN GW for an IES group interface. See 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462) .

j. Add a TMS interface to the IES. See 78.25 “To add a TMS interface to an IES” (p. 2467) .

k. Implement dual homing using SRR. See 78.26 “To implement dual homing using SRRP” (p. 2469) .

l. Create and enable a video interface for the IES site, if required. See 35.3 “To add a video interface to an IES or VPRN site” (p. 1275) .

5

As required, create an L3 access interface on an IES site. See 78.28 “To configure an L3 access interface on an IES site” (p. 2472) .

6

As required, perform one or more of the following on an IES L3 access interface.

a. Configure custom object attributes for AA reporting. See 78.4 “To configure an IES for AA reporting” (p. 2430) .

b. Bind the L3 access interface to a VPLS site. See 78.30 “To bind an IES L3 access interface to a VPLS site or VPLS I-site” (p. 2475) .

c. Apply OSPF, RIP, or IS-IS to an IES L3 interface. See 78.31 “To apply OSPF, RIP, or IS-IS to an IES L3 interface” (p. 2476) .

d. Assign ingress and egress QoS policies. See 78.32 “To assign ingress and egress QoS policies to an IES L3 access interface” (p. 2477) or 78.33 “To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site” (p. 2479).

e. Configure scheduling. See 78.34 “To configure scheduling on an IES L3 access interface” (p. 2481) .

f. Assign ingress and egress ACL filters. See 78.35 “To assign ingress and egress ACL filters to an IES L3 access interface” (p. 2483) .

g. Assign a virtual port to an IES L3 access interface. See 78.36 “To assign a virtual port to an IES L3 access interface” (p. 2484).

h. Associate a local DHCPv4 or DCHPv6 server. See 78.37 “To associate a local DHCPv4 or DHCPv6 server to an IES L3 access interface” (p. 2484) .

i. Assign an accounting policy. See 78.38 “To assign an accounting policy to an IES L3 access interface” (p. 2485) .

j. Assign an ANCP policy. See 78.41 “To assign an ANCP policy to an IES L3 access interface” (p. 2488) .

k. Assign a time of day suite. See 78.43 “To assign a time of day suite to an IES L3 access interface” (p. 2490) .

l. Configure residential subscriber management. See 78.44 “To configure residential subscriber management for an IES L3 access interface” (p. 2490) .

m. Assign a DoS protection policy or DDoS protection policy. See 78.45 “To assign a DoS protection policy or DDoS protection policy to an IES L3 access interface” (p. 2491) .

n. Assign an IP address. See 78.46 “To assign an IP address to an IES L3 access interface” (p. 2492) .

o. Configure BFD. See 78.47 “To configure BFD for an IES L3 access interface” (p. 2493) .

p. Configure ICMPv4. See 78.48 “To configure ICMPv4 for an IES L3 access interface” (p. 2494) .

q. Configure ICMPv6. See 78.49 “To configure ICMPv6 on an IES L3 access interface” (p. 2495) .

r. Assign an ICMP Ping template. See 78.50 “To assign an ICMP ping template to an IES L3 access interface” (p. 2495).

s. Configure ARP. See 78.51 “To configure ARP for an IES L3 access interface” (p. 2497) .

t. Configure neighbor discovery. See 78.52 “To configure neighbor discovery on an IES L3 access interface” (p. 2497) .

u. Configure DHCPv4. See 78.53 “To configure DHCPv4 for an IES L3 access interface” (p. 2498) .

v. Create a VRRP instance. See 78.54 “To create a VRRP instance on an IES L3 access interface for a virtual router” (p. 2500) .

w. Configure anti-spoofing filters. See 78.55 “To configure anti-spoofing filters for an IES L3 access interface” (p. 2501) .

x. Configure router advertisement. See 78.56 “To configure router advertisement on an IES L3 access interface” (p. 2502) .

y. Specify queue overrides. See 78.57 “To specify QoS policy overrides on an IES L3 access interface” (p. 2503) .

z. Configure DHCPv6. See 78.58 “To configure DHCPv6 on an IES L3 access interface” (p. 2504) .

aa. Configure an MLD interface on the L3 access interface. See 28.124 “To configure an MLD interface on an IES L3 access interface” (p. 1037) .

7

As required, view IES information.

a. View the IES service operational status. See 78.61 “To view the service operational status” (p. 2506) .

b. View the service topology map associated with an IES. See 78.62 “To view the service topology” (p. 2507) .

8

As required, modify an IES.

a. Using the Manage Services form. See 78.63 “To modify an IES” (p. 2508) .

b. Using the topology view. See 78.64 “To modify an IES using the topology view” (p. 2509) .

9

As required, delete an IES. See 78.65 “To delete an IES” (p. 2511) .

IES management procedures

78.3 To create an IES

78.3.1 Steps

1

Choose Create→Service→IES from the NFM-P main menu. The IES Service (Create) form opens.

2

Select a customer to associate with the IES.

3

Configure the required general parameters.

The Service ID and SVC Mgr Service ID parameters are configurable when the Auto-Assign ID parameter is disabled.

4

Save the changes and close the forms.

END OF STEPS

78.4 To configure an IES for AA reporting

78.4.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES in the list and click Properties. The IES Service (Edit) form opens.

3

Click on the Application Assurance tab, then on the NSP Analytics Parameters tab.

4

Configure the required parameters.

5

To configure one or more custom DCP groups, see the custom DCP group configuration procedure in Chapter 87, “Application assurance” .

6

To specify one or more application or application groups for IP detail reporting:

1. Click on the IP Detail Application or IP Detail Application Group tab, as required. The IP Detail Application or IP Detail Application Group (Create) form opens.
2. Select an application or application group from the Select Application or Select Application Group form.
3. Save the changes and close the forms.

7

Save the changes and close the forms.

Note: You can specify up to 10 applications and up to 10 application groups.

END OF STEPS

78.5 To configure an IES site

78.5.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, right-click on the Site icon and choose Create IES Site and select a site, or expand the Site icon and choose Properties. The IES Site (Create|Edit) form opens.

4

Configure the general site parameters.

Tunnel Fault Notification is configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept in order to receive the fault notification from the tunnel facility MEP.

5

Save the changes and close the forms.

END OF STEPS

78.6 To configure a GNE site on an IES service

78.6.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, choose a site by right-clicking on IES Service and choosing Create GNE Site. The GNE Site (Create) form opens.

4

Configure the general site parameters.

5

Click Apply.

6

To configure an interface for the GNE site:

1. Click on the GNE Service Interfaces tab and click Create. The GNE Service Interface (Create) form opens.
2. Configure the required parameters.
3. Click on the Ports tab and select a generic NE interface.
4. Configure the required parameters.
5. Save the changes and close the form.

7

Save the changes and close the forms.

You can use the topology maps to view the service. See Chapter 4, “Topology map management” for more information about service topology maps.

END OF STEPS

78.7 To apply OSPF, RIP, or IS-IS to an IES site

Note: OSPF, RIP, or IS-IS must be enabled at the routing instance level before you can apply OSPF, RIP, or IS-IS to an IES.

78.7.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES in the list and click Properties. The IES Site (Edit) form opens.

3

On the service tree, expand the Sites icon and click on a site. The IES Site (Edit) form opens.

4

Click on the Routing tab, then the Protocols tab.

5

Click Add. The Create Interface form opens.

6

Specify the interface type by configuring the What type of interface would you like to create? parameter.

7

Click OK. The Interface (Create) form opens for the selected protocol. See “Routing protocol configuration workflow and procedures” (p. 909) for information about configuring specific routing protocols.

END OF STEPS

78.8 To add an IGMP interface to an IES

Note: IGMP must be enabled on the NE routing instance before you can create an IGMP interface.

78.8.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand the Sites icon and click on a site. The IES Site (Edit) form opens.

4

Click on the Multicast tab, then the Interfaces tab.

5

Click Add. The Create Interface form opens.

6

Set the What type of interface would you like to create? parameter to IGMP.

7

Click OK. The IGMP Interface (Create) form opens.

8

Select an interface and configure the parameters.

9

Click on the Behavior tab and configure the parameters.

10

To add a multicast CAC policy:

1. Click on the Multicast CAC tab.
2. Select a multicast CAC policy in the Multicast CAC Policy panel.
3. Configure the required parameters.
4. Click on the Levels tab and click Create. The Multicast CAC Level, IGMP Interface (Create) form opens.
5. Configure the parameters and close the form.
6. Click on the LAG Port Down tab and click Create. The LAG Port Down, IGMP Interface (Create) form opens.
7. Configure the parameters and close the form.

11

To configure SSM translation:

1. Click on the SSM Translation tab and click Create. The SSM Translation, IGMP Interface (Create) form opens.

2. Configure the required parameters and close the form.

12

To add a static multicast group or source:

1. Click on the Static Group/Source tab and click Create. The StaticGrpSrc, Interface ID - 6, Routing Instance (Create) form opens.
2. Configure the required parameters, save the changes and close the form.

13

Save the changes and close the form.

14

On the multicast tab, click on the IGMP Group Interfaces tab to identify an IGMP group interface for the IES service.

1. Click Create. The IGMP Group Interface - Routing Instance (Create) form opens.
2. Configure the required parameters.
3. Select an IGMP group interface in the IGMP Group Interface panel and configure the parameters.
4. Click on the Behavior tab. select an import policy and configure the required parameters.
5. Click on the Multicast CAC tab to add a multicast CAC policy.
6. Select a multicast CAC policy in the Multicast CAC Policy panel and configure the required parameters.
7. Close the form.

15

Close the forms.

END OF STEPS

78.9 To add a PIM interface to an IES

Note: Before you can add a PIM interface to an IES, PIM must be applied to All or IES during the PIM configuration at the routing instance level. See 28.98 “To configure PIM on a routing instance” (p. 998) for more information.

78.9.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand the Sites icon and click on a site. The IES Site (Edit) form opens.

4

Click on the Multicast tab.

5

Click Add. The Create Interface form opens.

6

Set the What type of interface would you like to create? parameter to PIM.

7

Click OK. The PIM Interface (Create) form opens.

8

Select an interface and configure the required parameters.

9

Click on the Behavior tab and configure the required parameters.

The Sticky DR Priority parameter is configurable when the Sticky DR parameter is enabled.

10

To add a multicast CAC policy:

1. Click on the Multicast CAC tab.
2. Select a multicast CAC policy in the Multicast CAC Policy panel.
3. Configure the required parameters.
4. Click on the Levels tab and click Create. The Multicast CAC Level, IGMP Interface (Create) form opens.
5. Configure the parameters and close the form.
6. Click on the LAG Port Down tab and click Create. The LAG Port Down, IGMP Interface (Create) form opens.
7. Configure the parameters and close the form.

11

Click on the IPv6 Specifics tab and configure the IPv6 Administrative State parameter.

12

Close the forms.

END OF STEPS

78.10 To create an L2 SDP spoke termination on an IES service

78.10.1 Purpose

Ensure that a service and site are created in the IES. To terminate an L2 service on an IES SDP spoke, you must identify the VC and an interface that belong to the VC. The interface must not have an associated port.

78.10.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Sites→site, right-click on Spoke Sdp Bindings and choose Create Spoke SDP Binding. The Spoke SDP Binding (Create) form opens.

4

Select a source interface for the spoke SDP binding in the Source Interface panel.

5

Specify a destination NE for the spoke SDP binding:

a. If the destination NE is a managed NE, select an NE in the Tunnel Termination Site panel.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

6

Configure the required general parameters.

These parameters appear only when a source interface is assigned to the spoke SDP binding.

The range of the Ingress Label parameter depends on the parameter value set for the Static Label Range on the MPLS instance. See 31.6 “To configure an MPLS instance” (p. 1116) .

7

Configure the parameters in the Hash Label panel.

The Enable Hash Label and Enable Signal Capability parameters can only be configured for spoke-SDP bindings that are access interface terminated.

8

Perform one of the following to specify a transport tunnel for the spoke SDP binding.

a. Let the NFM-P configure the transport tunnel automatically by enabling the Auto-Select Transport Tunnel parameter and configuring the Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually by selecting a tunnel in the Tunnel panel.

c. Configure an MPLS-TP transport tunnel manually by selecting a tunnel in the tunnel panel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for more information about how to create an MPLS-TP service tunnel.

9

To specify an application profile for the spoke SDP binding, by selecting an application profile string.

Note: The Application Profile String: - Spoke SDP Binding - IES service form displays only local profiles on the NE.

10

Choose one of the following AA transit policy types, if required.

Note: You can only associate one AA transit policy type with a service object.
To bind a transit policy to an L3 access interface, a port must already exist on the interface.
You can bind a transit policy to only one L3 access interface or spoke SDP binding per NE.
The transit policy and the application profile must belong to the same application assurance group or partition.

a. Associate an AA transit IP policy with the service object by selecting a transit IP policy in the Transit IP Policy panel.

b. Associate an AA transit prefix policy with the service object by selecting a transit prefix policy in the Transit Prefix Policy panel.

11

To associate an AA redundant protocol with the service object, select an AA redundant protocol in the AA Redundant Protocol panel.

12

Configure the AARP Service Reference Type parameter.

13

To configure custom object attributes for AA reporting:

1. Click on the NSP Analytics Parameters tab, then the Reporting tab.
2. Click Create. The AA Reporting (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.
5. Click on the General tab.

14

To configure QoS:

1. Click on the QoS tab.
2. Select a network policy and an ingress queue group template policy in the Forwarding Plane Redirect panel.
3. Configure the Instance ID parameter in the Forwarding Plane Redirect panel.
4. Select a network policy and egress queue group template policy to assign to the Spoke SDP Binding in the Port Redirect panel.
5. Configure the Instance ID parameter in the Port Redirect panel.

15

Click on the States tab and configure the Administrative State parameter.

16

To assign ingress and egress ACL filters to the spoke SDP binding:

1. Click on the ACL tab.
2. Select an ingress ACL filter in the Ingress Filter panel.
3. Select an egress ACL filter in the Egress Filter panel.

17

To assign an accounting policy to the spoke SDP binding:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy.

18

Note: A default DoS protection policy is automatically assigned to the SAP.

To assign a DoS protection policy or DDoS protection policy to the SAP:

1. Click on the Security tab.
2. Select a DoS protection policy in the NE DoS Protection panel.
3. Select an NE DDoS protection policy.
4. Configure the required parameters.

19

To associate a MEP with an SDP binding:

1. Click on the OAM tab, then the ETH-CFM tab.
2. Click Create in the MEPs panel. The MEP (Create) form opens.
3. Select a maintenance entity group.
4. Configure the required parameters.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.
5. If the MD for the MEP has a Name Type of none and the associated MEG has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable. Click on the Y.1731 Tests tab. Otherwise, go to Step 20 .
6. Configure the required parameters.

   The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.
7. Click on the AIS tab and configure the required parameters.

20

Save the changes and close the forms.

END OF STEPS

78.11 To configure an MPLS-TP static pseudowire on an IES spoke SDP binding

78.11.1 Purpose

An MPLS-TP service tunnel must be used in the SDP binding, and the Control Word parameter for pseudowire OAM must be set to Preferred. See 78.10 “To create an L2 SDP spoke termination on an IES service” (p. 2437) .

78.11.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an IES and click Properties. The IES Service (Edit) form opens.

3

Click on the Control Channel tab.

4

Configure the required parameters.

5

Configure the static PW:

1. Click on the Static PW tab.
2. Click Create. The PW Path ID (Create) form opens.
3. Configure the Path AGI parameter.
4. Configure the parameters in the Source Attachment Individual Identifier panel.
5. Configure the parameters in the Target Attachment Individual Identifier panel.
6. Click OK. The PW Path ID (Create) form closes.

6

Save the changes and close the forms.

END OF STEPS

78.12 To configure BFD on an IES spoke SDP binding

78.12.1 Purpose

BFD is used over the VCCV control channel for PW fault detection. BFD carried over a PW associated channel enables the monitoring of the PW between the terminating PEs, regardless of whether the service spans multiple hops. This allows faults that are local to individual PWs to be detected.

78.12.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and configure BFD on the spoke SDP binding:

1. On the General tab, enable the Enable BFD parameter.
2. Choose a BFD template. To create a BFD template, see 28.25 “To configure a BFD template policy” (p. 911) .

   You must assign a BFD template to the spoke SDP binding if you enable BFD.
3. In the Failure Action dropdown, select the action to be performed on BFD failure.
4. Configure the Up-Timer parameter.

   Note:

   The Up-Timer parameter is applicable only when the value of Failure Action is set to down.

5

Save and close the forms.

END OF STEPS

78.13 To clear BFD sessions and statistics on an IES SDP binding

78.13.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to view the BFD session information.

4

Clear BFD sessions or statistics on the spoke SDP binding:

1. Click on the BFD tab, then on the BFD Session tab.
2. Click Clear All to clear all BFD sessions.
3. Click Clear All Statistics to clear all BFD statistics.

5

Close the forms.

END OF STEPS

78.14 To view the BFD session status on an IES SDP binding

78.14.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and view the status of the BFD session on the spoke SDP binding:

1. Click on the BFD Session tab.
2. Choose a BFD session and click Properties. The BFD Session (View) form opens.

5

Close the forms.

END OF STEPS

78.15 To view the last cleared BFD statistics and sessions on an IES site

78.15.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, click on the site on which you want to view the last cleared BFD statistics and sessions.

4

Click on the Clear Status tab and view the five last cleared BFD statistics and the five last cleared BFD sessions.

5

Close the forms.

END OF STEPS

78.16 To configure a subscriber interface on an IES

78.16.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Sites→Subscriber Interfaces, right-click on the Subscriber Interfaces object and choose Create IES Subscriber Interface or right-click on a subscriber interface and choose Properties. The IES Subscriber Interface (Create|Edit) form opens

4

Configure the required general parameters.

The Export Host Routes parameter is only available on retail subscriber interfaces when the Allow Unmatched Subnets parameter is enabled, or when the Interface Class parameter is set to Unnumbered.

The Tunnel Fault Notification parameter is configurable on interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept, in order to receive the fault notification from the tunnel facility MEP.

Note: The Name value for a subscriber interface must be unique in the NE.
You must configure the Default Primary DNS Server Address parameter before you can configure the Default Secondary DNS Server Address parameter.
The WAN Mode parameter is configurable only on subscriber interface creation. The parameter is read-only on existing interfaces. The WAN Mode parameter is validated on retail/wholesale subscriber interfaces, in which case the parameter value should be equal across interfaces.

5

Select a forwarding service in the Forwarding Service panel.

6

Select a forwarding subscriber interface in the Forwarding Subscriber Interface panel.

7

In the WLAN GW panel, set the Redundancy/Pool Manager parameter to Enabled if you need to configure WLAN GW redundancy or an IPv6 pool manager. The WLAN GW Redundancy and WLAN GW IPv6 Pool Manager tabs appear.

8

Click on the WLAN GW Redundancy tab to configure WLAN GW redundancy.

9

Click on the WLAN GW IPv6 Pool Manager tab to configure DHCPv6 functionality for the subscriber interface.

In order to support WLAN GW IPv6 address pools, the subscriber interface must have the IPv6 and WLAN Redundancy/Pool Manager parameters enabled on the General configuration tab, and must be configured with a soft GRE group interface.

1. Select a WLAN GW group.
2. Configure the required parameters.

   The DHCPv6 Lease Query parameter must be enabled before the DHCPv6 Lease Query Max Retry parameter can be configured.

   The SLAAC and IA-NA Administrative State parameters must be set to Out Of Service before their related parameters can be configured.

10

To configure IPv6 forwarding on the subscriber interface:

1. Configure the IPv6 Allowed parameter.
2. Configure the required IPv6 parameters.

   You must remove the check mark from the Default check box to access the Default Primary IPv6 DNS Server Address and Default Secondary IPv6 DNS Server Address parameters.
3. Click on the IPv6 Subscriber Prefixes tab.
4. Choose a subscriber prefix and click Properties, or click Create to create a new subscriber prefix. The Subscriber Prefix (Create|Edit) form opens.
5. Configure the required parameters.

11

To configure the interface class parameters:

1. Configure the Class parameter in the Interface Class panel.
2. If you selected a value of Unnumbered for the Class parameter, configure the Unnumbered Type parameter.
3. Configure the Interface Name parameter or the IP Address parameter.

12

Configure the DHCP server synchronization parameters in the DHCP Server Synchronization panel.

13

To create one or more IP addresses for the subscriber interface that are inherited by the SAPs in the group interfaces that are child objects of the subscriber interface:

1. Click on the Addresses tab.
2. Click Create. The IP Address (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

14

Click on the DHCP tab and configure the required parameters.

15

If the subscriber interface is part of a wholesale/retail configuration, click on the IPoE Linking tab and configure the required parameters.

The IPoE Linking tab is only configurable if a forwarding service is specified for the subscriber interface (see Step 5).

16

Click on the Server tab and configure the required parameters.

The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.

17

Click on the Client Applications tab and configure the Client Applications parameter. You can enable either or both of the PPPoE or DHCP choices.

18

Click on the PPPoE tab and configure the required parameters.

19

Save the changes and close the forms.

END OF STEPS

78.17 To add an AA interface to an IES or a VPRN site

78.17.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required IES or VPRN service and click Properties. The IES or VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Sites or Sites→Routing Instance, right-click on an AA Interface and choose Create IES or VPRN AA Interface. The IES/VPRN AA Interface (Create) form opens

4

Configure the required parameters on the General tab.

5

Click on the Port tab and click Select to choose a terminating port.

6

Save the changes. The IES/VPRN AA Interface (Create) form refreshes with addition tabs.

7

As required, click on the QoS tab to select an ingress policy or egress policy for the AA interface.

8

As required, click on the ACL tab to select an ingress filter for the AA interface.

9

Click on the Addresses tab and click Create. The IP Address, Service forms opens.

10

Configure the required parameters.

11

Save the changes and close the forms.

END OF STEPS

78.18 To add an AARP interface to an IES or a VPRN site

78.18.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required IES or VPRN service and click Properties. The IES or VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Sites or Sites→Routing Instance, right-click on an AARP Interface and choose Create AARP Interface. The AARP Interface (Create) form opens.

4

Configure the required parameters.

5

Save the changes and close the forms.

END OF STEPS

78.19 To configure a group interface on an IES

78.19.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears at the bottom of the Manage Services form.

3

Choose an IES and click Properties. The IES Service (Edit) form opens.

4

On the service navigation tree, expand Sites→site→Subscriber Interfaces, right-click on a subscriber interface and choose Create IES Group Interface, or expand the subscriber interface, right-click on a group interface and choose Properties. The IES Group Interface (Create|Edit) form opens.

5

Configure the required parameters.

The LNS parameter defines the type of group interface. This parameter is set at creation time and cannot be modified. Regular group interfaces cannot configure LNS attributes, and an LNS group interface does not allow PPPoE configuration or SAPs.

The Unicast RPF configuration is intended to enhance security for a MAC-only anti-spoofing configuration on the WLAN GW for soft GRE group interfaces; see 78.23 “To configure a WLAN GW on an IES group interface” (p. 2462) .

6

To configure cflowd sampling:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click Apply. The Cflowd sampling object appears in the Cflowd Sampling panel.
3. To associate a service template with the Cflowd sampling object, select the newly created Cflowd sampling object and click Properties. The Cflowd Sampling (Edit) form opens.
4. Click on the Template tab and select an associated template.
5. Save your changes.

7

Select a redundant interface.

8

Select a Diameter Application Policy.

9

Select a diameter authentication policy.

10

Configure the required parameters in the IPv4/IPv6 panel. To configure IPv6:

1. Enable the IPv6 Allowed parameter and configure the required parameters.
2. Click on the IPv6 Advertisement tab and configure the required parameters.

   If the No Default Router check box is enabled, the Router Lifetime parameter is not configurable.
3. Configure the parameters in the Prefix Options panel.

   If the either of the Infinite check boxes is enabled, the associated parameter is set to its maximum value.
4. Configure the parameters in the DNS Options panel.

   If the Infinite check box is enabled, the RDNSS Lifetime parameter is set to -1.

11

To configure IPv6 router solicitation on the group interface:

1. Ensure the IPv6 Allowed parameter is enabled on the General tab.
2. Click on the IPv6 Router Solicit tab.
3. Configure the Administrative State parameter.
4. Select a local user database.
5. Configure the parameters in the Minimum Authentication Time panel.
6. Configure the parameters in the Inactivity Timer panel.

   If you enable the Infinite check box, the Inactivity Timer parameters are disabled.

12

To configure DHCPv6 on the group interface:

1. Ensure the IPv6 Allowed parameter is enabled on the General tab.
2. Click on the DHCPv6 tab and configure the parameters.

   The Interface ID String parameter is configurable only if the Interface ID Option parameter is set to String.

   If the PD Managed Route parameter is enabled, the PD Managed Route Next Hop parameter is configurable.
3. Select a local user database.
4. Select a python policy.

5. Select a DHCPv6 filter.
6. Click on the Proxy Server tab and configure the parameters.

   The ID Type and Value parameters are configurable only if the Server ID Type parameter is set to Enterprise.
7. Click on the Relay tab and configure the parameters.
8. Configure the Server 1 through Server 8 parameters.
9. Select the interface name for each DHCPv6 server in the Zone Index panel.

   If you have entered a Unicast address, then the Interface Name parameter is not required.

13

To configure IPv6 router solicitation on the group interface:

1. Ensure the IPv6 Allowed parameter is enabled on the General tab.
2. Click on the IPv6 Router Solicit tab and configure the Administrative State parameter.
3. Select a local user database.
4. Select a python policy.
5. Configure the parameters in the Minimum Authentication Time panel.
6. Configure the parameters in the Inactivity Timer panel.

14

To configure WPP on the group interface:

1. Ensure the Enable WPP parameter is enabled on the General tab.
2. Click on the WPP tab.
3. Configure the Administrative State parameter.
4. Select a WPP site in the WPP Site panel.
5. Configure the Portal Name parameter.
6. Select an Initial Subscriber Profile.
7. Select an Initial SLA Profile.
8. Select an Initial Application Profile.
9. Configure the Restore Default Profile On Disconnection parameter.
10. Select a Local User DB.
11. Configure the Enable Triggered Hosts parameter.
12. Disable the Default check box and configure the Lease Time (Days, Hours, Minutes, and Seconds) parameters, if required.

15

To configure anti-spoofing:

1. Click on the Anti-Spoofing tab.

2. Configure the ARP Populate parameter.

16

To configure subscriber management:

1. Click on the Subscriber Management tab.
2. Configure the required parameters.

17

To configure ICMP for the group interface:

1. Click on the ICMP tab.
2. Configure the required parameters.

   The parameters on the SHCV panel are configurable when the SHCV Enabled parameter is configured.
3. On the Subscriber Host Connectivity Verification panel, select an IPv4/IPv6 policy, IPv4 policy, and IPv6 policy, as required.

18

To configure ARP for the group interface:

1. Click on the ARP tab.
2. Configure the Timeout (seconds) parameter.
3. Click on the Proxy ARP tab.
4. Configure the required parameters.

19

To configure DHCPv4 relay for the group interface:

1. Click on the DHCP tab.
2. Configure the required general parameters.

   When the Relay Unicast Message parameter is set to Renew or Release Update Source IP, the GI address parameter can be configured as any local configured address in the same routing instance as the GI address for the DHCP relayed messages. If the Relay Unicast Message parameter is set to None, the GI address is restricted to the IP address configured on the subscriber interface.
3. Select a subscriber authentication policy in the Subscriber Authentication panel.
4. Select a local user database in the Local User Database panel.
5. Select a filter policy in the DHCP Filter panel.
6. Click on the Server tab and configure the required parameters.

The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.

7. Click on the Client Applications tab and configure the Client Applications parameter.

20

To configure IPoE sessions for the group interface:

1. Click on the IPoE Session tab.
2. Configure the required parameters.

   The Stateless Redundancy parameter is configurable only on a regular (not soft GRE) group interface.

   Certain conditions and restrictions apply to an IPoE session on a soft GRE group interface:
   - RADIUS Session Timeout is set to Backwards Compatible by default
   - The IPoE session policy is set to a default policy (with SAP MAC address as key); cannot be changed
   - SAP Session Limit is set to 131071; cannot be changed
   - Administrative Sate is set to Enabled; cannot be changed.
3. Select an IPoE session policy.
4. Select a local user database.

21

To configure neighbor discovery on the group interface:

Neighbor discovery is configurable when the IPv6 Allowed parameter is enabled on the group interface.

1. Click on the Neighbor Discovery tab.
2. Configure the Maximum Number of Neighbors Learned and Duplicate Address Detection parameters.

22

To configure PPP for the group interface:

1. Click on the PPP tab.
2. Configure the Description and Administrative State parameters in the PPPoE panel.
3. Select a PPPoE policy in the PPPoE panel.
4. Select a PPPoE Local User DB in the PPPoE panel.
5. Select a python policy in the PPPoE panel.
6. Configure the required parameters in the PPPoE panel.
7. Configure the Description and Administrative State parameters in the PPP panel.
8. Select a PPP policy in the PPP panel.
9. Select a PPP Local User DB in the PPP panel.

10. Configure the Session Limit parameter in the PPP panel.

23

To configure IPoE linkage for the group interface:

1. Click on the IPoE Linkage tab.
2. Configure the required parameters.

24

To configure the ARP host for the group interface:

1. Click on the ARP Host Configuration tab.
2. Configure the required parameters.

25

To configure local address assignment for the group interface:

1. Click on the Local Address Assignment tab.
2. Configure the required parameters.
3. In the IPv4 Options panel, configure the PPP-v4 parameter and select a local DHCP server.
4. In the IPv6 Options panel, configure the IPoE SLAAC, IPoE WAN, and PPP SLAAC parameters and select a local DHCPv6 server.

26

To configure GTP for the group interface:

1. Ensure the Specific Type parameter is set to GTP on the General tab.
2. Click on the GTP tab.
3. Configure the Administrative State parameter.
4. Select a forward path extension.

27

To configure BRG for the group interface:

1. Click on the BRG Configuration tab.
2. Configure the Administrative State parameter.
3. Select a default BRG profile.
4. Enable the Authenticate BRG Only parameter, if only authenticated BRGs are permitted on the group interface.

28

After you create an LNS group interface, you must configure the L2TP tunnel group profile or tunnel profile to terminate sessions for the LNS group interface that you just created; see 28.89 “To configure L2TP on a routing instance” (p. 989) . You can also configure the termination of sessions on a group interface using a RADIUS server.

To configure LNS for the group interface:

1. Ensure that the LNS parameter is enabled on the General tab.
2. Click on the LNS tab.
3. Configure the Description parameter.
4. Select a subscriber profile in the Default Subscriber Profile panel.
5. Select an SLA profile in the Default SLA Profile panel.
6. Select a subscriber identification policy in the Subscriber Identification Policy panel.
7. Select an application profile in the Default Application Profile panel.
8. Configure the Default Subscriber Identification String parameter.

29

To configure SRRP for the group interface:

1. Click on the SRRP tab.
2. Configure the parameters in the Routing panel.
3. Click Create in the SRRP Instance panel. The SRRP Instance (Create) form opens.
4. Configure the required parameters in the SRRP Instance Information panel.
5. Select an operational group in the Operational Group panel.
6. Configure the Priority Step parameter.
7. Click on the Behavior tab.
8. Configure the required general parameters.
9. Select the SAP you need to use for the in-band messaging between the sites in the Message Path panel.
10. Select a policy pointer in the Policy Pointer 1 and Policy Pointer 2 panels.
11. Save the changes and close the form.

30

Save the changes and close the forms.

END OF STEPS

78.20 To create a bonding group interface on an IES

78.20.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces, right-click on a subscriber interface and choose Create IES Group Interface. The IES Group Interface (Create) form opens.

4

Configure the Name and Description parameters.

5

Select a redundant interface.

6

Select a diameter application policy.

7

Select a diameter authentication policy.

8

To configure bonding for the group interface:

1. Click on the Bonding tab.
2. On the General tab, configure the required parameters.
3. Select an FPE with subscriber management extensions enabled.
4. Click on the Connections tab.
5. Click Create. The Bonding Connection Object (Create) form opens.
6. Configure the required parameters, setting the Connection ID parameter to 1.
7. Save you changes and close the form.
8. Create a second bonding connection object, setting the Connection ID parameter to 2.

9

On the IES Group Interface, General tab set the Specific Type parameter to Bonding.

10

Save your changes and close the forms.

END OF STEPS

78.21 To configure a SAP on an IES group interface

78.21.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service navigation tree, expand Sites→site→Subscriber Interfaces→*subscriber_interface*, right-click on a group interface and choose Properties. The IES Group Interface (Edit) form opens.

4

Click on the Service Access Points tab.

5

Click Create. The IES Service Access Point (Create) form opens.

6

Configure the required general parameters.

7

Select a host lockout policy in the Host Lockout panel.

8

Configure a port:

1. Click on the Port tab.
2. Select a port for the L3 access interface from the Select Terminating Port form.

   The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.
3. Configure the required parameters.

   The Auto-Assign ID parameter is configurable if the port uses dot1q encapsulation. When

the parameter is enabled, the NFM-P automatically configures the Outer Encapsulation Value parameter using the lowest unassigned value.

You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter on the User Preferences form. Choose Application→ User Preferences from the main menu.

The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with QinQ encapsulation.

The Outer Encapsulation Value (VPI) and Inner Encapsulation Value (VCI) parameters are configurable only for ATM ports.

9

Items such as policies, schedulers, and filters can be applied later to multiple service components at once. Choose and right-click the components in the service navigation tree, choose Properties, and configure the parameters on the appropriate tab.

To assign ingress and egress QoS policies to the SAP:

1. Click on the QoS tab.
2. Configure the required parameters.

   The Ingress Match QinQ Dot1P and Egress Mark QinQ Top Bits Only parameters are configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.
3. Select a policy in the Ingress Policy and Egress policy panels.

   If you select an ingress or egress policy which has a forwarding class mapped to an ingress or egress queue group, you must ensure that the port you selected in Step 8 has the access ingress or egress queue group with the same name created on it.

   See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.
4. Select an HS secondary shaper in the HS Overrides panel, if required.

10

To specify that an aggregation scheduler policy is not applied to the interface:

1. Set the Aggregation parameter to Off.

   The Aggregation parameter is not configurable if the port you selected in Step 8 is an HSMDA port.
2. Configure the required parameters.

   The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

   The Frame-Based Accounting parameter is not configurable if the port you selected is an HSMDA port.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.

3. Select an ingress scheduler in the Ingress Scheduler panel.
4. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
5. If the port you selected in Step 8 is an HSMDA port, save the changes and close the forms.
6. Select an egress scheduler in the Egress Scheduler panel.
7. Select an egress policer control policy in the Egress Policer Control Policy panel.

11

To specify that an aggregation scheduler policy is applied to the interface:

1. Set the Aggregation parameter to On.

   You cannot specify an access scheduler policy if the port you selected in Step 8 is an HSMDA port.
2. Select an aggregation scheduler in the Aggregation Scheduler panel.

12

To assign ingress and egress ACL filters to the SAP:

1. Click on the ACL tab.
2. Select an ingress ACL filter in the Ingress Filter panel.
3. Select an egress ACL filter in the Egress Filter panel.

13

To assign an accounting policy to the SAP:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy.

14

To assign a virtual port to the SAP:

1. Click on the Virtual Port Name tab.
2. Configure the required parameters.

15

To configure anti-spoofing for the SAP, click on the Anti-Spoofing tab. Configure the Anti-Spoofing parameter and static hosts, as described in 74.24 “To create a static host for residential subscriber management on a SAP” (p. 2045) .

16

Note: A default DoS protection policy is automatically assigned to the SAP.

To assign a DoS protection policy or DDoS protection policy to the SAP:

1. Click on the Security tab.
2. Select a DoS protection policy in the NE DoS Protection panel.
3. Select an NE DDoS protection policy.
4. Configure the required parameters.

17

To define the levels of the ETH-CFM PDUs that are discarded on ingress into the SAP or to configure LMM or ETH-LMM frame loss measurement statistics collection:

1. Click on the OAM tab, then on the ETH-CFM tab.
2. Configure the Tunnel Fault Notification parameter.
3. In the MEPs panel, click Create. The MEP (Create) form opens.
4. Configure the required parameters.
5. Save the changes and close the form.
6. Configure the parameters in the Squelch Ingress Level panel.

   Levels must be assigned contiguously from Level 0. If you select a level greater than 0, then all levels lower than the one you chose will automatically be selected.
7. Configure the required parameters in the LMM Session Stats Collection panel.

18

To configure residential subscriber management on the SAP:

1. Click on the Subscriber Management tab, then on the IGMP Host Tracking tab.
2. Select the import policy used to filter IGMP packets.
3. Configure the required parameters.
4. You can click on the Host Tracking Info tab to view a list of hosts that are being tracked on this service access point.
5. Click on the Profiles tab.
6. Configure the required general parameters.
7. Select a default subscriber profile, subscriber authentication policy, and default application profile in the Policies panel.
8. Configure the Profiled Traffic Only parameter in the Single Subscriber Configuration panel.
9. Select a non-subscriber subscriber profile for the SAP in the Non-Subscriber Traffic Subscriber Profile panel.
10. Select a non-subscriber traffic SLA profile for the SAP in the Non-Subscriber Traffic SLA Profile panel.

11. Select a non-subscriber traffic application profile for the SAP in the Non-Subscriber Traffic Application Profile panel.
12. To view active hosts for the subscriber instance, click on the Subscriber Hosts tab.

19

To configure a default host in cases where the associated DHCP filter policy is configured with the Bypass Host Creation action:

1. Click on the Default Hosts tab.
2. Depending on the required addressing format, click on either the IPv4 or IPv6 tab.
3. Click Create. The (IPv4|IPv6) Default Host (Create) form opens.
4. In the Default Host Configuration panel, Select an (IPv4|IPv6) address to associate with the subscriber interface SAP.

   The address must be one of the addresses on the subscriber interface that the SAP belongs to, and not another default host.
5. Configure the Next Hop (IPv4|IPv6) Address parameter.

   The next hop address can be duplicated on the same SAP, but not on another SAP default host on the same routing instance.
6. Save your changes and close the form.

   Default host configuration settings cannot be changed once the default host is created.

20

To assign an ANCP policy to the SAP:

1. Click on the ANCP Static Map tab.
2. Click Create. The ANCP Static Map (Create) form opens.
3. Configure the ANCP String parameter.
4. Select an ANCP policy in the ANCP Policy panel.
5. Save the changes and close the forms.

21

To configure ATM on the SAP:

1. Click on the ATM tab.
2. Configure the required parameters.
3. Select an ingress ATM policy in the Ingress ATM Policy panel.
4. Select an egress ATM policy in the Egress ATM Policy panel.

22

Save the changes and close the forms.

END OF STEPS

78.22 To configure LAG per-link hashing on an IES group interface SAP

78.22.1 Purpose

You can configure weighted per-link hashing on a SAP on an IES group interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member.

78.22.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES (Edit) form opens.

3

On the service tree, expand Site→Subscriber Interfaces→*Interface*→*Group Interface*→*SAP*.

4

Right-click on the SAP you want to modify and choose Properties. The IES Service Access Point (Edit) form opens.

5

Click on the LAG Per Link Hash tab.

6

Configure the Class and Weight parameters.

7

Save your changes and close the forms.

END OF STEPS

78.23 To configure a WLAN GW on an IES group interface

78.23.1 Purpose

Before you can configure a WLAN GW for an IES group interface, you must configure an IES group interface with the Specific Type parameter set to Soft GRE. See 78.19 “To configure a group interface on an IES” (p. 2449) .

78.23.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES (Edit) form opens.

3

On the service tree, expand Sites→*routing_instance*→Subscriber Interfaces→*subscriber_interface*.

4

Right-click on a soft GRE group interface. The IES Group Interface (Edit) form opens.

5

Click on the WLAN GW tab.

6

Configure the required parameters.

7

Select a WLAN GW group.

8

Select a tunnel router instance.

9

Select a default retail IES.

10

Configure egress QoS.

1. In the Egress QoS panel, configure the required parameters.

   Do not configure the Aggregate Rate Limit parameter with a positive value if you specify a scheduler policy in 3 .
2. Select a QoS policy.
3. Select a Scheduler Policy.

11

If the WLAN GW is part of the home LAN extension configuration, configure the Maximum Number of Bridged Domains parameter on the LAN Extension panel.

This parameter can only be configured when the administrative state of the WLAN GW is Out of Service.

12

Select an authentication policy.

13

Configure an authentication hold time.

1. Enable the check box in the Authentication Hold Time panel.
2. Specify an authentication hold time in the Hours, Minutes, and Seconds fields.

14

Click on the SAP Parameters tab to configure SAP parameters.

1. Configure the required parameters.

   If the Default ID Type parameter is set to Use String, configure the Default ID String parameter.

   If the WLAN GW is supporting managed routes, set the Anti-Spoofing parameter to Next Hop IP and MAC Address. To enhance security for MAC-only anti-spoofing, configure unicast RPF on the group interface; see 78.19 “To configure a group interface on an IES” (p. 2449) .
2. Select a subscriber profile.
3. Select an SLA profile.
4. Select an application profile.
5. Select a subscriber identification policy.

15

Click on the Gateway Addresses tab to configure gateway addresses for the WLAN GW.

1. Click Create or choose a gateway address entry and click Properties. The Gateway Address (Create|Edit) form opens.
2. Type a gateway IP address in the field and click OK.

   You can configure up to 10 gateway addresses (six IPv6 and four IPv4).

   One IPv6 address can be configured with the XConnect option.

16

Click on the VLAN Tag Ranges tab to configure VLAN tag ranges.

1. Click Create or choose a VLAN tag range entry and click Properties. The VLAN Tag Range (Create|Edit) form opens.
2. Configure the required parameters.
3. Select a retail service.
4. Select a NAT policy.
5. Select an HTTP redirect policy.
6. Select an authentication policy.
7. Select a RADIUS proxy cache server.
8. If the VLAN tag range is part of an L2 wholesale-retail configuration, select a VPLS site in the L2 Retail panel, and configure the Administrative State and Description parameters as required.
9. Click on the DHCP tab and configure DHCP parameters, as required.
10. Click on the Distributed Subscriber Management tab to configure distributed subscriber management.
11. Configure the required parameters.
12. Select an accounting policy.
13. Select an ingress policer policy.
14. Select an egress policer policy.
15. Select an IP filter policy.
16. Select an application profile string, if required.
17. Click on the SLAAC tab to configure lifetimes for SLAAC-configured IPv6 hosts.

    The Administrative State parameter must be Up to configure parameters on the SLAAC tab.
18. Disable the Default check box and configure the lifetime parameters (in hours, minutes, and seconds) for any of the following categories:
    - Preferred Lifetime (Initial State)
    - Preferred Lifetime (Active State)
    - Valid Lifetime (Initial State)
    - Valid Lifetime (Active State)
19. Click on the DHCP6 tab to configure lifetimes for DHCPv6-configured IPv6 hosts.

    The Administrative State parameter must be Up to configure parameters on the DHCP6 tab.
20. Disable the Default check box and configure the lifetime parameters (in hours, minutes, and seconds) for any of the following categories:
    - Preferred Lifetime (Initial State)
    - Preferred Lifetime (Active State)
    - Valid Lifetime (Initial State)
    - Valid Lifetime (Active State)

21. If the VLAN tag range is part of a home LAN extension configuration, click on the Cross Connect tab.

    Configure the required parameters and select an ISA RADIUS policy.

    You must enable Authentication on DHCP on the VLAN Tag Range - General tab in order to enable the Cross Connect configuration.
22. If the VLAN tag range is part of a home LAN extension configuration, click on the Home LAN Extension tab.

    Configure the required parameters.

    You must enable Authentication on DHCP on the VLAN Tag Range - General tab, and enable the Administrative State of the BRG on the VLAN Tag Range - BRG Configuration tab in order to enable the home LAN extension configuration.
23. Save your changes and close the form.

17

Click on the L2 AP tab to configure an L2 access point.

1. Click Create or select an existing L2 AP entry and click Properties. The WLAN GW L2 Access Point (Create|Edit) form opens.
2. Select a port.
3. Configure the required parameters.

18

Save your changes and close the forms.

END OF STEPS

78.24 To configure a PIM interface on an IES group interface

78.24.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES (Edit) form opens.

3

On the service tree, expand Sites→*routing_instance*→Subscriber Interfaces→*subscriber_interface*.

4

Right-click on a group interface. The IES Group Interface (Edit) form opens.

5

Click on the PIM Interfaces tab.

6

Click on the Add button. The PIM Interface (Create) form opens.

7

Configure the required parameters.

8

Click on the Behavior tab.

1. Configure the required parameters.
2. To change the IPv4 or IPv6 administrative state for the PIM interface, click on the IPv4 and IPv6 tabs and configure the Administrative State parameter, as required.

9

Click on the Multicast CAC tab to configure multicast CAC on the PIM interface.

1. Select a multicast CAC policy.
2. Configure the Unconstrained Bandwidth and Mandatory Bandwidth parameters.
3. Click on the Levels tab.
4. Click Create. The PIM Interface Multicast CAC level form opens.
5. Configure the Level ID and Bandwidth parameters.
6. Save your changes and close the form.
7. Click on the LAG Port Down tab.
8. Click Create. The PIM Interface Multicast CAC LAG Port Down form opens.
9. Configure the Number of Ports Down and Level parameters.

10

Save your changes and close the forms.

END OF STEPS

78.25 To add a TMS interface to an IES

78.25.1 Purpose

The 7750 SR-7 and 7750 SR-12 devices support the configuration of a TMS interface in an IES.

The following conditions must be met before you can configure a TMS interface.

- You must have one IES configured.The TMS interface is configured on this on-ramp IES.

- You must have one VPRN configured for the off-ramp VPRN.
- Optionally, you can have one VPRN configured for the management VPRN.
- Each service must be associated within the same site on the same NE.
- An XP-IOM-3 is required.
- The ISA-TMS daughter card must be installed.

78.25.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service Subscriber (Edit) form opens.

3

On the service navigation tree, expand Sites→site, right-click on TMS Interfaces and choose Create TMS Interface. The TMS Interface (Create) form opens.

4

Configure the required parameters.

The Name value for a TMS interface must be unique in the context of the IES Site. That is, there cannot be another L3 access interface, subscriber interface, group interface or TMS interface with the same name on the same IES site.

5

To assign an off-ramp VPRN to the TMS interface, select a VPRN in the Off-Ramp VPRN panel.

6

Set the off-ramp ingress routing context. Configure the Routing Instance parameter.

7

To assign a management VPRN to the TMS interface, select a VPRN in the Management VPRN panel.

8

To associate an ISA-TMS daughter card with the TMS interface, select an ISA-TMS in the TMS Info panel.

Note: An ISA-TMS card can only be assigned to one TMS interface.

9

Configure the ISA-TMS Authentication parameter.

10

To assign an IP address to the TMS interface:

1. Click on the Addresses tab.
2. Click Create. The IP Address (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

11

Save the changes and close the forms.

END OF STEPS

78.26 To implement dual homing using SRRP

78.26.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Create the redundant interface used for SRRP out-of-band messaging between the two NEs:

Note: The two sites that participate in the dual homing configuration do not have to be part of the same service.
Ensure that the pair of sites each contains a properly configured subscriber interface and SAPs underneath the group interface that are participating in the redundant configuration.

Ensure that all subscriber interface IP addresses have a gateway address configured on them.

1. On the service navigation tree, click on the site to which you need to add the redundant pair; expand the entries for that site.
2. Right-click on Redundant Interfaces for one site of the redundant pair, and choose Create Redundant Interface. The Redundant Interface (Create) form opens.
3. Configure the required parameters.
4. Click on the Addresses tab.
5. Click Create. The IP Address (Create) form opens.

6. To specify IP addresses for the redundant interface on the current and remote sites, configure the required parameters.

   The remote IP address must be on the same subnet as the redundant interface IP address of the current site. For example, if the IP address of the current site is 7.7.7.7, with a prefix length of 24, then the redundant interface IP address of the remote site must be 7.7.7.d, where d is a value from 0 to 255, excluding 7.

7. Save the changes and close the forms.

4

Create an SDP spoke binding between the current and remote sites. The Source Interface is the Redundant Interface you created in Step 3 and the Tunnel Termination Site is the remote site. The Return Tunnel must come from the remote site.

5

Assign the Redundant Interface to the Group Interface for the current site.

1. On the service navigation tree under the Subscriber Interface, right-click on the Group Interface and choose Properties. The IES Group Interface (Edit) form opens.
2. Select a redundant interface you created in Step 3 in the Redundant Interface panel.
3. Save the changes and close the form.

6

Create an SRRP Instance for the current site.

1. On the service navigation tree under Group Interfaces, right-click on the SRRP Instances item for the current site, and choose Create SRRP Instance. The SRRP Instance (Create) form opens.
2. Configure the required parameters.

   The SRRP ID value must be the same for the current and remote sites.

3. Click on the Behavior tab and configure the general parameters.
4. Select the SAP you need to use for the in-band messaging between the sites in the Port field.
5. Configure the Policy Pointers for the SRRP Instance.
6. Save the changes and close the form.

7

Click Turn Up to activate the SRRP instance.

8

Repeat Step 3 to Step 7 for the remote site.

Note: When you repeat Step 3 to Step 7 for the remote site, that site becomes the current site and the previously configured site is the remote site.

After the two sites have been properly set up, you can examine the SRRP peer associations at any time by right-clicking an SRRP Instance in the service navigation tree. This opens the SRRP Instance - Edit form, which contains a read-only field called SRRP Peer. The Site ID, Service ID, and Operational State of the associated peer appear in this field.
You can also examine the state of an SRRP Instance by checking the Operational Flags field. The flags indicate specific problems that might occur with the SRRP Instance, as follows:

- Duplicate Subscriber IF Address: one of the local subscriber IP addresses is the same as a subscriber IP address on the remote node.
- Redundant Interface Mismatch: the local SRRP instance and remote SRRP instance have mismatched redundant interfaces.
- SAP Mismatch: the local SRRP instance is backing a different set of SAPs than the peer.
- Subnet Mismatch: one of the subnets that SRRP is backing up does not have a match with the peer.
- Dual Master: both SRRP instances are master at the same time.
- SAP Tag Mismatch: the local SRRP instance is backing a set of SAPs with different remote and local tags.
- SRRP ID Mismatch: the peer has a different SRRP instance ID backing the same subnet.

9

Save the changes and close the forms.

END OF STEPS

78.27 To configure IGMP host tracking on an IES site

78.27.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Sites icon in the service navigation tree, right-click on a Site icon and choose Properties. The IES Site (Edit) form opens.

4

Click on the Multicast tab, then on the IGMP Host Tracking tab.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

78.28 To configure an L3 access interface on an IES site

78.28.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service navigation tree, expand Sites→site, right-click on L3 Access Interfaces and choose Create IES L3 Access Interface, or expand L3 Access Interfaces, right-click on an access interface and choose Properties. The IES L3 Access Interface (Create|Edit) form opens.

4

Configure the required parameters.

The Unnumbered Type parameter is configurable when the Class parameter is set to Unnumbered.

The Unnumbered IP Address parameter is configurable when the Unnumbered Type parameter is set to IP Address.

The Unnumbered Interface Name parameter is configurable when the Unnumbered Type parameter is set to Name.

For a 7705 SAR, if a port or LAG is associated with the L3 access interface, the loopback parameter cannot be enabled. If the loopback is enabled, a port or LAG cannot be associated with the L3 access interface.

The Admin Link Local Address parameters are only configurable when the IPv6 Allowed parameter is enabled.

The Enable Ingress FlowSpec IPv4 and Enable Ingress FlowSpec IPv6 parameters appear only when a port is configured on the interface.

The Tunnel Fault Notification parameter is configurable on interfaces where the device has

ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept, in order to receive the fault notification from the tunnel facility MEP.

5

Configure the Configured IP MTU (Octets) parameter in the Frame Size Constraints panel.

6

Configure the parameters in the Unicast RPF panel.

The URPF Check State IPv6 and URPF Check Mode IPv6 parameters are only configurable when the IPv6 Allowed parameter is enabled.

7

To associate a host lockout policy with the L3 access interface, select a policy in the Host Lockout panel.

8

Configure the PTP HW Assist parameter in the PTP HW panel.

You can only enable the PTP HW Assist parameter after the port is associated to the interface, and an IP address is configured on the interface.

9

To configure cflowd sampling:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click Apply. The Cflowd sampling object appears in the Cflowd Sampling panel.
3. Save your changes.

10

Select an operational group in the Operational Group panel.

11

Associate a port with the L3 access interface:

1. Click on the Port tab.

   If the Loopback Enabled parameter in Step 4 is enabled, you cannot associate a port with the L3 access interface.
2. Select a port in the Terminating Port panel.

   Only ports in access or hybrid mode, or PW ports are listed. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. The port is listed when you click Search.

The form only lists PW ports that are bound to a service tunnel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for information about creating a service tunnel with a PW port binding.

You can select a bundle in the Terminating Port panel.

3. Configure the required parameters.

   The Auto-Assign ID parameter is configurable if the port uses dot1q encapsulation. When the parameter is enabled, the NFM-P automatically configures the Outer Encapsulation Value parameter using the lowest unassigned value.

   You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter on the User Preferences form. Choose Application→ User Preferences from the main menu.

   The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with QinQ encapsulation.

   The Outer Encapsulation Value (VPI) and Inner Encapsulation Value (VCI) parameters are configurable only for ATM ports.

4. Configure the required parameters in the Properties panel.
5. Select an NE DoS protection policy in the Security panel and configure the required parameters.

12

Save the changes and close the forms.

END OF STEPS

78.29 To configure LAG per-link hashing on an IES L3 access interface

78.29.1 Purpose

You can configure weighted per-link hashing on an IES L3 access interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member.

78.29.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the LAG Per Link Hash tab.

5

Configure the Class and Weight parameters.

6

Save your changes and close the forms.

END OF STEPS

78.30 To bind an IES L3 access interface to a VPLS site or VPLS I-site

78.30.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Routed VPLS tab.

Note: You can create and manage a routed VPLS connector from the navigation tree on the Composite Service (Edit) form.

5

Enter a VPLS site name or select a VPLS site.

The site that you select must be enabled for RVPLS; see 77.33 “To configure a VPLS site” (p. 2294).

The operational state of the IP interface binding will not be turned up until the Enable IP Interface Binding parameter is set to true.

6

Expand the Ingress or Egress panels as needed.

7

To configure the filter policies:

1. Select an IPv4 filter in the IPv4 Filter panel.
2. Select an IPv6 filter in the IPv6 Filter panel.

8

To enable table-based color-aware ingress classification, select the Enable Table Classification For VPLS parameter, then select a 7210/7250 DSCP classification policy as the Routed Override QoS policy. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

9

Select an egress QoS policy in the Egress - QoS Policy panel.

10

Save your changes and close the forms.

END OF STEPS

78.31 To apply OSPF, RIP, or IS-IS to an IES L3 interface

Note: OSPF, RIP, or IS-IS must be enabled at the routing instance level before you can apply OSPF, RIP, or IS-IS to an L3 interface.

78.31.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

Click on the Interfaces tab, then on the L3 Access Interfaces tab.

4

Choose an L3 interface in the list and click Properties. The L3 Access Interface (Edit) form opens.

5

Click on the Protocols tab and click Create. The Create Interface form opens.

6

Specify the interface type by configuring the What type of interface would you like to create? parameter.

7

Click OK. The Interface (Create) form opens. See “Routing protocol configuration workflow and procedures” (p. 909) for information about configuring specific routing protocols.

END OF STEPS

78.32 To assign ingress and egress QoS policies to an IES L3 access interface

78.32.1 Before you begin

The available panels and parameters vary depending on the NE, chassis type, and release.

78.32.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the QoS tab and configure the Ingress Match QinQ Dot1P parameter.

The QoS tab is configurable only if a port is assigned to the interface.

The Ingress Match QinQ Dot1P parameter is configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.

5

Select an ingress QoS policy in the Ingress Policy panel.

Note: If you select an ingress policy which has a forwarding class mapped to an ingress queue group, you must ensure that the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) has the access ingress queue group with the same name created on it.

See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

6

Select an ingress queue group template policy in the Forwarding Plane Redirect panel.

7

Configure the Instance ID and the Egress Mark QinQ Top Bits Only parameters.

The Egress Mark QinQ Top Bits Only parameter is configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.

8

Configure the required parameters in the Aggregate Rate Limit panel.

9

Select a redirect list policy in the Ingress SAP Queue Group Redirect List panel.

See 50.76 “To configure a queue group redirect list policy” (p. 1624) for more information about how to configure a redirect list policy.

10

Select an egress policy in the Egress Policy panel.

Note: If you select an egress policy which has a forwarding class mapped to an egress queue group, you must ensure that the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) has the access egress queue group with the same name created on it.
See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.
Queue Group Template policies are not applicable to L3 interfaces associated with HSMDA ports.

11

Select an Egress Queue Group Template Policy in the Port Redirect panel.

Note: Selecting an Egress Queue Group Template Policy here permits the redirection of Ethernet traffic packets to a queue ID specified in the egress port queue group of the SAP. The following properties and restrictions apply:

- If an Egress Queue Group Template Policy is specified here, the policy must have port redirection enabled.
- You cannot use policy-based redirection with the queue group when the SAP has SAP-based redirection enabled.

- Port access egress redirection is only supported on Ethernet/LAG ports. It is not supported on SAPs bound on non-Ethernet, Eth-tunnel, or CCAG ports.
- Supported ports include access, hybrid, and HSMDA.
- Queue groups can be applied to SAPs that incorporate LAGs. The LAGs can include port members from just a single card or from multiple cards.
- If you edit a LAG incorporated by the SAP, you cannot remove the last LAG member if a queue group reference exists to the containing SAP.
- You cannot add a secondary LAG member that has a queue group mismatch with primary LAG member.

12

Select a redirect list policy in the Egress SAP Queue Group Redirect List panel.

See 50.76 “To configure a queue group redirect list policy” (p. 1624) for more information about how to configure a redirect list policy.

13

In the Shaper Group panel, select a Shaper Group for the access ingress port or access egress port.

14

In the IXR Specific panel, select an Egress Remark policy, Egress VLAN QoS policy, and shared policer, and configure all other parameters, as required. See 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634) and 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594).

15

Select an HS secondary shaper in the HS Overrides panel, if required.

16

Save the changes and close the forms.

END OF STEPS

78.33 To assign ingress and egress QoS policies to an IES L3 access interface on a 7210 SAS site

Note: The available parameters and policies vary depending on the device type and chassis variant. The configurations that are supported on the site NE are shown on the form.

78.33.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the QoS tab and expand the 7210/7250 Specific panel.

5

Select a SAP Access Ingress policy in the Ingress Policy panel.

When you assign an ingress policy, the Color Mode parameter setting for the meters in the policy must coincide with the Enable DEI parameter setting on the physical port. When Enable DEI is selected, the Color Mode for meters must be set to Color Aware. When Enable DEI is not selected, the Color Mode for meters must be set to Color Blind. See 16.24 “To configure Ethernet ports” (p. 599) and 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .

To support H-metering, you must choose an ingress policy with all meter rate modes set to trTCM (RFC 4115).

For 7250 IXR sites, the selected SAP Access Ingress policy must contain an assigned 7250 Ingress CoS policy.

6

Select an Egress policy.

7

To enable table-based color-aware ingress classification, select the Enable Table Classification parameter. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

8

Select an Egress Remarking policy in the Egress Remark Policy panel.

9

Configure the required parameters in the Aggregate Rate Limit panel.

You can configure the Ingress Meter parameter only during SAP creation. The parameter must be set to true to support H-metering.

You can configure the Ingress Meter Rate (kbps) and Ingress Meter Burst parameters only after SAP creation

You can configure the Egress Meter Rate and Egress Meter Burst parameters only when resources are allocated to the SAP Egress Aggregate Meter parameter in the system resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380). To allocate resources on the 7210 SAS-R, configure the Egress SAP Aggregate Meter parameter in the system resource profile policy assigned to the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).

You must also enable port-based scheduling on 7210 SAS-Mxp and 7210 SAS-R NEs; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384).

The Enable Egress Meter Stats parameter is available when a value is configured for the Egress Meter Rate parameter.

10

Save the changes and close the forms.

END OF STEPS

78.34 To configure scheduling on an IES L3 access interface

78.34.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Schedulers tab and configure the required parameters.

Note: The Schedulers tab is configurable only if a port is assigned to the SAP in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) .

5

To configure scheduling on a 7210 SAS site, enable and configure the required parameters in the Egress Aggregate Rate Limit panel and go to Step 9 .

6

To configure scheduling on a 7705 SAR site:

Note: For the 7705 SAR, scheduler behavior is determined by the scheduler mode, which is 4-Priority by default. You can only configure the Egress and Ingress Aggregate Rate Limit parameters when the Scheduler Mode parameter is set to 16-Priority. You can set the Scheduler Mode to 16-Priority only when the port is on an MDA that supports 16-Priority. See the 7705 SAR documentation for more information.

If you change the Scheduler Mode parameter from 16-Priority to 4-Priority, the NFM-P automatically restores the default settings for the Egress Aggregate Rate Limit and Ingress Aggregate Rate Limit panels when you click Apply or OK.

1. In the Egress Scheduler panel, configure the Scheduler Mode parameter.
2. In the Ingress Scheduler panel, configure the Scheduler Mode parameter.
3. If you set the Scheduler Mode parameter to 16-Priority in the Egress Scheduler panel or Ingress Scheduler panel, configure the parameters in the Egress Aggregate Rate Limit panel and Ingress Aggregate Rate Limit panel.

7

To specify that an aggregation scheduler policy is not applied to the interface:

1. Set the Aggregation parameter to Off.

   The Aggregation parameter is not configurable if the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) is an HSMDA port.
2. Configure the required parameters.

   The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

   The Frame-Based Accounting parameter is not configurable if the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) is an HSMDA port.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.
3. Select an ingress scheduler in the Ingress Scheduler panel.
4. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
5. If the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) is an HSMDA port, go to Step 8 .
6. Select an egress scheduler in the Egress Scheduler panel.
7. Select an egress policer control policy in the Egress Policer Control Policy panel.
8. Go to Step 9 .

8

Note: You cannot specify an access scheduler policy if the port you selected in 78.28 “To configure an L3 access interface on an IES site” (p. 2472) is an HSMDA port. Go to Step 9 .

To specify that an aggregation scheduler policy is applied to the interface:

1. Set the Aggregation parameter to On.
2. Select an aggregation scheduler in the Aggregation Scheduler panel.

9

Save the changes and close the forms.

END OF STEPS

78.35 To assign ingress and egress ACL filters to an IES L3 access interface

78.35.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the ACL tab.

Note: When you configure ACL filters on a 7210 SAS NE, you must configure the system resource profile appropriately. See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) for more information.

5

Select the required ACL filter policies.

Note: Select the Ingress IP Filter Pair option to configure IP and IPv6 filters simultaneously to filter IP and IPv6 traffic respectively. Once the filters are paired, they can

not be used separately or paired with another filter. This option is supported only for IES services on 7250 IXR NE, 22.2 R1 or later.

6

Save the changes and close the forms.

END OF STEPS

78.36 To assign a virtual port to an IES L3 access interface

78.36.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Virtual Port Name tab.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

78.37 To associate a local DHCPv4 or DHCPv6 server to an IES L3 access interface

78.37.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Local DHCP tab:

a. For local DHCPv4 servers, select a local DHCP server in the Local DHCP Server panel.

Note: You cannot associate a local DHCPv4 server to the L3 group Interface if the Administrative State parameter in the Local Proxy Service panel is up.

b. For local DHCPv6 servers, select a local DHCPv6 server in the Local DHCPv6 Server panel.

Note: To associate local DHCPv6 servers the IPv6 Allowed parameter must be enabled on the General tab.

5

Save the changes and close the forms.

END OF STEPS

78.38 To assign an accounting policy to an IES L3 access interface

78.38.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Sites icon in the service navigation tree.

4

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

5

Click on the Accounting tab.

6

Select an accounting policy.

7

Configure the required parameters.

The Collect Egress Queue Statistics parameter can be configured only during IES L3 access interface creation. See 78.28 “To configure an L3 access interface on an IES site” (p. 2472).

8

Save the changes and close the forms.

END OF STEPS

78.39 To assign an accounting template policy to an IES interface

78.39.1 Purpose

This procedure applies to IES L3 access interfaces, tunnel interfaces, or subscriber group interfaces.

For information about creating an accounting template policy, see 54.12 “To configure an accounting template policy” (p. 1757).

78.39.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Sites icon in the service navigation tree.

4

Expand the Routing Instance icon in the service navigation tree on which you need to configure the accounting template policy.

5

Right-click on the L3 access interface, tunnel interface, or subscriber group interface you need to configure and choose Properties.

6

Click on the Policies tab.

7

Select an accounting template policy in the Accounting Template panel.

8

Save and close the form.

END OF STEPS

78.40 To configure application assurance on an L3 access interface

78.40.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand Sites→Routing Instance in the service navigation tree and right-click on the L3 access interface and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Application Assurance tab.

5

To bind an application profile to the IES L3 access interface, select an application profile string.

Note: Only local profiles that already exist on the NE can be selected.

6

To choose one of the following AA transit policy types:

Note: You can only associate one AA transit policy type with a service object.
To bind a transit policy to an L3 access interface, a port must already exist on the interface.
You can bind a transit policy to only one L3 access interface or spoke SDP binding per NE.
The transit policy and the application profile must belong to the same application assurance group or partition.

a. Select an AA transit IP policy in the Transit IP Policy panel.

b. Select an AA transit prefix policy in the Transit Prefix Policy panel.

7

To associate an AA redundant protocol with the service object, select an AA redundant protocol in the AA Redundant Protocol panel.

8

Save your changes and close the forms.

END OF STEPS

78.41 To assign an ANCP policy to an IES L3 access interface

78.41.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Sites icon in the service navigation tree.

4

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

5

Click on the ANCP Static Map tab.

6

Click Create. The ANCP Static Map (Create) form opens.

7

Configure the ANCP String parameter.

8

Select an ANCP policy in the ANCP Policy panel.

9

Save the changes and close the forms.

End of steps

78.42 To associate a security zone policy with an IES L3 access interface on a 7705 SAR

78.42.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Zone tab.

5

Select a security zone policy.

You can create a security zone policy by clicking Create.

You can associate a security zone policy with an IES service that is configured with a tunnel port as a SAP.

6

Configure the ByPass Zone Config parameter.

7

Save the changes and close the forms.

End of steps

78.43 To assign a time of day suite to an IES L3 access interface

78.43.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the TOD Suite tab.

5

Select a time of day suite.

Note: SapEgrQosPlcyStats and SapIngQosPlcyStats statistics are collected only if a Time Of Day Suite is applied on the SAP.

6

Save the changes and close the forms.

END OF STEPS

78.44 To configure residential subscriber management for an IES L3 access interface

78.44.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Subscriber Management tab.

5

Enable the SHCV Enabled parameter.

6

Configure the required parameters on the Subscriber Host Connectivity Verification panel.

7

Select an SHCV IPv4 policy.

8

Save the changes and close the forms.

END OF STEPS

78.45 To assign a DoS protection policy or DDoS protection policy to an IES L3 access interface

Note: A default DoS protection policy is automatically assigned to the interface.
See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for more information about configuring and applying NE DoS protection policies.

78.45.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Security tab.

Note: If you are assigning a DoS protection policy for 7210 SAS-R6 or 7210 SAS-R12 nodes, click on the Port tab.

5

Select a DoS protection policy in the NE DoS Protection panel.

6

Select an NE DDoS protection policy.

7

Save the changes and close the forms.

END OF STEPS

78.46 To assign an IP address to an IES L3 access interface

78.46.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Addresses tab.

5

Click Create. The IP Address (Create) form opens.

6

Configure the required parameters.

The Broadcast Address Format parameter only appears if the IP Address parameter is set to an IPv4 address.

The parameters in the IPv6 panel only appear if the IP Address parameter is set to an IPv6 address.

7

Select a Track SRRP Instance, if required.

8

Save the changes and close the forms.

END OF STEPS

78.47 To configure BFD for an IES L3 access interface

78.47.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the BFD tab.

5

In the IPv4 BFD Configuration panel, set the Admin Status parameter to Up and configure the required parameters.

6

In the IPv6 BFD Configuration panel, set the Admin Status parameter to Up and configure the required parameters.

7

To view local and remote session peers that are managed by the NFM-P, click on the BFD Session tab. A list of BFD current sessions on a router interface or an L3 interface appears.

8

Click on a session. The properties form for the session opens. View the following:

- BFD status
- protocol used

- local address
- remote address
- operational status and statistics

9

Save the changes and close the forms.

END OF STEPS

78.48 To configure ICMPv4 for an IES L3 access interface

78.48.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose an IES service and click Properties. The IES Service (Edit) form opens.

4

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

5

Click on the ICMP tab and configure the required parameters.

6

Configure the required parameters in the Redirects panel.

7

Configure the required parameters in the Unreachables panel.

8

Configure the required parameters in the TTL Expired panel.

9

Save the changes and close the forms.

END OF STEPS

78.49 To configure ICMPv6 on an IES L3 access interface

78.49.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Ensure that the IPv6 Allowed parameter is enabled.

5

Click on the ICMPv6 tab.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

78.50 To assign an ICMP ping template to an IES L3 access interface

78.50.1 Before you begin

ICMP ping templates are used to populate values for ICMP ping tests when the tests are used on L3 interfaces to control the operational state. ICMP ping templates are supported for IPv4 only. See 90.51 “To configure an ICMP Ping template” (p. 3079) for information about configuring ICMP ping templates.

Note: Nokia recommends that you assign an NE DoS protection policy, configured with a protocol of ICMP-Ping-Check, to the affected interfaces; see 78.45 “To assign a DoS protection policy or DDoS protection policy to an IES L3 access interface” (p. 2491).

78.50.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you need to modify, and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the ICMP Ping Template tab.

5

Click Create, or choose the template in the list and click Properties. The Virtual Interface ICMP Template Config form opens.

6

Click Select to choose a template pointer for the template. The Select Template Pointer form opens.

7

Choose an ICMP Ping template from the list and click OK to close the Select Template Pointer form.

8

On the Virtual Interface ICMP Template Config form, configure the Admin State and Destination Address parameters for the ICMP Ping test.

9

Save the changes and close the forms.

END OF STEPS

78.51 To configure ARP for an IES L3 access interface

78.51.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the ARP tab and configure the required parameters.

5

To configure static ARP:

1. Click Create. The Static ARP (Create) form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

6

Click on the Proxy ARP tab and configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

78.52 To configure neighbor discovery on an IES L3 access interface

78.52.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Set the IPv6 Allowed parameter to enabled.

5

Click on the Neighbor Discovery tab and click Create. The Neighbor Discovery (Create) form opens.

6

Configure the required parameters.

7

Save the changes and close the form.

8

Click on the Proxy ND tab and configure the required parameters.

Note: Do not leave an empty policy parameter between two configured policy parameters. For example, do not configure the Policy 1 and Policy 3 parameters and leave the Policy 2 parameter unconfigured, or the NFM-P reorders the policies and moves the policy specified for the Policy 3 parameter to the Policy 2 parameter.

9

Click on the Secure ND tab and configure the required parameters.

10

Save the changes and close the forms.

END OF STEPS

78.53 To configure DHCPv4 for an IES L3 access interface

78.53.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the DHCP tab and configure the required parameters.

The Lease Populate parameter is configurable when the Enable parameter is enabled.

5

Select a subscriber authentication policy in the Subscriber Authentication Policy panel.

6

Configure the required parameters in the Option 82 panel.

The Remote ID String parameter is configurable when the Remote ID is set to Remote ID String.

When the Relay Unicast Message parameter is set to Renew or Release Update Source IP, the GI address parameter can be configured as any local configured address in the same routing instance as the GI address for the DHCP relayed messages. If the Relay Unicast Message parameter is set to None, the GI address is restricted to the IP address configured on the subscriber interface.

7

Configure the parameters in the Vendor Specific Option panel.

8

Configure the parameters in the GI-Address panel.

9

Select a Python policy, if required.

10

Click on the Server tab and configure the required parameters in the DHCP Servers and Local Proxy Server panels.

The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.

11

Depending on the type and release of device that you are configuring, the Subscriber Authentication Policy panel is present. Select a Subscriber Authentication policy.

12

Save the changes and close the forms.

END OF STEPS

78.54 To create a VRRP instance on an IES L3 access interface for a virtual router

78.54.1 Prerequisites

You must know the VRID for an existing virtual router and ensure that the interface is a member of the same subnet as the virtual router.

See Chapter 37, “VRRP” for configuration information about VRRP instances.

The following configurations are required for the operation of the IPv6 VRRP instance:

- You can only create an IPv6 VRRP Instance if you enable the IPv6 Allowed parameter on the General tab of the IES L3 Access Interface (Edit) form.
- The Link Local Address on the parent interface must be set to preferred or to disable DAD, depending on your node release, and configured as one of the backup addresses (or same subnet) for the IPv6 VRRP instance. The Admin Link Local Address related parameters on the General tab of the IES L3 Access Interface (Edit) form must be set accordingly.
- The IPv6 address on the parent interface must be set to preferred or disable DAD to be used as a backup address (on same subnet) for the IPv6 VRRP instance. The IPv6 parameters in Step 6 of 78.46 “To assign an IP address to an IES L3 access interface” (p. 2492) must be set accordingly.
- The Send Advertisement and Use Virtual MAC Address parameters must be enabled in Step 6 of 78.56 “To configure router advertisement on an IES L3 access interface” (p. 2502) for the router advertisement on the parent interface.

78.54.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the VRRP tab.

5

Click Create. The VRRP Instance (Create) form opens.

6

Configure the Virtual Router ID parameter.

7

Perform Step 3 to Step 15 of 37.4 “To create and configure a VRRP instance” (p. 1283) .

You can use the VR Instances tab to create, modify, and view VR instances.

8

Save the changes and close the forms.

END OF STEPS

78.55 To configure anti-spoofing filters for an IES L3 access interface

78.55.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Anti-Spoofing tab and configure the required parameters.

5

Configure the required general parameters.

The ARP Populate parameter is configurable when all of the IP addresses of the defined static hosts on the interface are in one of the subnets configured for the interface.

6

To configure static subscriber host entries, if subscriber entries are not available through DHCP lease management:

1. Click on the Static Hosts tab.

2. Click Create. The Access Interface Anti-Spoofing Static Host Display (Create) form opens.
3. Configure the required parameters.

   Specify at least one IP address or MAC address for each static host. The values specified for the Anti-Spoofing and ARP Populate parameters determine the type of address entry that is required for the static host. For example, when you set the Anti-Spoofing parameter to Source Ip Addr, you must specify at least the IP address for the static host.

   You can configure a static host on a SAP only when no static ARP entries exist on the IP interface.

   When the ARP Populate parameter is enabled, the IP address of the new static host must be in one of the subnets that is configured for the interface in Procedure 78.46 “To assign an IP address to an IES L3 access interface” (p. 2492) .

7

Save the changes and close the forms.

END OF STEPS

78.56 To configure router advertisement on an IES L3 access interface

78.56.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you want to modify and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Advertisement tab.

5

Click Create to add a router advertisement entry. The Router Advertisement (Create) form opens.

6

Configure the required parameters.

If you are configuring the L3 interface for an IPv6 VRRP instance, then the Send Advertisement and Use Virtual MAC Address parameters must both be enabled.

7

Click on the Prefix tab.

8

Click Create. The Router Advertisement Prefix (Create) form opens.

9

Configure the required parameters.

10

Save the changes and close the forms.

END OF STEPS

78.57 To specify QoS policy overrides on an IES L3 access interface

78.57.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Site icon in the service navigation tree on which you need to configure the L3 access interface, right-click on the L3 Access Interface icon and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Click on the Override Policy Items tab.

Note: The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.

- If you configured a non-HSMDA port, the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.
- If you configured an HSMDA port, the Access Ingress Queues, Access Egress HSMDA Queues, and Ingress Policer tabs are active.

Configure the policy overrides, as described in 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) .

To configure meter overrides on a 7210 SAS, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) . To configure queue overrides on a 7210 SAS, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

5

Save the changes and close the forms.

END OF STEPS

78.58 To configure DHCPv6 on an IES L3 access interface

78.58.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Site icon in the service navigation tree on which you need to configure the L3 access interface, right-click on the L3 Access Interface icon and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Ensure that the IPv6 Allowed parameter is enabled.

5

Click on the DHCPv6 tab and configure the required parameters on the DHCPv6 Relay — General tab.

6

Select a python policy, if required.

7

Select a local user database, if required.

8

Click on the Server tab.

9

Configure the Server 1 through Server 8 parameters.

10

Configure the interface name for each DHCPv6 server that you configured in Step 9 by selecting the an interface in the Zone Index panel.

11

Click on the DHCPv6-Prefix tab.

12

Click Create. The DhcpRelayV6PrefixDelegation (Create) form opens.

13

Configure the required parameters.

14

Save the changes and close the forms.

END OF STEPS

78.59 To associate a Multi-Chassis shunting profile to an IES L3 access interface

78.59.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES service and click Properties. The IES Service (Edit) form opens.

3

Expand the Site icon in the service navigation tree on which you need to configure the L3 access interface, right-click on the L3 Access Interface icon and choose Properties. The IES L3 Access Interface (Edit) form opens.

4

Select a profile in the Multi-Chassis Shunting Profile panel; see 27.31 “To configure a Multi-Chassis shunting profile on a base routing instance or VPRN routing instance” (p. 876) .

Note: Only the shunting profile created on a base routing instance can be selected.

5

Save the changes and close the forms.

END OF STEPS

78.60 To start or stop the ignore SAP port state tool on an IES interface

78.60.1 Purpose

Use the ignore SAP port state tool to bypass the checking of the physical port operational state if it is down during operational checks on the NE. The tool is available for 7x50 NEs on L3 and subscriber interfaces of IES and VPRN services. The interface can be IPV4/IPV6 or dual stack. An SAP must be attached.

If this procedure is performed when the IP interface is operationally up the command will be accepted but will enter a pending state. It will not become active unless the port state becomes non-operational.

78.60.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The IES Service (Edit) form opens.

3

Perform one of the following:

a. On the service navigation tree, expand Site→L3 Access Interfaces and right-click on an L3 access interface.

b. On the service navigation tree, expand Site→Subscriber Interfaces→Group Interface and right-click on a group interface.

4

Choose Start Ignore SAP Port State or Stop Ignore SAP Port State.

END OF STEPS

78.61 To view the service operational status

78.61.1 Purpose

The Aggregated Operational State and State Cause indicators on the General tab of the service management form display information about service faults.

78.61.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The IES Service (Edit) form opens.

3

View the Aggregated Service Site Operational State and State Cause indicators. When the Aggregated Operational State is Down or Partially Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The Object Alarms tab is displayed.

6

Click on the Aggregated Alarms tab to view the aggregated alarms for the object. The Aggregated Alarms tab is displayed.

7

Close the forms.

END OF STEPS

78.62 To view the service topology

78.62.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Topology View. A Topology View dialog box appears.

3

Click Yes to proceed. The Service Topology - map opens.

See Chapter 4, “Topology map management” for more information about service topology views.

END OF STEPS

78.63 To modify an IES

CAUTION

Service Disruption

Modifying parameters can be service-affecting.

78.63.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click Properties. The IES Service (Edit) form opens.

The following tabs list the service elements that can be individually or collectively selected and configured:

- General tab — displays general customer and service information
- Sites tab — lists the sites that are included in the service
- Interfaces tab — lists the interfaces, such as L3 access, that are included in the service, and allows the creation and configuration of subscriber, group, and redundant interfaces
- SDP Bindings tab — displays the spoke SDP bindings that are associated with the service
- Addresses tab — lists the IP addresses that are associated with the service

  You cannot remove an IP address from an interface when the IP address of a static host is defined in the subnet of the interface IP address and the ARP Populate parameter is enabled on the Anti-Spoofing tab.
- Faults tab — displays the faults associated with the service

3

Modify the parameters for the service, as required.

4

Save the changes and close the forms.

END OF STEPS

78.64 To modify an IES using the topology view

78.64.1 Purpose

The topology view for a service provides a graphical representation of the various components and their interconnections. You can also use this view to add, modify, or just navigate to service components. This provides an alternative approach to performing these functions from the navigation tree view.

Working from the topology view can expedite the creation of the components, since many of the fields you would ordinarily have to set in the configuration forms will be automatically populated using this approach. The configuration forms can also be accessed directly at any time from this view by right-clicking a component. This allows quick access to conduct more detailed component configuration.

78.64.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an IES and click on Topology View. The Service Topology map opens.

The remainder of this procedure contains two sub-procedures describing the components that can be created and modified from the topology view. These include:

- Creating a new site. Go to Step 3 .
- Creating spoke SDP bindings. Go to Step 8 .

Adding a new site

3

Right-click on any blank space in the service topology map. A contextual menu is displayed. Choose the Create IES Site option.

The Select Network Elements form appears.

4

Select one or more sites to add to the service and click OK. The IES Site (Create) form for the new site is displayed. If you selected more than one site, the IES Site (Multiple Instances) (Create) form for the new sites is displayed.

5

Click OK. The IES Site (Create) (or IES Site (Multiple Instances) (Create)) form closes and the new site (or sites) is displayed on the map.

6

If you need to perform detailed configuration of site properties for the new site, right-click the site icon and choose Properties. The Site (Edit) form opens. See 78.5 “To configure an IES site” (p. 2431) for detailed site configuration information.

7

Go to Step 8 if you need to create spoke SDP bindings or go to Step 15 to finish.

Creating spoke SDP bindings

8

Select the sites you need to connect in the service topology map and right-click on any one of them. A contextual menu is displayed.

Note: When you create a spoke binding between two sites, the order in which you select them is important. The first site you select will become the source site and the second site will become the destination site. Therefore, it is not recommended that you do a marquee-select in the topology view, since you will not be sure of this hierarchy. Instead, select the sites individually, and hold down the Shift or Ctrl key after your first selection.

9

Select Connect and choose the Create Spoke SDP Binding option.

The Spoke SDP Binding (Create) form is displayed.

Note: For this function, it is assumed that you clicked on the source site first and then held down the Shift or Ctrl key while right-clicking on the destination site to display the contextual menu.

10

Enable the Auto-Select Transport Tunnel parameter.

11

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new binding between the two sites is displayed in the topology view. See 78.10 “To create an L2 SDP spoke termination on an IES service” (p. 2437) for more detailed information on creating and configuring spoke SDP bindings, if required.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

12

Assuming that the spoke SDP binding was successfully created in Step 11 , select the same two sites again in the topology view, although this time do so in the opposite order that you originally selected them. This will allow you to create a spoke binding for the return tunnel.

13

Right-click on the second site you selected and choose the Create Spoke SDP Binding ... option. The Spoke SDP Binding (Create) form is displayed.

14

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new return binding between the two sites is displayed in the topology view.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information on how to create the required tunnel. Once the tunnel is created, you can repeat this sub-procedure.

15

Close the Service Topology forms.

END OF STEPS

78.65 To delete an IES

CAUTION

Service Disruption

Deleting a service may result in a service disruption for customers. Consider the implication of deleting the service before proceeding.

78.65.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service or a range of services from the list.

3

Click Delete. A warning form appears. This form is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 4 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 4 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 4 .

4

For all services regardless of how their priority is configured, acknowledge the check box that prompts you confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

5

Click Yes to confirm the action. The service is deleted and removed from the list.

6

Close the forms.

END OF STEPS

79 VPRN service management

VPRN service management

79.1 Overview

79.1.1 General information

The NFM-P supports the creation of VPRN services using the 7450 ESS in mixed mode, the 7750 SR, the 7705 SAR, and the 7950 XRS as a PE and provider core (P) router. VPRNs, also called IP VPNs or BGP/MPLS VPNs, are defined in RFC 2547bis. This standard describes a method of forwarding data and distributing routing information across an IP/MPLS service provider core network.

The NFM-P does not support the configuration of CE routers or devices.

The NFM-P supports end-to-end VPRN configuration using tabbed configuration forms with an embedded navigation tree.

The NFM-P supports the configuration in a VPRN of an L3 aggregation mechanism called routed CO. Routed CO uses DHCP relay to manage dynamic subscriber hosts; the network resources for static subscriber hosts are explicitly provisioned. Routed CO supports all residential subscriber management functions of the NFM-P. See Chapter 74, “Residential subscriber management” for more information about residential subscriber management and routed CO.

Routed CO uses a subscriber interface that defines up to 256 subnets. A subscriber interface has child objects called group interfaces. A group interface supports the configuration of multiple SAPs as child objects. A SAP in a group interface supports all residential subscriber management functions. A group interface does not allow the specification of IP subnets or addresses, but inherits the addressing scheme of the parent subscriber interface. The NFM-P service topology map displays VPRN subscriber interfaces, group interfaces, and the associated SAPs.

You can configure NAT for dynamic subscriber hosts in a routed CO deployment. NAT implementation in a VPRN service requires a NAT configuration on the VPRN routing instance and a NAT policy that is associated with a subscriber profile. See Chapter 30, “NAT” for information about configuring and deploying NAT, and about configuring a NAT policy. See 64.4 “To configure a subscriber profile” (p. 1840) for information about associating a NAT policy with a subscriber profile. See 30.9 “To configure NAT on a routing instance” (p. 1086) for information about how to configure NAT on a VPRN site.

A VPRN routed CO allows a service provider to resell wholesale carrier services while providing direct DSLAM connectivity. You can create a VPRN service for the retailer and also define subscriber access and configuration information for the retailer network. See 79.37 “To configure a group interface on a VPRN” (p. 2586) for more information on how to define a wholesale and retail VPRN configuration.

The General tab of the NFM-P service management form displays useful information about the operational state of the service and its sites through the Aggregated Operational State and State Cause indicators.

When you use the NFM-P to create or discover a service, the NFM-P assigns a default Service Tier value to the service. The Service Tier parameter value is relevant only in the context of composite service topology views. See Chapter 85, “Composite service management” for more information about the hierarchical organization of composite services.

Route advertisement

VPRN services use BGP to exchange the VPRN routes among the PE routers that participate in the VPRN. This is done in a way that ensures that routes from different VPRNs remain distinct and separate, even if two VPRNs have an overlapping address space. PE routers distribute routes to CE routers in the VPRN. Since the CE routers do not peer with each other, there is no overlay visible to the routing algorithm of the VPRN. The PE routers use BGP, RIP, OSPFv2, or OSPFv3 as the IGP to distribute internal routes to the CE routers.

Each route in a VPRN service is assigned an MPLS label. When BGP distributes a VPRN route, it also distributes an MPLS label for that route. Before a customer data packet travels across the backbone network, it is encapsulated with the MPLS label that corresponds, in the customer VPRN, to the route that best matches the destination address of the packet.

The MPLS packet is further encapsulated with either another MPLS label or with an IP or GRE tunnel header, so that it gets tunneled across the backbone to the proper PE router. Each route exchanged by the MP-BGP protocol includes an RD that identifies the VPRN association. Thus the backbone core routers do not need to know the VPRN routes.

In addition to exchanging routes using BGP, other types of routes can be leaked from the GRT to a VPRN instance so that they can be used for traffic forwarding or to redistribute those routes to other routers within the routing instance using either BGP or the VPRN IGP instance. Leaking of routes is administered using routing policies. A set of leak export routing policies is added to the NE routing interface. These policies determine the pool of GRT routes that can be imported by a VPRN instance. In the VPRN routing instance, the set of import policies determines which GRT routes from the pool will be added to the route table of the VPRN.

The following types of routes can be leaked from the GRT to a VPRN:

- Direct or local routes
- Static routes
- OSPF routes
- IS-IS routes
- RIP routes

79.1.2 VPRN routing instances

VPRN routing instances are a representation of the objects assigned to a VPRN service such as the assigned interfaces, protocols, confederations, and spoke SDP bindings managed by the NFM-P. The supported objects are displayed in the service navigation tree as child objects to the VPRN routing instance.

You do not create VPRN routing instances; they are created automatically by the NFM-P when you create a VPRN service. You cannot view VPRN routing instances in the navigation tree routing view. You can view them using the Manage→Service→Services main menu, selecting a pre-

configured VPRN service, then choose Properties. The VPRN routing instances are displayed under the Site icon in the service navigation tree.

79.1.3 VPRN service policies

Common to all device services, such as VPRN, are policies that are assigned to the service. Policies are defined at a global level and can then be applied to components of the service, such as interfaces or circuits, when the service is configured or modified. The following policies are common to all device services:

- QoS policies define ingress classification, policing, shaping, and marking on the ingress side of the interface. QoS policies are configured using the SAP Access Ingress Policy form, the SAP Access Egress Policy form, and the ATM QoS Policy form.
- Policer control policies to control access ingress policers and access egress policers under a common hierarchy. Policer control policies are configured using the Policer Control Policy Manager.
- Scheduling policies define hierarchical rate limiting and scheduling to govern the scheduling of queues. Scheduler policies are configured using the Scheduler Policy and HSMDA Scheduler Policy forms.
- Port scheduler policies define hierarchical bandwidth allocation and scheduling at the egress port level. Port scheduler policies are configured using the Port Scheduler Policy form.
- Filter policies control network traffic into or out of an interface or circuit based on DHCP, IP, or MAC matching criteria. Filter policies are configured using the ACL IP Filter form, the ACL MAC Filter form, and the DHCP Filter form.
- Accounting policies measure the traffic on a service to ensure proper billing and enforcement of SLAs. Accounting policies are configured using the Accounting Policy form.
- ANCP policies provide status and control information based on port-up and port-down messages and current line rate changes between the edge device and the access node. ANCP policies are configured using the Manage Subscriber Policies form.
- Routing policies control the size and content of the routing tables, the routes that are advertised, and the best route to take to reach a destination. Routing policies are configured using the Routing Policy Manager.

See Chapter 49, “Policies overview” for more information about policies.

79.1.4 VPRN service validation

The NFM-P provides OAM tools for service validation and for troubleshooting service and network transport issues. You can run an OAM Validation test suite for the service by clicking Validate, or by clicking More Actions and choosing Validate. Alternatively, you can also perform a One Time Validation. If a check mark appears beside the OAM Validation Failed state cause indicator, the test has failed. The Tested Entity Result tab on the Tests tab displays detailed information about the OAM test result. See Chapter 90, “OAM diagnostic tests” for general information about alarm

management using OAM tools. See Chapter 89, “Service Test Manager” for more information about how to configure OAM validation test suites.

The Aggregated Operational State indicator has four possible values: Up, Down, Partially Down, and Unknown. The value is derived from the operational states of the sites that are part of the service, as follows:

- Up—all sites are operationally up
- Partially Down—at least one site is operationally down
- Down—all sites are operationally down
- Unknown—the service has no provisioned sites

When the Aggregated Service Site Operational State is Partially Down or Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P operator. You can view alarms on the Faults page.

When the Aggregated Operational State is Partially Down or Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P operator.

79.1.5 VPRN service routers

A VPRN service consists of CE routers or devices connected to PE routers. PE routers connected to P routers transport data across the IP/MPLS provider core network in service tunnels.

Packets that arrive at an edge 7450 ESS, 7705 SAR, 7750 SR, or 7950 XRS are associated with a VPRN service based on the access interface on which they arrive. An access interface is uniquely identified by the following parameters.

- physical Ethernet port or POS port and channel
- encapsulation type
- encapsulation identifier (if required)

The following table describes the general functions performed by PE, P, and CE routers in a VPRN. See Figure 79-4, “Sample VPRN Configuration” (p. 2525) in this chapter for a sample VPRN. See the appropriate hardware services guide for more detailed information about VPRN functionality on the supported managed devices.

Table 79-1 VPRN router functionality

Router type	Functionality
PE	• Are directly connected to PE, CE, and P routers • Learn VPRN routes from CE devices using e-BGP, RIP, OSPFv2 or OSPFv3, or static routes • Maintain a separate routing table, called a VRF, for each service • Exchange multicast VPRN route information with PE routers in other autonomous systems using MP-BGP • Distribute MPLS inner labels using MP-BGP. Before data traverses the IP/MPLS backbone, it is encapsulated with the MPLS label that corresponds, within the VPRN, to the route that best matches the packet's destination address. • Distribute MPLS outer labels using RSVP-TE or LDP. Before the MPLS packet traverses the IP/MPLS backbone, it is further encapsulated with either another MPLS label or with a GRE or MPLS LSP service tunnel header, so that it is tunneled across the backbone to the appropriate PE router. • Use RDs to identify the VPRN associations • Use RTs to determine when a received route is destined for a VPRN • Terminate RFC 2684-encapsulated IPv4 traffic from ATM access network on SAPs
P	• Are directly connected to PE and P routers • Act as transit LSRs • Maintain routes to PE routers and are unaware of specific VPRN routing information
CE	• Are directly connected to PE routers • Provide customer access to the VPRN

79.1.6 Inter-AS connections

You can connect VPRN service sites (or VRFs) on multiple ASs using EBGP. ASs set up mutual connections by exchanging routing information, such as routes and labels. Labeled VPN-IPv4 routes are distributed within an AS on a PE router using IGBP and between ASs using EGBP on an ASBR. The ASBR redistributes VPN-IPv4 routes to an ASBR in another AS, which in turn distributes the routes to PE routers in its own AS or to an ASBR.

When a VPRN inter-AS connection is between two service providers, the ASs must be on private peering points. For an LSP to operate between ASBRs on the AS borders, EBGP peering must be set up between the ASBRs and MPLS label exchange must be supported. Furthermore, an LSP must run from a packet's ingress PE router to its egress PE router.

You can enable inter-AS connections from the BGP settings in 79.22 "To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance" (p. 2555) .

79.1.7 MP-BGP Multicast IPv4

The MP-BGP multicast extension allows for a network topology that supports both multicast and unicast routing. Routes from the unicast routing table can be imported into the multicast routing table, and routes from the multicast routing table can be imported into the unicast routing table. An ASBR can be configured to advertise VPRN routes to peers in other ASs, redistributing unicast routes learned by BGP into MP-BGP routes, and MP-BGP routes into unicast routes. This configuration enables the support of the two sets of routing information.

The MP-BGP multicast extension specifies that BGP can exchange routing information for the

multicast IPv4 address family within and between BGP ASs. All configurations entered in a multicast IPv4 address family for a BGP instance affect multicast services and are applied to the multicast routing table. See Chapter 28, “Routing protocol configuration” for more information about the configuration of BGP and the MP-BGP multicast extension.

79.1.8 IPv6 support

To configure IPv6 in a VPRN, you must first enable VPN IPv6 for BGP on the base routing instance of each device that acts as a site in the VPRN.

A customer can use an SNMP utility to manage the IPv6 objects in a VPRN service. SNMP mediation of VPRN objects requires the configuration of a community string on each site in the VPRN, regardless of the IP or SNMP version. SNMPv3 mediation of VPRN IPv6 objects, however, requires the additional configuration of an SNMP context for the VPRN using a CLI. See Chapter 9, “Device discovery” for information about configuring an SNMPv3 context for a VPRN.

79.1.9 PIM for VPRN

The PIM protocol can be applied to a VPRN service to create a private multicast distribution network. PIM uses an MDT group address to identify multicast traffic for the VPRN instance to prevent flooding of multicast packets to PE devices in the VPRN. VRFs with the same MDT address are members of that group and receive multicast traffic from each other. The MDT address cannot be in the SSM range.

By default, the PIM protocol only uses the information in the unicast routing table to determine the RPF interface. PIM can be configured to use the separate multicast and unicast routing tables built by MP-BGP to perform RPF lookups for multicast-capable sources to build and maintain distribution trees for multicast traffic forwarding. See Chapter 28, “Routing protocol configuration” for more information about configuring PIM to use multicast or unicast routing tables in RPF lookups.

Data-MDT

A data-MDT is a tunnel for high-bandwidth source traffic through the P-network to interested PE routers. Data-MDTs do not broadcast customer multicast traffic to all PE routers in a multicast domain.

Note: Data-MDTs are only supported for VPRN services.

Multicast data transmission from a CE router is typically delivered to all CE routers in the same multicast group. Some CE routers do not require the delivery of a specific multicast stream because there no downstream receivers for the multicast group. You can prune a PE router from the MDT if the router does not deliver multicast traffic to the attached CE routers. This task is beneficial for high-traffic multicast applications.

A data-MDT allows you to configure a traffic threshold in Kb/s. The NFM-P signals the data-MDTs when the bandwidth for the SSM group exceeds the configured threshold. The PE router sends an MDT join TLV, at 60 s intervals, over the default MDT to all PE routers. The routers respond with the following actions:

- PE routers that require the SSM group specified in the MDT join TLV; join the data-MDT used by the PE router to transmit the SSM group

- PE routers that do not require the SSM group specified in the MDT join TLV; do not join the data-MDT, pruning the PEs from the MDT

The transmitting PE router switches the multicast stream to the data-MDT after allowing the PE routers to join the data MDT. You can configure the data-MDT delay interval using the NFM-P.

The PE router stops sending the MDT join TLV when the transmission bandwidth no longer exceeds the configured threshold. The PE routers using the data-MDT leave the group and transmission resumes over the default MDT.

79.1.10 OSPF sham link support

You can use the OSPF protocol to connect CE routers to PE routers over an MPLS VPN backbone. This can be useful for customers who subscribe to a VPN service and need to use OSPF as their intra-site routing protocol to exchange routing information between their sites. However, there is a potential configuration issue associated with this approach.

OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone. OSPF treats a link through a Layer 3 VPN as an inter-area link. However, other paths between VPN sites may also exist. For example, in the following figure, the link between CE-3 and CE-4 (two CE routers in the same OSPF area) might be a low-speed OC3/STM1 intra-area backup link. OSPF preferentially utilizes intra-area links over inter-area links, and since it establishes an intra-area route connection between CE-3 and CE-4, the potentially high-speed PE-1 to PE-2 primary connection is not utilized.

Figure 79-1 Sham link configuration example

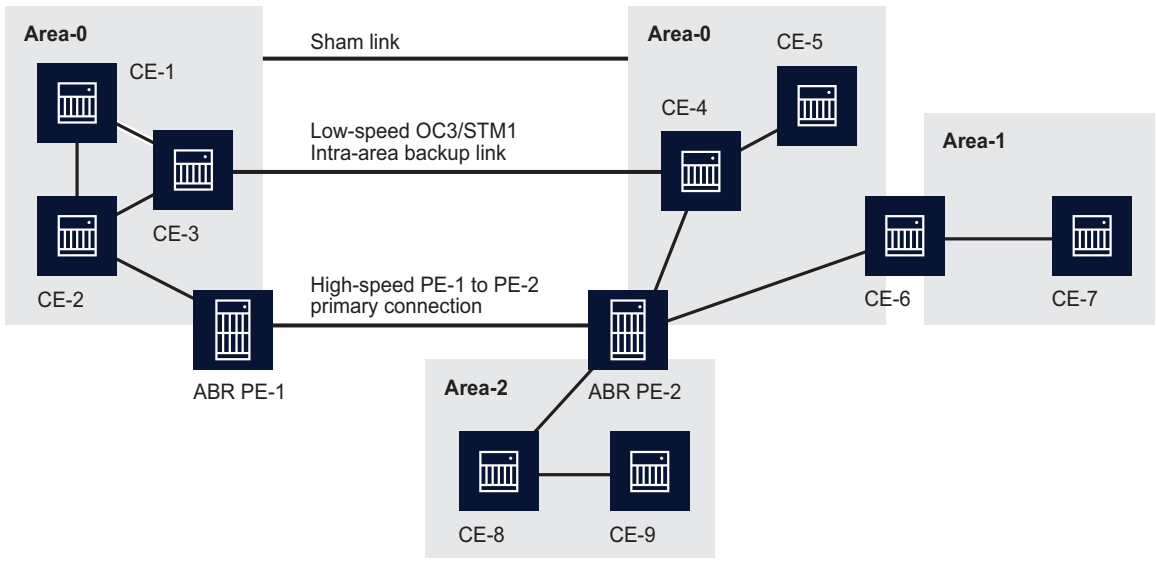

OSPF sham links can be created to resolve this problem. By creating and configuring a sham link as an intra-area link between PE-1 and PE-2, a normal OSPF adjacency is formed, and the link-state database is exchanged across the MPLS VPN. As a result, the desired intra-area connectivity

is created between PE-1 and PE-2. In addition, the cost of the CE-3/CE-4 and PE-1/PE-2 links can be managed by the use of a numerical metric. You could then configure the service so that the CE-3/CE-4 link becomes a standby link only in event that the VPN fails.

79.1.11 ATM SAP terminations for VPRN

CE routers that have access to an ATM network can connect with a VPRN using ATM SAP terminations on a 7750 SR. The interconnection between ATM point-to-point and L3 services uses RFC 2684-encapsulated IPv4 traffic over an ATM PVC that terminates on a specially configured SAP. All RFC 2684- encapsulated traffic can be routed over ATM networks, frame relay, or directly through ATM connections.

The following figure shows how CPE Router A in an ATM network can access L3 IP services, such as a VPRN, using a statically configured ATM PVC on a 7750 SR (SR#1). A SAP is configured on SR #1 to serve a specific VPRN as identified by the VRF. Destination CPE router D can receive RFC 2684-encapsulated traffic over an IP network through a frame relay over 7750 SR 2.

Figure 79-2 ATM SAP network connection to a VPRN

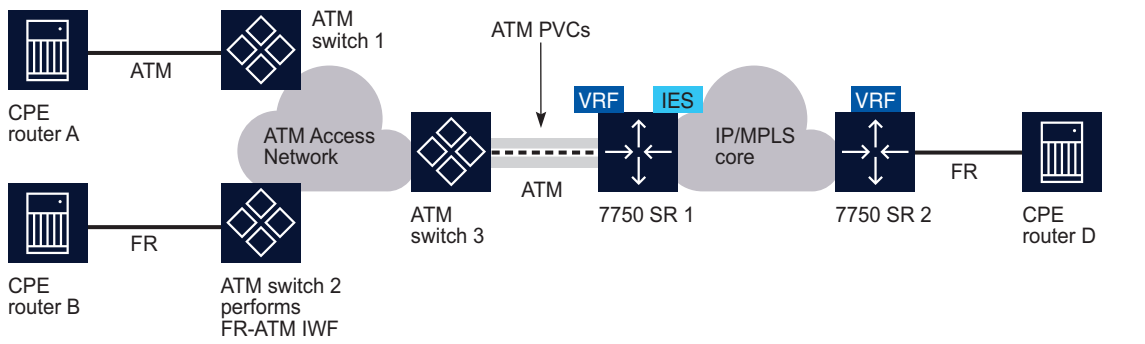

The two connection methods used between ATM and VPRN, which appear in NFM-P as AAL5 Encapsulation parameters. LLC/SNAP encapsulation and VC-multiplexing.

79.1.12 Epipe SDP spoke termination on VPRN services

A VLL Epipe service can terminate directly on a VPRN service using an SDP spoke on the 7750 SR, 7450 ESS, or 7950 XRS. Traffic that terminates on a VPRN service is identified by the interface ID of the SDP on the L2 access router and the VC ID label in the service packet. All routing protocols supported by VPRN are also supported for spoke SDP termination.

The following figure shows a spoke SDP terminating directly on a VPRN. The spoke SDP could be tied to an Epipe or VPLS. No configuration is required for the CE-to-PE connection on the SAP.

Figure 79-3 SDP spoke termination on an L2 service

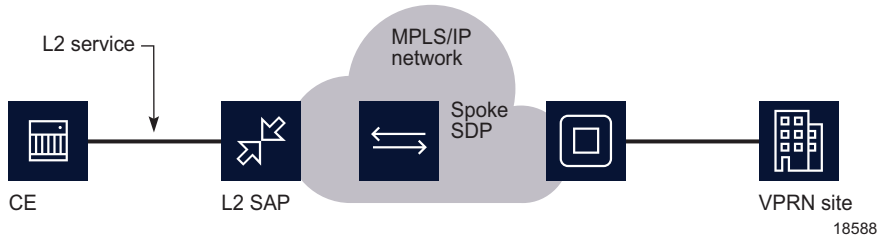

79.1.13 Routed CO dual homing using SRRP

Subscriber Router Redundancy Protocol (SRRP) allows two separate connections to an access NE such as DSLAM to operate in an active/standby configuration similar to the way in which VRRP interfaces operate. SRRP is a collection of functions and messaging protocols that allows a system to create a set of redundant gateway IP addresses that are shared by a local and remote NE.

Each SRRP instance is created within the context of a subscriber group IP interface and is identified by a unique SRRP instance ID, which must be unique within the NE. This SRRP instance controls the redundant routing for all subscriber subnets configured or associated with the group interface. One SRRP instance is supported for each group interface and the SRRP ID must be the same as the SRRP instance ID on the group IP interface on the redundant NE.

A subscriber subnet redundant gateway IP host address is assigned at the subscriber IP interface level and is used for each SRRP instance associated with the subscriber subnet. The redundant IP host address must be configured for a subscriber subnet before it can be associated with an SRRP instance.

When SRRP is active on a group interface, the SRRP instance advertises to a remote NE using in-band messaging on the group-interface SAPs and out-of-band messaging on the group-interface redundant interface. If the remote NE uses the same SRRP instance ID, one NE enters a master state, while the other NE enters a backup state. Since the NEs share a common SRRP gateway MAC address (used for the SRRP gateway IP address and for proxy ARP functions), either NE can act as the default gateway for the attached subscriber hosts. This functionality helps to preserve subscriber QoS enforcement. The master state allows routing to and from the subscriber hosts associated with the group IP interface. The backup state stops ingress forwarding for packets destined to the SRRP gateway MAC and causes all packets destined to subscriber hosts on the group IP interface to be forwarded to a redundant IP interface associated with the group IP interface.

Normally, when anti-spoofing is enabled on a group-interface SAP, the SAP drops SRRP packets because they do not contain a subscriber MAC or IP address. However, you can use a configuration option to enable anti-spoofing for subscriber hosts on a group-interface SAP that participates in SRRP advertisements.

The underlying mechanism that controls state transitions is based on a dynamic priority level that an SRRP instance maintains. The SRRP instance with the highest priority level assumes the master operating state. An SRRP instance with a higher current priority level always preempts an SRRP instance with a lower priority level. If the priority levels are equal, the SRRP instance with the

lowest source SRRP host IP address assumes the master state. The local SRRP instance priority may also be controlled by associating the instance with an existing VRRP policy.

To prevent a flood of AccessInterfaceDown alarms that an SRRP fault or link failure may generate for LAG-based MSAPs, the NFM-P performs alarm suppression. See Chapter 74, “Residential subscriber management” for more information.

The redundant IP interface is a special interface that connects two systems with one or more common SRRP instances. The interface is configured with a /31 address and a spoke SDP binding, creating an Ethernet pseudowire shortcut between the redundant NEs. When the SRRP instance is in backup state, the group interface associated with this instance is not allowed to forward or route traffic downstream towards the subscriber. As a result of this, the packets are shunted across the redundant interface so that the active group interface does the forwarding or routing.

If the redundant IP interface goes down, the system allows the group IP interfaces associated with the down interface to forward locally downstream, when they are in the backup SRRP state. While forwarding downstream in the backup state, the system uses the MAC address associated with the group IP interface, not the SRRP redundant gateway MAC address.

SRRP is supported on the 7450 ESS in mixed mode and 7750 SR.

79.1.14 DoS protection

To protect a VPRN from a high incoming packet rate that characterizes a DoS attack, you can use the NFM-P to create DoS protection policies for the VPRN L3 access interfaces. A DoS protection policy limits the number of control-plane packets that an interface receives each second, and optionally logs a violation notification if a policy limit is exceeded. You can use the NE System Security form to view the violations for a specific NE.

You can configure a DoS protection policy to control the following on a VPRN L3 access interface:

- the control-plane packet arrival rate per subscriber host on the interface
- the overall control-plane packet arrival rate for the interface
- whether an NE sends a notification trap if a policy limit is exceeded

Each VPRN L3 access interface on an NE that supports DoS protection is automatically assigned a default DoS protection policy. This default policy limits only the overall packet arrival rate for the interface, and cannot be deleted or modified. See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for information about creating a DoS protection policy.

79.1.15 DDoS protection

To protect a VPRN from a high incoming packet rate that characterizes a DDoS attack, you can use the NFM-P to configure TMS interfaces to route malicious traffic through the ISA-TMS MDA where the malicious traffic is cleaned before being released to the network.

The TMS interface on a VPRN consists of three VPRNs as follows:

- The management VPRN communicates with the TMS server to determine whether incoming packets are malicious. Configuration of this VPRN is optional.
- The Off-Ramp VPRN receives the malicious traffic and routes it through the ISA-TMS MDA where the traffic is scrubbed. Configuration of this VPRN is optional.

- The On-Ramp VPRN returns the cleaned traffic to the network. Configuration of this VPRN is mandatory.

See 79.49 “To add a TMS interface to a VPRN” (p. 2614) for information about creating a TMS interface.

You can configure a DDoS protection policy on a VPRN group interface SAP, network interface, or L3 access interface. See the procedure to configure an NE DDoS protection policy in the *NSP System Administrator Guide* for more information.

79.1.16 Local DHCP servers

A local DHCP server can be associated with an L3 access interface on a VPRN service. See 74.2 “Residential subscriber components” (p. 1992) .

79.1.17 Local user database

A local user database can be associated with a local DHCP server and PPPoE configurations on group interfaces. See 74.2 “Residential subscriber components” (p. 1992) .

79.1.18 PPPoE protocol on VPRN services

A VPRN service can be configured to run PPPoE protocol. PPPoE is used in subscriber networks to encapsulate PPP frames inside Ethernet frames. PPPoE combines the point-to-point protocol used with DSL sessions with the Ethernet protocol used to support multiple subscribers in a local area network. From the group interface configuration form you can assign a PPPoE policy and a local user database to authenticate PPPoE subscribers.

PPPoE termination in a business VPRN environment is also supported. This ability targets applications such as PPPoE VPRN with IP overlap, where there are two participants in the service:

- The “Wholesale VPRN”, which is a VPRN that provides access to the SAP.
- The “Retail VPRN”, which is a business VPRN that routes the packets belonging to the PPPoE sessions terminating in it. The Retail VPRNs may have overlapping IP addresses.

In this configuration, the PPPoE subscriber host terminates in a Retail VPRN and provides a routed path to the customer site. The VPRN service-id that carries it is determined by the service configuration, specifically:

- If a local user database is used, the Retail Service ID property that you specify in the PPPoE host configuration provides a reference to the VPRN service-id that should be used.
- If RADIUS is used for authentication, the retailer service-id is provided by an Nokia VSA.
- If MSAP is used, the SAP is created in the wholesale VPRN using the information from RADIUS.

The PPPoE session is negotiated with the parameters defined by the Wholesale VPRN interface. Since the IP address space of the subscriber management host may overlap between VPRN services, the node anti-spoofs the packets at access ingress with the session-id.

79.1.19 L2TP on VPRN services

The NFM-P supports the configuration of L2TP on a 7750 SR or 7950 XRS, and on the 7450 ESS in mixed mode. L2TP is a session-layer protocol that extends the PPP model by allowing L2 and

PPP endpoints to reside on different devices that are interconnected by a PSN. L2TP extends the PPP sessions between the CPE and PPP/L2TP termination points on the L2TP network server (LNS), via an intermediate L2TP access concentrator (LAC). The LAC is the initiator of session-generated L2TP tunnels; the LNS is the server that waits for new tunnels. Manually configured and initiated L2TP tunnels can be initiated or stopped from either the LNS or LAC.

At least one ISA-LNS group must be configured for the LNS NE.

On an LNS NE, L2TP destinations configured for L2TP tunnel profiles can include the following:

- loopback L3 access interfaces for a VPRN or IES service
- loopback interfaces configured for a base routing instance

See Chapter 13, “Logical group object configuration” for more information about ISA-LNS groups. See 13.12 “To configure an ISA-LNS group” (p. 426) for information about how to create and configure an ISA-LNS group. See Chapter 28, “Routing protocol configuration” for more information about L2TP.

See 79.22 “To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance” (p. 2555) for information about enabling L2TP on a VPRN router instance site. See 79.37 “To configure a group interface on a VPRN” (p. 2586) for information about configuring a VPRN group interface to terminate LNS PPP sessions.

79.1.20 IPsec

You can configure a VPRN with a tunnel interface for secure and encrypted tunneling between sites. An IPsec VPRN allows you to share secure and encrypted VPN traffic among multiple sites.

IPsec VPRN services include:

- NAT traversal
- DES, 3DES, AES-128, AES-192 and AES-256 encryption methods
- HMAC-MD5 and HMAC-SHA1 authentication and hashing methods
- Diffie-Hellman key generation algorithms
- Pre-shared keys and IKE shared secret with PFS key management authentication methods

See Chapter 34, “IPsec” for more information about IPsec configuration.

79.2 Sample VPRN service configuration

79.2.1 Sample VPRN service

The following figure shows a sample VPRN service configuration.

Figure 79-4 Sample VPRN Configuration

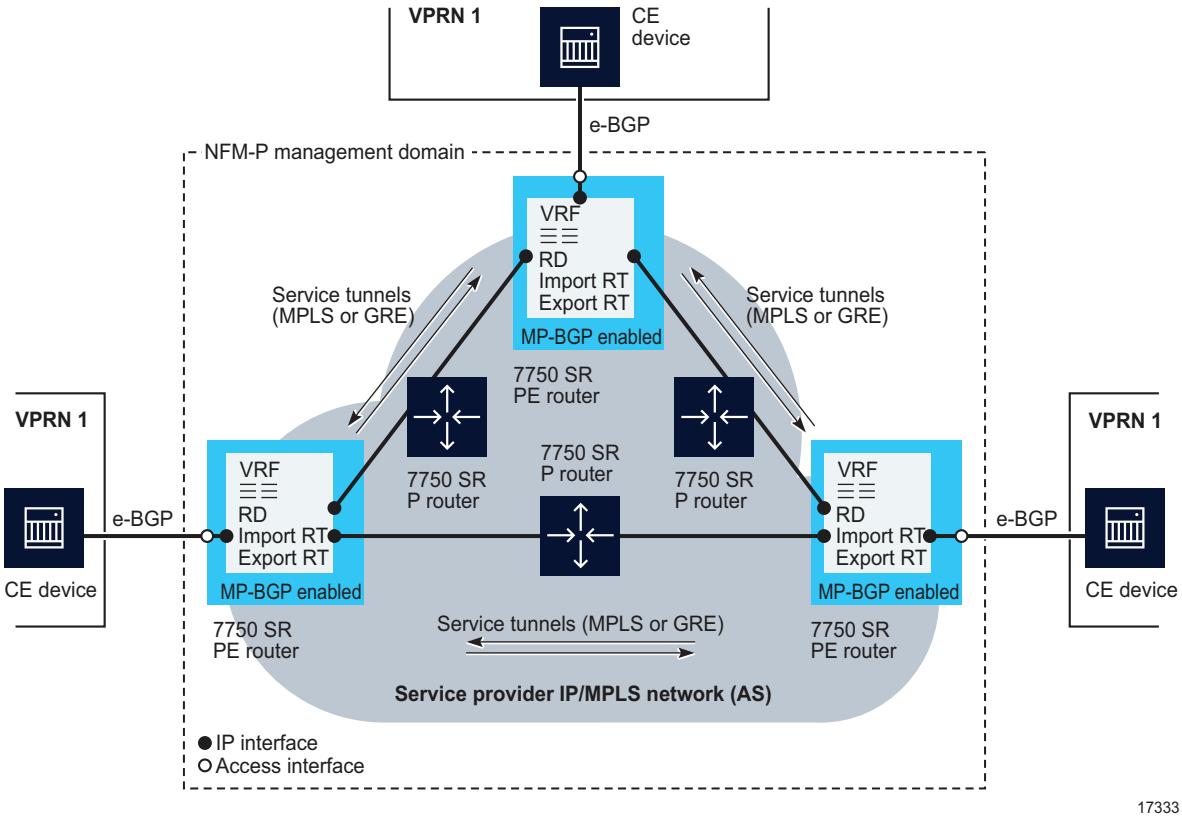

79.2.2 Configuration steps

Assuming the core IP/MPLS or GRE network is already configured, the following high-level tasks are required to configure the sample VPRN service.

1

Configure policies as required, for example, access ingress and egress, Routing, Scheduler, ACL IP, Accounting, and ANCP.

2

Configure ports as access ports for use in the service.

3

Configure service tunnels as required.

4

Configure MP-BGP for PE-to-PE routing.

5

Create and configure customers.

6

Create and configure VPRN 1.

79.3 Sample hub-and-spoke VPRN configuration

79.3.1 Sample hub-and-spoke VPRN

In a hub-and-spoke VPRN, the majority of the traffic is exchanged between the hub (for example, headquarters) and spoke sites (for example, branches). Traffic between spoke sites passes through the hub site. Spoke sites advertise their routes to the hub site, which in turn advertises these routes to the other spoke sites.

The following figure shows a sample hub-and-spoke VPRN service configuration. Your configuration will vary depending on your network requirements.

Figure 79-5 Sample hub-and-spoke VPRN Configuration

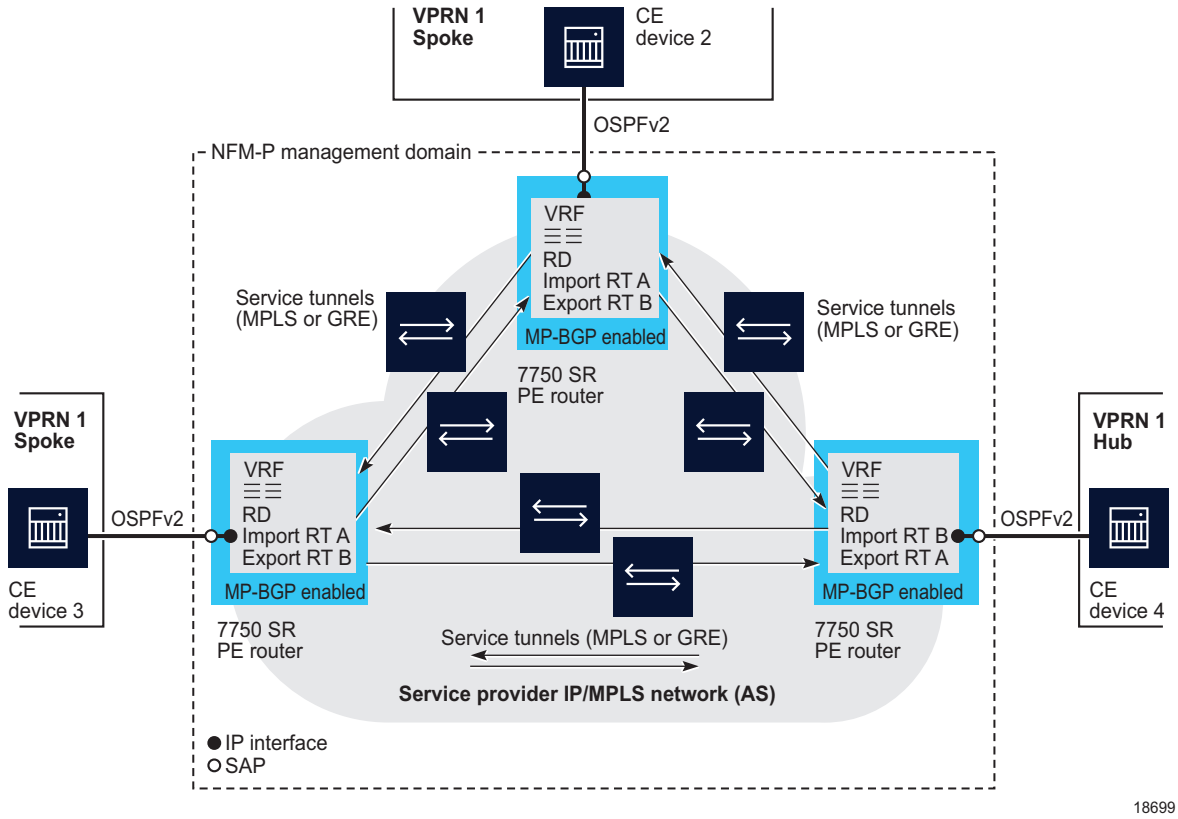

79.3.2 Configuration steps

Assuming the core IP/MPLS or GRE network is already configured, the following table high-level tasks are required to configure this sample hub-and-spoke VPRN service.

1

Configure policies as required, for example, access ingress and egress, Routing, Scheduler, ACL IP, Accounting, ANCP, and DoS protection.

2

Configure ports as access ports for use in the service.

3

Configure service tunnels as required.

4

Configure MP-BGP for PE-to-PE routing.

5

Create and configure customers.

6

Create and configure VPRN 1.

VPRN service management procedures

79.4 Workflow to create a VPRN service

79.4.1 Overview

The following workflow lists the high-level steps required to create a VPRN service. As a prerequisite for creating a VPRN service, this workflow assumes the following:

- a group or customer with the required user access privileges has been configured.
- the IP or IP/MPLS core network exists.
- any required service tunnels are created including the static, dynamic or SR-TE LSP required to create the service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
- the access ports for the service are created; see Chapter 16, “Port and channel object configuration” .
- any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite policies are created; see Chapter 49, “Policies overview” . You do not have to create pre-defined policies if policies are created on a per-service basis.
- any required MP-BGP for PE-to-PE routing is configured; see Chapter 28, “Routing protocol configuration” for more information about protocol configuration.

79.4.2 Stages

1

Create a VPRN service; see 79.5 “To create a VPRN service” (p. 2534) .

2

Create a VPRN site; see 79.11 “To configure a VPRN site” (p. 2545) .

3

Configure the VPRN site TCP/IP assignments, IP addressing, and timing protocols as required.

a. Configure NAT on a VPRN site; see 30.9 “To configure NAT on a routing instance” (p. 1086) .

b. Configure DNS for a VPRN site; see 79.12 “To configure DNS for a VPRN site” (p. 2546) .

c. Configure QoS for self-generated traffic on a VPRN site; see 79.13 “To configure QoS for self-generated traffic on a VPRN site” (p. 2546) .

d. Configure NTP on a VPRN site; see 79.14 “To configure NTP on a VPRN site” (p. 2547) .

e. Configure PTP on a VPRN site; see 79.15 “To configure PTP on a VPRN site” (p. 2548) .

f. Configure an SNMP community for mediating the IPv6 objects on a VPRN site; see 79.17 “To configure an SNMP community on a VPRN site” (p. 2551) .

g. Configure IGMP host tracking on a VPRN sit; see 79.18 “To configure IGMP host tracking on a VPRN site” (p. 2552) .

h. Configure an override source IP address on a VPRN site; see 79.19 “To configure an override source IP address on a VPRN site” (p. 2553) .

4

Enable one or more routing protocols on a VPRN site as required.

a. Enable routing protocols on a VPRN site; see 79.21 “To enable routing protocols on a VPRN site” (p. 2554) .

b. Configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance; see 79.22 “To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance” (p. 2555) .

c. Configure BGP on a VPRN routing instance; see 79.23 “To configure BGP on a VPRN routing instance” (p. 2557) .

d. Configure IGMP on a VPRN routing instance; see 79.24 “To configure IGMP on a VPRN routing instance” (p. 2560) .

5

Configure or add the following objects or instances on a VPRN site or service.

a. Routing instance; see 79.26 “To configure a routing instance on a VPRN site” (p. 2564) .

b. VRF instance; see 79.27 “To configure a VRF instance on a VPRN site” (p. 2567) .

c. MVPN VRF instance; see 79.28 “To configure an MVPN VRF instance on a VPRN site” (p. 2570) .

d. BGP confederation; see 79.29 “To configure a BGP confederation on a VPRN site” (p. 2574) .

e. GSMP group; see 79.30 “To configure a GSMP group on a VPRN site” (p. 2575) .

f. GNE site; see 79.31 “To configure a GNE site and GNE service interface on a VPRN service” (p. 2576) .

6

Configure the following server types for a VPRN site or service as required.

a. Local DHCPv4 server; see 79.33 “To configure a local DHCPv4 server on a VPRN site” (p. 2579) .

b. Local DHCPv6 server; see 79.34 “To configure a local DHCPv6 server on a VPRN site” (p. 2581) .

c. RADIUS server; see 27.8 “To configure a RADIUS server on a routing instance” (p. 846) .

d. RADIUS proxy server; see 27.9 “To configure a RADIUS proxy server on a routing instance” (p. 847) .

7

Create or add the following interfaces to a VPRN site as required.

a. AA interfaces; see 78.17 “To add an AA interface to an IES or a VPRN site” (p. 2447) .

b. AARP interfaces; see 78.18 “To add an AARP interface to an IES or a VPRN site” (p. 2448) .

c. GNE service interfaces; see 79.31 “To configure a GNE site and GNE service interface on a VPRN service” (p. 2576) .

d. Group interfaces; see 79.37 “To configure a group interface on a VPRN” (p. 2586) .

   In addition, perform the following as required on VPRN group interfaces:

   - Configure a SAP on a VPRN group interface; see 79.39 “To configure a SAP on a VPRN group interface” (p. 2594) .
   - Configure LAG per-link hashing on a VPRN group interface SAP; see 79.40 “To configure LAG per-link hashing on a VPRN group interface SAP” (p. 2599) .
   - Configure a WLAN GW for a VPRN group interface; see 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631) .

e. IGMP interfaces; see 79.41 “To add an IGMP interface to a VPRN” (p. 2600) .

f. IGMP Group interfaces; see 79.42 “To add an IGMP group interface to a VPRN” (p. 2601) .

g. IP mirror interfaces; see 79.43 “To add an IP mirror interface to a VPRN” (p. 2602) .

h. ISIS interfaces; see 28.60 “To configure IS-IS on a routing instance” (p. 953) .

i. L3 access interfaces; see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) . See the workflow in 79.82 “Workflow to configure VPRN L3 access interfaces” (p. 2653) for detailed instructions/tasks that apply to VPRN L3 access interfaces.

j. MLD interfaces; see 28.123 “To configure an MLD interface on a base routing instance or VPRN routing instance” (p. 1035) .

k. Network interfaces; see 79.44 “To configure a network interface on a VPRN site” (p. 2603) .

l. PIM interfaces; see 79.45 “To add a PIM interface to a VPRN” (p. 2606) .

m. Redundant interfaces; see 79.46 “To implement dual homing using SRRP” (p. 2607) .

n. Subscriber interfaces; see 79.47 “To configure a subscriber interface on a VPRN” (p. 2610) .

   In addition, perform the following as required on VPRN subscriber interfaces:

   - Force a WLAN GW switchover to standby management on a VPRN subscriber interface; see 79.48 “To force a WLAN GW switchover to standby management on a VPRN subscriber interface” (p. 2613) .

o. TMS interfaces; see 79.49 “To add a TMS interface to a VPRN” (p. 2614) .

p. Tunnel interfaces; see 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249) .

q. Video interfaces; see 35.3 “To add a video interface to an IES or VPRN site” (p. 1275) .

8

Create and assign policies to VPRN sites as required:

a. Configure VRF import and export policies on a VPRN site; see 79.50 “To configure VRF import and export policies on a VPRN site” (p. 2615) .

b. Configure a range policy to reserve route targets and route distinguishers on a VPRN site; see 79.51 “To configure a policy to reserve an RT and RD range for VPRN services” (p. 2616) .

c. Automatically assign RT policies and RD configuration to VPRN sites; see 79.52 “To automatically assign RT policies and RD configuration to VPRN sites” (p. 2617) .

d. Configure a network ingress filter policy on a VPRN site; see 79.53 “To configure a network ingress filter policy on a VPRN site” (p. 2618) .

e. Configure ingress QoS control on a VPRN site; see 79.54 “To configure ingress QoS policies on a VPRN site” (p. 2619) .

9

If GRT route leaking is required, configure the following:

a. Configure routing policies as required; see 54.5 “To configure a routing policy statement” (p. 1745).

b. Configure leak export policies on the NE routing instance; see 27.2 “To configure a routing instance or a VRF instance” (p. 826).

c. Configure leak import policies on the VPRN routing instance; see 79.26 “To configure a routing instance on a VPRN site” (p. 2564).

10

Configure the VPRN service connections/routes/targets/links/tunnels between devices as required.

a. Bind a VPRN site to service tunnels; see 79.56 “To bind a VPRN site to service tunnels” (p. 2620) .

b. Configure static routes on a VPRN site; see 79.57 “To configure static routes on a VPRN site” (p. 2621) .

c. Configure route aggregates on a VPRN site; see 79.58 “To configure route aggregates on a VPRN site” (p. 2622) .

d. Enable a tunnel facility MEP on the VPRN site; see 79.59 “To enable a tunnel facility MEP on the VPRN site” (p. 2623) .

e. Enable the automatic selection of an RD on a VPRN site; see 79.60 “To enable the automatic selection of a route distinguisher on a VPRN site” (p. 2624) .

f. Add a Global Route Table to a VPRN site; see 79.61 “To add a Global Route Table to a VPRN site” (p. 2625) .

g. Create an intra-area OSPF sham link between two VPRN sites; see 79.62 “To create an OSPF sham link between two VPRN sites” (p. 2626) .

h. Reserve route targets for specific VPRN services; see 79.63 “To reserve route targets for specific VPRN services” (p. 2628) .

i. Create a VXLAN tunnel termination; see 79.64 “To configure a VXLAN termination on a VPRN” (p. 2629).

11

Configure WLAN Gateway functionality on a VPRN site as required.

a. Configure WLAN GW functionality on a VPRN site; see 79.65 “To configure WLAN GW functionality on a VPRN site” (p. 2630) .

b. Configure a WLAN GW for a VPRN group interface; see 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631) .

c. Resync WLAN GW tunnels on a VPRN site; see 79.67 “To resync WLAN GW tunnels on a VPRN site” (p. 2635) .

12

Configure or add Spoke SDP Binding to a VPRN site as required.

a. Create a VPRN spoke SDP binding; see 79.68 “To configure a VPRN spoke SDP binding” (p. 2635) .

b. Create an L2 SDP spoke termination on a VPRN service; see 79.69 “To create an L2 SDP spoke termination on a VPRN service” (p. 2638) .

c. Configure an MPLS-TP static pseudowire on a VPRN spoke SDP binding; see 79.70 “To configure an MPLS-TP static pseudowire on a VPRN spoke SDP binding” (p. 2641) .

d. Configure BFD on a VPRN spoke SDP binding; see 79.71 “To configure BFD on a VPRN spoke SDP binding” (p. 2642) .

e. Clear BFD sessions and statistics on a VPRN spoke SDP binding; see 79.72 “To clear BFD sessions and statistics on a VPRN spoke SDP binding” (p. 2643) .

f. View the BFD session status on a VPRN SDP spoke binding; see 79.73 “To view the BFD session status on a VPRN SDP spoke binding” (p. 2644) .

13

Perform service assurance/OAM related task as required on VPRN services or sites:

a. Run an OAM validation test for a VPRN; see 79.74 “To run an OAM validation test for a VPRN service” (p. 2644) .

b. Configure OAM components on a VPRN site; see 79.75 “To configure OAM components on a VPRN site” (p. 2646) .

c. Create a TWAMP Light reflector on a VPRN site; see 79.76 “To create a TWAMP Light reflector on a VPRN site” (p. 2647) .

d. Assign an ICMP Ping template. See 79.105 “To assign an ICMP ping template to a VPRN L3 access interface” (p. 2679).

14

View VPRN service information assigned to a VPRN site; as required:

a. MVPN extranet objects; see 27.28 “To list MVPN Extranet objects for a NE” (p. 873) .

b. Last cleared BFD statistics and sessions; see 79.77 “To view the last cleared BFD statistics and sessions on a VPRN site” (p. 2648) .

c. VPRN services that use an auto-assigned RT and RD or reservation table; see 79.78 “To view VPRN services that use an auto-assigned RT and RD or reservation table” (p. 2649) .

d. DHCPv6 prefixes that the server has given out; see 79.79 “To view DHCPv6 leases” (p. 2649) .

e. Lease not owner and pool unknown log events for local DHCPv6 servers; see 79.80 “To view DHCPv6 log events” (p. 2651) .

f. Service topology map associated with a VPRN service; see 79.81 “To view the service topology map associated with a VPRN service” (p. 2652) .

g. VPRN service operational status; see 79.6 “To view the VPRN service operational status” (p. 2535) .

15

As required, modify the VPRN service:

a. Using the Show Info form. See 79.8 “To view VPRN service contents” (p. 2537) .

b. Using the Manage Services form; see 79.7 “To modify a VPRN service” (p. 2536) .

c. Using the topology view; see 79.9 “To modify a VPRN service using the topology view” (p. 2538) .

16

As required, delete the VPRN service; see 79.10 “To delete a VPRN service” (p. 2544) .

79.5 To create a VPRN service

79.5.1 Steps

1

Choose Create→Service→VPRN from the NFM-P main menu. The VPRN Services (Create) form opens.

2

Select a customer in the Customer panel to associate with the VPRN.

3

Configure the general VPRN service parameters.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Save and close the forms.

END OF STEPS

79.6 To view the VPRN service operational status

79.6.1 Purpose

The Aggregated Operational State and State Cause indicators on the General tab of the service management form display information about service faults.

79.6.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VPRN Service (Edit) form opens.

3

View the Aggregated Operational State and State Cause indicators. When the Aggregated Operational State is Down or Partially Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The Object Alarms tab is displayed.

6

Click on the Aggregated Alarms tab to view the aggregated alarms for the object. The Aggregated Alarms tab is displayed.

7

Close the forms.

END OF STEPS

79.7 To modify a VPRN service

CAUTION

Service Disruption

Modifying parameters can be service-affecting.

The behavior of the VPRN service may become unpredictable if modifications to the configuration affect the IPsec portion of the service configuration. For example, if a VPRN service is configured with IPsec tunnels, IPsec SAPS, and policies are deleted, the service is not deleted from the NFM-P and the service will be in an inconsistent state. The IPSEC portion of the VPRN configuration must be deleted using CLI scripts or the CLI before the VPRN service can be deleted from the NFM-P.

79.7.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

The following tabs list the service elements that can be individually or collectively selected and configured:

- Tests tab — allows the creation and execution of service-specific diagnostic tests
- Sites tab — lists the sites that are included in the service
- L3 Access Interfaces tab — lists the L3 access interfaces that are included in the service
- Spoke SDP Bindings tab — displays the spoke SDP bindings that are associated with the service
- Addresses tab — lists the IP addresses that are associated with the service

  Note:

  You cannot remove an IP address from an interface when the IP address of a static host is defined in the subnet of the interface IP address and the ARP Populate parameter is enabled on the Anti-Spoofing tab.
- Faults tab — displays the faults associated with the service

3

Modify the parameters for the service, as required.

To configure items in the tabs that contain lists of service elements, select the items and click Properties.

4

Save the changes and close the forms.

END OF STEPS

79.8 To view VPRN service contents

79.8.1 Purpose

Use this procedure to view various contents of a VPRN service and any modifications you make to it before deploying the changes.

Note: The procedure also provides information on the following policy types associated with the service:

- QoS SAP ingress and QoS SAP egress polices
- ACL IP/IPv6 filters and ACL MAC filter policies
- VRF Import and Export policies

Two information views are available. The Committed info view displays the current contents of a service, and the Committed menu item is always enabled. The Modified info view allows you to review any changes you make to a service before committing them. Modified, created, and deleted attributes and objects are displayed.

79.8.2 Steps

1

Choose Manage Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search and select the service that you want to view or modify.

3

On the service navigation tree, expand the Sites icon and click on the required routing instance. The site properties form is displayed.

4

Perform one of the following:

a. To view the currently committed service contents, go to Step 5 .

b. To view modified service contents before committing any changes, go to Step 6 .

5

Click on the Show Info and choose the Committed menu item. A Committed Values form is displayed and shows various policy attributes that have been previously configured and saved

in the local policy. The information displayed in the form is similar to the information retrieved in CLI by running the “config>service>*Service Type#* info” command, where *Service Type#* is the type of service (followed by its Service ID, for example: vprn77) that you want to query.

6

Modify any service parameters or objects, as required. Otherwise go to Step 9 .

7

Click on the Show Info and choose the Modified menu item. A Modified Values form is displayed. The table lists modified, created, and deleted actions, as well as specific attributes and objects, along with their old value, new value, and tab location. The Attribute Title corresponds to the attribute or object name acted upon by your current modifications. For created objects, the values of mandatory attributes are shown in comma-separated format.

8

Select an item in the Modified Values form and then click Show on Form. The service form tab containing the changed item is displayed and the modified attribute is highlighted in blue.

9

Save your changes if required, and close the form.

END OF STEPS

79.9 To modify a VPRN service using the topology view

79.9.1 Purpose

The topology view for a service provides a graphical representation of the various components and their interconnections. You can also use this view to add, modify, or just navigate to service components. This provides an alternative approach to performing these functions from the navigation tree view.

Working from the topology view can expedite the creation of the components, since many of the fields you would ordinarily have to set in the configuration forms will be automatically populated using this approach. The configuration forms can also be accessed directly at any time from this view by right-clicking a component. This allows quick access to conduct more detailed component configuration.

79.9.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose a VPRN service and click Topology View. The Service Topology map opens.

The remainder of this procedure contains a number of sub-procedures describing the various components that can be viewed, created, or modified from the topology view. These include:

- Creating a new site. Go to Step 4 .
- Creating site components. Go to Step 9 .
- Creating spoke SDP bindings. Go to Step 36 .
- Viewing the route target topology. Go to Step 44 .

Adding a new site

4

Right-click on any blank space in the service topology map. A contextual menu is displayed. Choose the Create VPRN Site option.

The Select Network Elements form opens.

5

Select one or more sites to add to the service and click OK. The VPRN Site (Create) form for the new site is displayed. If you selected more than one site, the VPRN Site (Multiple Instances) (Create) form for the new sites is displayed.

6

Click on OK. The VPRN Site (Create) (or VPRN Site (Multiple Instances) (Create)) form closes and the new site (or sites) is displayed on the map.

7

If you need to perform detailed configuration of site properties for the new site, right-click the site icon and select Properties. The VPRN Site (Edit) form opens. See 79.5 “To create a VPRN service” (p. 2534) for detailed site configuration information.

8

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 46 to finish.

Creating site components

9

Right-click on any site icon in the service topology map. A contextual menu is displayed. You can choose to create one of the following:

- VPRN L3 Access Interface. Go to Step 10 .
- Tunnel Interface. Go to Step 15 .

- VPRN Subscriber Interface. Go to Step 20 .
- Redundant Interface. Go to Step 24 .
- IP Mirror Interface. Go to Step 28 .
- Video Interface. Go to Step 32 .

10

If you choose to create a VPRN L3 Access Interface, then the VPRN L3 Access Interface (Create) form is displayed.

11

Configure the Name parameter on the General tab.

12

Click on the Port tab and assign a port to the interface.

See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) for detailed information on further configuring the interface, if required.

13

Click OK. The VPRN L3 Access Interface (Create) form closes and the new L3 access interface is displayed in the topology view.

14

Go to Step 35 .

15

If you choose to create a tunnel interface, then the Tunnel Interface (Create) form is displayed.

16

Configure the Name parameter on the General tab.

17

Click on the Port tab and assign a port to the interface.

See 34.20 “To configure a tunnel interface on an IES or VPRN” (p. 1249) in Chapter 34, “IPsec” for detailed information on further configuring the interface, if required.

18

Click OK. The Tunnel Interface (Create) form closes and the new tunnel interface is displayed in the topology view.

19

Go to Step 35 .

20

If you choose to create a VPRN Subscriber Interface, then the VPRN Subscriber Interface (Create) form is displayed.

21

Configure the Name parameter for the interface.

See 79.47 “To configure a subscriber interface on a VPRN” (p. 2610) for detailed information on further configuring the interface, if required.

22

Click OK. The VPRN Subscriber Interface (Create) form closes and the new subscriber interface is displayed in the topology view.

23

Go to Step 35 .

24

If you choose to create a Redundant Interface, then the Redundant Interface (Create) form is displayed.

25

Configure the Name parameter for the interface.

See 79.46 “To implement dual homing using SRRP” (p. 2607) for detailed information on further configuring the interface, if required.

26

Click OK. The Redundant Interface (Create) form closes and the new redundant interface is displayed in the topology view.

27

Go to Step 35 .

28

If you choose to create an IP Mirror Interface, then the IP Mirror Interface (Create) form is displayed.

29

Configure the Name parameter for the interface.

See 79.43 “To add an IP mirror interface to a VPRN” (p. 2602) for detailed information on further configuring the interface, if required.

30

Click OK. The IP Mirror Interface (Create) form closes.

31

Go to Step 35 .

32

If you choose to create a Video Interface, then the Video Interface (Create) form is displayed.

33

Configure the Name parameter for the interface.

See 35.3 “To add a video interface to an IES or VPRN site” (p. 1275) in Chapter 35, “ISA-Video” for detailed information on further configuring the interface, if required.

34

Click OK. The Video Interface (Create) form closes and the new video interface is displayed in the topology view.

35

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 46 to finish.

Creating spoke SDP bindings

36

Select the sites you need to connect in the service topology map and right-click on any one of them. A contextual menu is displayed.

Note: When you create a spoke binding between two sites, the order in which you select them is important. The first site you select will become the source site and the second site will become the destination site. Therefore, it is not recommended that you do a marquee-select in the topology view, since you will not be sure of this hierarchy. Instead, select the sites individually, and hold down the Shift or Ctrl key after your first selection.

37

Select Connect and choose the Create Spoke SDP Binding option.

The Spoke SDP Binding (Create) form is displayed.

Note: For this function, it is assumed that you clicked on the source site first and then held down the Shift or Ctrl key while right-clicking on the destination site to display the contextual menu.

38

Enable the Auto-Select Transport Tunnel parameter.

39

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new binding between the two sites is displayed in the topology view. See 79.68 “To configure a VPRN spoke SDP binding” (p. 2635) for more detailed information on creating and configuring spoke SDP bindings, if required.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information about how to create the required tunnel. After the tunnel is created, you can repeat this sub-procedure.

40

Assuming that the spoke SDP binding was successfully created in Step 39 , select the same two sites again in the topology view, although this time do so in the opposite order that you originally selected them. This allows you to create a spoke binding for the return tunnel.

41

Right-click on the second site you selected and choose the Create Spoke SDP Binding ... option. The Spoke SDP Binding (Create) form is displayed.

42

You can manually configure other parameters here if required, or click OK. One of the following will result:

- If an available transport tunnel exists between the two sites, then the Spoke SDP Binding (Create) form closes and the new return binding between the two sites is displayed in the topology view.
- If an available transport tunnel does not exist between the two sites, then an error message is displayed to that affect. See Chapter 33, “Service tunnels” for information about how to create the required tunnel. After the tunnel is created, you can repeat this sub-procedure.

43

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 46 to finish.

Viewing the route target topology

44

Right-click on any blank space in the service topology map. A contextual menu is displayed. Choose the Highlight Route Target Topology option.

This option essentially draws dotted lines to represent the route target topology between NEs. For example, for two NEs A and B, if the import and export route targets match, then two dotted lines will be drawn on the map. One will represent the route target going from A to B, and the other from B to A.

45

Return to Step 3 for a list of other functions you can perform from the topology view or go to Step 46 to finish.

46

Close the Service Topology form.

47

Close the Manage Services form.

END OF STEPS

79.10 To delete a VPRN service

CAUTION

Service disruption

Deleting a service may result in a service disruption for customers.

Consider the implication of deleting the service before proceeding.

79.10.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

As required, configure the filter criteria to locate the service or range of services to be deleted. A list of services appears.

3

Choose a service or a range of services from the list.

Note: You cannot delete a VPRN service when IPsec security policies, interfaces, or tunnels are configured. See the appropriate node documentation for more information about modifying or deleting IPsec configurations.
For example, if a VPRN service is configured with IPsec tunnels, IPsec SAPS, and policies is deleted, the service is not deleted from the NFM-P and the service will be in an inconsistent state. The IPSEC portion of the VPRN configuration must be deleted using the CLI scripts or through the CLI before the VPRN service can be deleted from the NFM-P.
The L3 interface properties for the IPSEC SAP cannot be configured using the NFM-P.

4

Click Delete. A warning form opens. This form is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 5 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 5 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 5 .

5

For all services regardless of how their priority is configured, acknowledge the check box that prompts you confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

6

Click Yes to confirm the action. The service is deleted and removed from the list.

7

Close the forms.

END OF STEPS

79.11 To configure a VPRN site

79.11.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, right-click on the Sites icon and choose Create VPRN Site, or expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Create|Edit) form opens.

4

Configure the general site parameters.

5

Save and close the forms.

END OF STEPS

79.12 To configure DNS for a VPRN site

79.12.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Enable the Enable DNS parameter.

5

Click on the DNS tab and configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

79.13 To configure QoS for self-generated traffic on a VPRN site

79.13.1 Before you begin

NEs produce SGT for various applications, for example Telnet, SNMP, SSH. For each application, you can configure the DSCP or dot1p value for the traffic generated by that application. You can also map DSCP values to forwarding classes.

79.13.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance, and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Self Generated Traffic tab and perform the following as required:

a. Configure DSCP values for the required SGT applications.
   1. Click on the DSCP Marking tab and select an application in the list, then click Properties. The Application DSCP Marking form opens.
   2. Configure the DSCP parameter.
   3. Save your changes and close the form.

b. Map DSCP names to forwarding classes.
   1. Click on the DSCP Mapping tab and select a DSCP name in the list, then click Properties. The Application DCSP Mapping form opens.
   2. Configure the Forwarding Class parameter.
   3. Save your changes and close the form.

c. Configure dot1p values for the required SGT applications.
   1. Click on the Dot1p Marking tab and select an application in the list, then click Properties. The Application Dot1p Marking form opens.
   2. Configure the Dot1p parameter.
   3. Save your changes and close the form.

5

Save your changes and close the forms.

END OF STEPS

79.14 To configure NTP on a VPRN site

Note: NTP must be enabled on the base router before it can be configured on a VPRN site. See 15.34 “To configure NTP on supported devices” (p. 490) for information about configuring NTP on the 7705 SAR and 7750 SR.

79.14.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Set the NTP State parameter to Enabled.

5

Click on the NTP tab and configure the required parameters.

6

To create an authentication key:

1. Click on the Authentication tab.
2. Click Create. The Network Time Protocol Authentication (New Instance) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

7

To configure an NTP broadcast source:

1. Click on the Broadcast tab.
2. Click Create. The NTP Broadcast (New Instance) form opens.
3. Select a source interface (NTP Broadcast).
4. Configure the required parameters.
5. Save the changes and close the form.

8

Save the changes and close the forms.

END OF STEPS

79.15 To configure PTP on a VPRN site

Note: PTP must be enabled on the base router before it can be configured on a VPRN site. See 15.21 “To configure the IEEE 1588 PTP clock on a 7210 SAS, 7250 IXR, 7450 ESS, or 7750 SR” (p. 478) for information about configuring PTP on the 7705 SAR and 7750 SR.

79.15.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Set the PTP State parameter to Enabled.

5

Click on the PTP tab and configure the required general parameters.

6

To create PTP peers:

1. Click on the Peers tab.
2. Click Create. The IEEE 1588 PTP Peer (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

7

Save the changes and close the forms.

END OF STEPS

79.16 To configure GTP on a VPRN site

Note: GTP must be enabled on the base router before it can be configured on a VPRN site; 27.14 “To configure GTP on a routing instance” (p. 853).

79.16.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Set the Enable GTP Uplink parameter to Enabled.

5

Click on the GTP tab to configure S11 and uplink interfaces.

6

Click on the On the S11 tab to configure S11 interfaces and peer profile mappings:

1. On the Interfaces tab, click Create. The S11 Interface (Create) form opens.
2. Select an interface and an APN policy.
3. Save the changes and close the form.
4. Click on the Peer Profile Mapping tab and click Create. The S11 Peer Profile Mapping (Create) form opens.
5. Configure the Address Prefix and Prefix Length parameters.
6. Select a Peer Profile.
7. Save the changes and close the form.

7

Click on the On the Uplink tab to configure an uplink interface and peer profile mappings:

1. On the General tab, configure the APN Network Identifier and PDN Type parameters.
2. Click on the Peer Profile Mapping tab and click Create. The Uplink Peer Profile Mapping (Create) form opens.
3. Configure the Address Prefix and Prefix Length parameters.
4. Select a Peer Profile.
5. Save the changes and close the form.

8

Save the changes and close the forms.

END OF STEPS

79.17 To configure an SNMP community on a VPRN site

79.17.1 Purpose

If the VPRN site is to support IPv6 addressing (for router advertisement or neighbor discovery) or VRRP objects, perform the following procedure to configure an SNMP community for mediating the IPv6 objects. You can configure multiple read-only and read-write SNMP communities on a VPRN site for virtual CPE management.

When the NFM-P uses SNMPv2 for device mediation, you must configure one and only one SNMP community string for the VPRN site. Otherwise, there is no mediation of the VPRN IPv6 objects on the site, and the NFM-P raises an alarm. The alarm is cleared and mediation resumes after the configuration is modified so that exactly one SNMP community string is associated with the VPRN site. The SNMP community must be set before an ISIS site can be added to the VPRN.

You can associate a list of IP addresses with an SNMP community. Only addresses that appear on the list have permission to use the associated SNMP community. Creating and configuring the list is performed using CLI; see the appropriate node documentation for more information about using SNMP community configuration commands.

79.17.2 Source access lists

You can use SNMP source access lists configured on the NE when configuring an SNMP community on a VPRN service. You can add SNMP source access lists to a VPRN site SNMP community string. Source access control lists cannot be configured through SNMP. They need to be configured on the NE using the CLI. You cannot modify or delete SNMP source access lists in the NFM-P.

79.17.3 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Configure the SNMP community:

1. Click on the SNMP Community tab.
2. Click Create. The SNMP Community (Create) form opens.
3. Configure the SNMP Community String parameter.

4. Configure the Read Only parameter to specify whether the community string of the VPRN is read-only or read-write for virtual CPE management.

   Note:

   You must also enable the Allow SNMP Access parameter on the VPRN routing instance. See Step 7 of 79.26 “To configure a routing instance on a VPRN site” (p. 2564) .

5. Select an SNMP source access list configured on the NE, if required.
6. Save your changes and close the form.

5

Save and close the forms.

END OF STEPS

79.18 To configure IGMP host tracking on a VPRN site

79.18.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Multicast tab, then on the IGMP Host Tracking tab.

5

Configure the required parameters.

6

To clear host tracking info or host tracking statistics, click Clear.

7

Save the changes and close the forms.

END OF STEPS

79.19 To configure an override source IP address on a VPRN site

79.19.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Source Addresses tab and click Create. The Source Address (Create) form opens.

5

Configure the Source IP Application parameter.

6

Set the Source Address Termination parameter to IP Address.

7

Configure the Source IP Address parameter. Enter a valid IP address that the VPRN virtual routing instance will identify.

Note: An IPv6 address option for the Source IP Address parameter is available when the IPv6 Allowed parameter is set during network interface creation.
You must select an IPv6 Source IP Address before you can select IPv6 Source IP Applications.

8

Save the changes and close the forms.

END OF STEPS

79.20 To configure unequal ECMP on a VPRN site

79.20.1 Purpose

Perform this procedure to apply a weighted ECMP distribution of flows for EVPN-IFL and EVPN-IFF routes.

79.20.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Segment Routing V6 tab, then on the BGP-EVPN tab.

5

Create or select an entry and click Properties. The SRv6 BGP EVPN (Create/Edit) form opens.

6

Configure the parameters in the EVPN Link Bandwidth panel, as required.

7

Save the changes and close the forms.

END OF STEPS

79.21 To enable routing protocols on a VPRN site

79.21.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab, then the Protocols tab.

5

Enable the required protocol parameters.

When you choose a protocol, the site and the protocols that are enabled on the site appear in the list panel.

Note: If both the OSPFv2 Enabled parameter and the OSPFv3 Enabled parameter are disabled, navigate to the Components view of the VPRN, right-click on Protocols, then select Create OSPF Site. The OSPF Site (Create) form opens. In the OSPF Instance panel, configure the Version parameter.

6

To configure multicast for the site, click on the Multicast tab and enable the required protocol parameters.

When you choose a protocol, the site and the protocols that are enabled on the site appear in the list panel.

7

Save and close the forms.

END OF STEPS

79.22 To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance

Note: To configure BGP on a VPRN routing instance, see 79.23 “To configure BGP on a VPRN routing instance” (p. 2557) .

To configure IGMP on a VPRN routing instance, see 79.24 “To configure IGMP on a VPRN routing instance” (p. 2560).

To configure MSDP on a VPRN routing instance, see 79.25 “To configure MSDP on a VPRN routing instance” (p. 2562).

79.22.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the VPRN service on which you want to configure the protocols and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Routing Instance→Routing Instance→Protocols.

4

Right-click on the Protocols icon and choose:

a. Create ISIS Site to configure ISIS on the VPRN routing instance, and see 28.60 “To configure IS-IS on a routing instance” (p. 953) .

b. Create L2TP Site to configure L2TP on the VPRN routing instance, and see 28.89 “To configure L2TP on a routing instance” (p. 989) .

c. Create OSPF Site to configure OSPFv2 on the VPRN routing instance, and configure the parameters as described in 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969) .

d. Create OSPF Site to configure OSPFv3 on the VPRN routing instance, and configure the parameters as described in 28.72 “To configure OSPF on a default routing instance or a VRF routing instance” (p. 969) .

e. Create RIP Site to configure RIP on the VPRN routing instance, and see 28.45 “To configure global-level RIP or RIPng” (p. 937) .

f. Create WPP Site to configure WPP on the VPRN routing instance, and see 28.134 “To create a web portal routing instance” (p. 1051) .

g. Create PIM Site to configure PIM on the VPRN routing instance, and see 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003) .

h. Create MLD Site to configure MLD on the VPRN routing instance, and see 28.121 “To enable MLD on a base routing instance” (p. 1034) .

Note: You must ensure that the VPRN routing instance has an associated L3 access interface, and that the IPv6 Allowed parameter on the interface is enabled. You must add an IPv6 address to the L3 access interface. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) and 79.102 “To assign an IP address to a VPRN L3 access interface” (p. 2677) .

5

Save the changes and close the forms.

END OF STEPS

79.23 To configure BGP on a VPRN routing instance

79.23.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance, right-click on the site on which you need to configure BGP and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab.

5

Enable the BGP Enabled parameter on the Protocols tab. A list of BGP routing instances appears.

6

Select a routing instance and click Properties. The BGP Site, routing instance (Create) form opens.

7

Configure the required parameters.

8

To assign a TCP key chain, select a TCP key chain in the KeyChain panel.

Note: You can assign a TCP key chain to a BGP site, group, or peer on a 7450 ESS in mixed mode, 7705 SAR, 7750 SR, or a 7950 XRS.

9

Configure the parameters in the Graceful Restart panel.

The graceful restart parameters are configurable when you enable the Graceful Restart parameter.

10

Configure the parameters in the MED panel.

These parameters are used to find a way to leave the AS when there are multiple methods of leaving the AS.

11

Configure the parameters in the MultiPath panel.

Set the Multi Path parameter to 1 to disable. When set from 2 to 16, multipath is enabled and BGP load shares traffic across the number of links specified. If the equal cost routes available are greater than the configured value, then routes with the lowest next hop IP address are chosen.

12

Configure the EIBGP Load Balance parameter in the EIBGP LoadBalance panel.

13

Configure the Backup Path parameter in the Backup Path panel.

14

Configure the parameters in the Best Path Selection panel.

When you set the MED Compare parameter to off, the Strict parameter is set to true and is not configurable. When you set the MED Compare parameter to on, the Strict parameter is set to false and is not configurable.

15

Click on the Behavior tab and configure the required parameters.

16

Configure the parameters in the Hold Time panel.

17

Configure the Fault Tolerance parameter in the Error Handling panel. The Fault Tolerance parameter allows BGP routers to be more tolerant of some Update message errors, use less disruptive error recovery mechanisms, and provide better operational and diagnostics information.

18

Configure the parameters in the Damp Peer Oscillations panel.

19

Click on the AS Properties tab and configure the Disable 4Byte ASN parameter.

20

Configure the parameters in the Local AS panel.

The No Prepend Global AS is configurable only if you set the Local AS parameter to a value other than 0.

The Local AS parameters are used to configure a virtual AS. A virtual AS is used when a router (RTA) is moved from one AS (AS1) to another AS (AS2). However, the customer router (CR1) is configured to belong to the AS1. To avoid reconfiguring CR1 to belong to AS2, CR1 can continue to belong to AS1, but RTA has its local AS value set to AS1. RTA can advertise AS1 for routes advertised to CR1.

21

Configure the parameters in the Remove Private AS panel.

22

Click on the VPN tab and configure the required parameters.

Note: The Apply Import Route Policies and Apply Export Route Policies parameters specify whether to apply the existing import and export route policies configured on the Import Policies and Export Policies tabs.

23

To add a group:

1. Click on the Groups tab.
2. Click Create. The Peer Group (Create) form opens.
3. Perform Step 3 to Step 23 of 28.32 “To configure peer-group-level BGP” (p. 922) .

24

To add a peer:

1. Click on the Peers tab.
2. Click Create. The Peer (Create) form opens.
3. Perform Step 3 to Step 19 of 28.33 “To configure peer-level BGP” (p. 927) .

25

Click on the Next-Hop Resolution tab and configure the Policy parameter.

Configure the parameters in the Use Leaked Routes panel.

26

Click on the Import Policies tab and configure the required parameters.

Configure the import route policies to determine which routes are accepted from peers. These policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . The router does not validate to ensure that the policies match.

27

Click on the Export Policies tab and configure the required parameters.

Configure the export route policies to determine which routes are advertised to peers. These policies should match the policies you configure using the Routing Policy Manager, as described in Chapter 27, “NE routing and forwarding” . The router does not validate to ensure that the policies match.

28

Click on the RIB Management tab to assign leak import and export policies:

1. On the IPv4 Leak Import Policies, IPv6 Leak Import Policies, Label IPv4 Leak Import Policies, and Label IPv6 Leak Import Policies tabs, assign up to 15 leak import policies to the BGP site as required.
2. Click on the IPv4 Leak Export Policies tab and assign up to 15 leak export policies to the BGP site.
3. Click on the Route Table Import tab and configure the required parameters.

29

To configure interface instance, perform the following steps:

1. Click on the Dynamic Peer Interfaces tab.
2. Create or select an entry and click Properties. The Dynamic Peer Interface (Create/Edit) form opens.
3. Configure the required parameters in General tab.
4. Click on the AS Allowed Ranges tab and configure the Minimum and Maximum values.

   Note:

   Upto 32 AS Allowed Range can be configured.
5. Save the changes and close the form.

30

Click on the Authentication tab and configure the required parameters.

31

Save the changes and close the forms.

END OF STEPS

79.24 To configure IGMP on a VPRN routing instance

Note: PIM-SSM for IPv6 is not supported in VPRN services.

79.24.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Routing Instance→Routing Instance→Protocols, right-click on Protocols and choose Create IGMP Site. The IGMP Site (Edit) form opens.

4

Configure the required general parameters.

5

To configure SSM:

1. Click on the SSM Translation tab.
2. Click Create to create a new entry. The SSM Translation (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

6

To add an interface:

1. Click on the Interfaces tab.
2. Click Create. The IGMP Interface (Create) form open.
3. Select an interface and configure the required parameters.
4. Save the changes and close the form.

7

To identify an IGMP group interface for a VPRN service:

1. Click on the Group Interfaces tab
2. Click Create. The IGMP Group Interface (Create) form opens.
3. Configure the required parameters.
4. Select an IGMP group interface in the IGMP Group Interface panel and configure the required parameters.

5. To configure behavior, click on the Behavior tab, select an import policy and configure the required parameters.
6. To configure multicast CAC, click on the Multicast CAC tab, select a multicast CAC policy and configure the required parameters.
7. Save the changes and close the form.

8

Save the changes and close the forms.

END OF STEPS

79.25 To configure MSDP on a VPRN routing instance

79.25.1 MSDP on VPRN Service object

You can configure MSDP on the VPRN Service object. The configuration forms are identical to those used to configure MSDP on the routing instance, and the behavior of the VPRN MSPD is the same as that of the base router MDSP.

Note: You can enable and configure MSDP without PIM. However, you need to configure PIM before MSDP can be operationally up. At a minimum, you need to configure the PIM Candidate BootStrap Router and Candidate Rendezvous Point (C-RP) on a pair of NEs. The MSDP peer is the C-RP Address parameter. See 28.99 “To create a PIM site on a VPRN routing instance” (p. 1003) for details.

79.25.2 Before you begin

Before you configure MSDP on a VPRN routing instance, you need to enable MSDP:

1. On the service navigation tree, expand VPRN Site→Routing Instance→Routing Instance *name*, right-click on Routing Instance *name* and choose Properties. The VPRN Site (Edit) form opens.
2. Click on the Multicast tab, and select the MSDP Enabled parameter.
3. Save your changes and close the form.

79.25.3 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand VPRN Site→Routing Instance - NE System ID→Routing Instance→Routing Instance *VPRNinstance*→Protocols, right-click on Protocols and choose Create MSDP Site. The MSDP Site (Create) form opens.

4

Configure the required parameters in the General panel.

5

Configure the parameters in the Receive Message Rate Parameters panel.

6

Configure the State parameter.

7

Configure an MSDP peer group.

1. Click on the Group tab and then click Create. The MSDP Peer Group (Create) form opens.
2. Configure the required parameters on the General tab.
3. Click on the Import Policies and configure the import route policies to determine which routes are accepted from peers.

   The policies need to match the policies that you set when configuring the Routing Policy Manager. The router performs no validation to ensure that the policies match.
4. Click on the Export Policies tab and configure the export route policies to determine which routes are advertised to peers.

   The policies need to match the policies that you set when configuring the Routing Policy Manager. The router performs no validation to ensure that the policies match.
5. Save the changes and close the form.

You can also configure an MSDP peer group separately, after you completed the configuration of the MSDP site. On the service navigation tree, expand VPRN Site→Routing Instance - NE System ID→Routing Instance→Routing Instance *VPRNinstance*→Protocols→MSDP, right-click on MSDP and choose Create MSDP Peer Group. The MSDP Peer Group (Create) form opens.

8

Configure an MSDP peer.

1. Click on the Peer tab and then click Create. The MSDP Peer (Create) form opens.
2. Configure the required parameters on the General tab.
3. Click on the Import Policies and configure the import route policies to determine which routes are accepted from peers.

   The policies need to match the policies that you set when configuring the Routing Policy Manager. The router performs no validation to ensure that the policies match.

4. Click on the Export Policies tab and configure the export route policies to determine which routes are advertised to peers.

   The policies need to match the policies that you set when configuring the Routing Policy Manager. The router performs no validation to ensure that the policies match.
5. Click on the Authentication tab and configure the required parameters.
6. Save the changes and close the form.

You can also configure an MSDP peer separately, after you completed the configuration of the MSDP site. On the service navigation tree, expand VPRN Site→Routing Instance - NE System ID→Routing Instance→Routing Instance *VPRNinstance*→Protocols→MSDP, right-click on MSDP and choose Create MSDP Peer. The MSDP Peer (Create) form opens.

9

Click on the Source tab to create an MSDP multicast source.

1. Click Create. The MSDP Source (Create) form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

10

Click on the following tabs to review MSDP Site information, as required:

- Data Source Active
- Data Source Active Rejected
- Statistics
- Faults

11

Save the changes and close the form.

END OF STEPS

79.26 To configure a routing instance on a VPRN site

Note: The tabs and parameters that are configurable vary depending on the NE.

79.26.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab.

5

Configure the required parameters.

You cannot configure the Carrier Carrier VPN parameter on a VPRN site if interfaces are configured on the site.

When you set the Route Distinguisher Type parameter to Type 0, the following configurable parameters appear:

- Type 0 Administrative Value
- Type 0 Assigned Value

You can click on Suggest Value to let the NFM-P assign these values. Choose Generate Unique RD from the drop-down menu.

The Type 0 Administrative Value parameter is the Autonomous System (AS) number of the PE node. For the Type 0 Assigned Value parameter, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

When you set the Route Distinguisher Type parameter to Type 1, the following configurable parameters appear:

- Type 1 IP Address
- Type 1 Assigned Value

When you set the Route Distinguisher Type parameter to Type 2, the following configurable parameters appear:

- Type 2 Administrative Value
- Type 2 Assigned Value

If you enabled the Enable RT/RD Auto Assignment parameter when you created the VPRN service, the NFM-P automatically configures the Route Distinguisher parameter with the values in the VPRN RT/RD Auto Assignment policy. See 79.51 “To configure a policy to reserve an RT and RD range for VPRN services” (p. 2616) for information about how to create a VPRN RT/RD Auto Assignment range policy.

6

To configure importing of leaked GRT routes, add routing policies in the GRT Import Policies sub-tab.

7

To specify whether SNMP requests are allowed using the VPRN interface, configure the Allow SNMP Access parameter.

8

Configure the parameters in the GRT panel.

9

Configure the Enforce Maximum Number Of Routes parameter in the Maximum Number Of Routes panel.

Note: When you select the Enforce Maximum Number Of Routes parameter, the following parameters are configurable:

- Maximum Number Of Routes
- Log Only
- Threshold (%)

10

Configure the Enforce Maximum Number Of IPv6 Routes parameter in the Maximum Number Of IPv6 Routes panel.

Note: When you select the Enforce Maximum Number Of IPv6 Routes parameter, the following parameters are configurable:

- Maximum Number Of IPv6 Routes
- Log Only
- Threshold (%)

11

Configure the Enable Backup BGP-VPN Routes parameter in the Backup BGP-VPN Routes panel.

12

Configure the parameters in the D-Path panel.

13

Configure the Single SFM Overload Admin State parameter in the Single SFM Overload panel.

The Hold-Off Time (seconds) parameter is configurable when the Single SFM Overload Admin State parameter is set to Up.

14

Configure the parameters in the Static Route Hold-Down Time panel.

The Initial Hold-Down Time, Hold-Down Time Multiplier, and Maximum Hold-Down Time parameters are configurable when you enable the Enable Hold-Down Time parameter.

15

Configure the IPv4 Max Size and IPv6 Max Size parameters in the Embedded Flowspec panel, as required.

16

Configure the Enforce Maximum Number Of Multicast Routes parameter in the Maximum Number Of Multicast Routes panel.

Note: When you enable the Enforce Maximum Number Of Multicast Routes parameter, the following configurable parameters appear:

- Maximum Number Of Multicast Routes
- Log Only
- Threshold (%)

17

To configure ingress multicast forwarding:

1. Click on the Mcast Path Mgmt tab.
2. Select an ingress info policy.

   Note:

   The Mcast Path Mgmt Channels tab displays data on the operational channels after actual traffic from a specific multicast source for a specific multicast group passes through the virtual router. You must click Search to refresh the data. See Chapter 49, “Policies overview” for a listing of the displayed operational channel parameters.

18

Save and close the forms.

END OF STEPS

79.27 To configure a VRF instance on a VPRN site

79.27.1 Purpose

Route targets are used to identify the VRFs of a VPRN. A PE router that is not a route reflector or an AS border router installs a VPRN route only when its import target matches the target of the route.

A fully-meshed VPRN requires one target for all participating VRFs. A hub-and-spoke VPRN requires VRF import and export targets. The export target of the hub VRF must be the same as the import target of the spoke VRFs. The import target of the hub VRF must be the same as the export target of the spoke VRF. VPRN VRF targets must not overlap.

79.27.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the VRF Target tab and configure the VRF Target Type parameter.

Set the VRF Target Type parameter to Define Default if you need to specify a default VRF target for the site. Otherwise, go to Step 6 .

5

To specify a two-byte AS number for the default target, set the Target Format parameter to AS and configure the required parameters.

a. You can click Suggest Value to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

   The Target AS Value parameter is the Autonomous System number of the PE NE. For the Target Extended Community Value parameter, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

b. To specify an IP Address for the default target, set the Target Format parameter to IP Address and configure the required parameters.

c. To specify a four-byte AS number for the default target, set the Target Format parameter to AS-4Byte Address and configure the required parameters.

   You can click Suggest Value to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

   The Target AS Value (4Byte) is the Autonomous System number of the PE NE. For the Target Community Value parameter, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

6

Set the VRF Target Type parameter to Define Import and Export if you need to specify import and export VRF targets for the service site.

7

To specify no VRF import target format for the site, set the Import Target Format parameter to None and go to Step 10 . Otherwise, go to Step 8 .

8

To specify a two-byte import target format for the site, set the Import Target Format parameter to AS and configure the required parameters. Otherwise, go to Step 9 .

You can click on the Suggest Value button to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

The Import Target AS Value parameter is the two-byte Autonomous System number of the PE node. For the Import Target Extended Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

9

To specify an IP Address as the import target, set the Import Target Format parameter to IP Address and configure the required parameters. Otherwise, go to Step 11 .

10

To specify a four-byte AS number for the import target format, set the Import Target Format parameter to AS-4Byte Address and configure the required parameters.

You can click on the Suggest Value button to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

The Import Target AS Value (4Byte) is the four-byte Autonomous System number of the PE node. For the Import Target Community Value parameter, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

11

To specify no VRF export target format for the site, set the Export Target Format parameter to None and go to Step 15 . Otherwise, go to Step 12 .

12

To specify a two-byte export target format for the site, set the Export Target Format parameter to AS and configure the required parameters. Otherwise, go to Step 13 .

You can click on the Suggest Value button to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

The Export Target AS Value parameter is the two-byte Autonomous System number of the PE node. For the Export Target Extended Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

13

To specify an IP Address as the export target, set the Export Target Format parameter to IP Address and configure the required parameters. Otherwise, go to Step 14 .

14

To specify a four-byte AS number for the export target format, set the Export Target Format parameter to AS-4Byte Address and configure the required parameters.

You can click on the Suggest Value button to let the NFM-P assign these values. Choose Generate Unique VRF Target from the drop-down menu.

The Export Target AS Value (4Byte) is the four-byte Autonomous System number of the PE node. For the Import Target Community Value parameter, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

15

Save and close the forms.

END OF STEPS

79.28 To configure an MVPN VRF instance on a VPRN site

79.28.1 Before you begin

For 7750 SR, 7450 ESS, and 7950 XRS NEs starting in Release 16.0 R1, PIM must be enabled before MVPN can be configured; see 79.21 “To enable routing protocols on a VPRN site” (p. 2554). If PIM is disabled, MVPN configuration is removed.

79.28.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

If the site you are configuring is a 7210 SAS-R, you must configure the device properties. See 12.55 “To configure 7210 SAS-R device properties for MVPN” (p. 386).

4

Click on the Multicast tab, the MVPN tab, then the Routing tab.

5

Configure the MVPN VRF Target Type parameter:

Note: Route targets are used to identify the VRFs of a VPRN. A PE router that is not a route reflector or an AS border router installs a VPRN route only when its import target matches the target of the route.
A fully-meshed VPRN requires one target for all participating VRFs. A hub-and-spoke VPRN requires VRF import and export targets. The export target of the hub VRF must be the same as the import target of the spoke VRFs. The import target of the hub VRF must be the same as the export target of the spoke VRF. VPRN VRF targets must not overlap.

a. If you do not need to specify a VRF target for the site, set the MVPN VRF Target Type parameter to None. Go to Step 9 .

b. If you need to specify a default VRF target for the site, set the MVPN VRF Target Type. Go to Step 6 .

c. If you need to specify import and export VRF targets for the service site, set the MVPN VRF Target Type parameter to Define Import and Export. Go to Step 7 .

6

Configure the Target Format parameter by performing one of the following steps:

a. To specify a two-byte AS number for the default target:
   1. Configure the required parameters.
      - Target AS Value
      - Target Extended Community Value

      The Target AS Value is the Autonomous System number of the PE node. For the Target Extended Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.
   2. Go to Step 11 .

b. To specify an IP Address for the default target:
   1. Configure the required parameters.
      - Target IP Address
      - Target Community Value
   2. Go to Step 11 .

c. To specify a four-byte AS number for the default target:
   1. Configure the required parameters.
      - Target AS Value (4Byte)
      - Target Community Value

      The Target AS Value (4Byte) is the Autonomous System number of the PE node. For the Target Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.
   2. Go to Step 11 .

7

Configure the Import Target Format parameter by performing one of the following steps:

a. Choose None to specify no VRF import target format for the site, then go to Step 8 .

b. Choose AS as the two-byte import target format for the site.

1. Configure the required parameters.
   - Import Target AS Value
   - Import Target Extended Community Value

   The Import Target AS Value is the two-byte Autonomous System number of the PE node. For the Import Target Extended Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.
2. Go to Step 8 .

c. Choose IP Address as the import target format for the site.

1. Configure the required parameters.
   - Import Target IP Address
   - Import Target Community Value
2. Go to Step 8 .

d. Choose AS-4Byte as the four-byte import target format for the site.

1. Configure the required parameters.
   - Import Target AS Value (4Byte)
   - Import Target Community Value

   The Import Target AS Value (4Byte) is the four-byte Autonomous System number of the PE node. For the Import Target Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.
2. Go to Step 8 .

8

Configure the Export Target Format parameter by performing one of the following steps:

a. Choose None to specify no VRF export target format for the site, then go to Step 9 .

b. Choose AS as the two-byte export target format for the site.

1. Configure the required parameters.
   - Export Target AS Value
   - Export Target Extended Community Value

   The Export Target AS Value parameter is the two-byte Autonomous System number of the PE node. For the Export Target Extended Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.

2. Go to Step 9 .

c. Choose IP Address as the export target format for the site.

1. Configure the required parameters.
   - Export Target IP Address
   - Export Target Community Value
2. Go to Step 9 .

d. Choose AS-4Byte as the four-byte export target format for the site.

1. Configure the required parameters.
   - Export Target AS Value (4Byte)
   - Export Target Community Value

   The Export Target AS Value (4Byte) parameter is the four-byte Autonomous System number of the PE node. For the Export Target Community Value, the NFM-P assigns the next unused number within the valid range. When all of the numbers within the valid range have been used once, the system starts at the beginning of the range and assigns the next available number.
2. Go to Step 9 .

9

To configure import policies:

1. Click on the Import Policies tab and set the Import Unicast parameter to False.
2. Select the required unicast import policies 1 through 15. You can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression, as described in 28.36 “To create a BGP policy expression” (p. 931).

10

To configure export policies:

1. Click on the Export Policies tab and set the Export Unicast parameter to False.
2. Select the required unicast export policies 1 through 15. You can use the Select button to choose a single policy, or click on the Create Expression button to use a logical expression, as described in 28.36 “To create a BGP policy expression” (p. 931).

11

Save and close the forms.

END OF STEPS

79.29 To configure a BGP confederation on a VPRN site

79.29.1 Purpose

A BGP confederation allows for the full distribution of external routing information within an AS.

You can configure BGP confederations only when:

- no VRF import or export policies are configured on the VPRN site
- no VRF target is configured on the VPRN site
- no VPRN GRT lookup or export policies are configured on the VPRN site

You cannot configure VRF import or export policies, VRF target, or GRT lookup or export policies when a BGP confederation is configured.

79.29.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab, then on the General tab.

5

Configure the Confederation Autonomous System parameter. When you set the Confederation Autonomous System to a value other than 0, the BGP Confederations tab appears.

6

To configure one or more BGP confederations members:

1. Click on the BGP Confederations tab and click Create. The Confederation (Create) form opens.
2. Click on the Members tab and click Create. The Confederation Member (Create) form opens.
3. Configure the Member AS parameter.
4. Save the changes and close the forms.

7

Save the changes and close the forms.

END OF STEPS

79.30 To configure a GSMP group on a VPRN site

79.30.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Multicast tab, then on the GSMP tab.

5

Click Create. The GSMP Group (Create) form opens.

6

Configure the required parameters.

7

To configure ANCP, configure the parameters in the ANCP panel.

8

To configure one or more GSMP group neighbours:

1. Click on the GSMP Group Neighbor tab and click Create. The GSMP (Create) form opens.
2. Configure the required parameters.

   The Priority Dscp parameter is configurable when you set the Priority Type to Dscp.

   The Priority Precedence parameter is configurable when you set the Priority Type to Precedence.
3. Save the changes and close the form.

9

Save the changes and close the forms.

END OF STEPS

79.31 To configure a GNE site and GNE service interface on a VPRN service

79.31.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, right-click on VPRN Service and choose Create GNE Site.

4

Select a site.

5

Configure the required parameters.

6

Click Apply.

7

To configure an interface for the GNE site:

1. Click on the GNE Service Interfaces tab and click Create. The GNE Service Interface (Create) form opens.
2. Configure the required parameters.
3. Click on the Ports tab and select a Generic NE Interface.
4. Configure the required parameters.
5. Save the changes and close the form.

8

Save the changes and close the forms.

END OF STEPS

79.32 To configure a RADIUS proxy server on a VPRN site

79.32.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the RADIUS/DHCP/Diameter tab, then on the RADIUS tab.

5

On the RADIUS Proxy Servers tab, click Create or select a RADIUS proxy server and click Properties. The RADIUS Proxy Server (Create|Edit).

6

Configure the parameters as required.

7

Select a Python Policy, if required.

If you intend to configure a WLAN GW group with a distributed Python policy in Step 8, the Python policy you select in this step must be the same Python policy that is distributed to the WLAN GW group.

8

Click Select in the WLAN GW panel to assign a ISA-WLAN GW Group, if required.

9

Configure default RADIUS policies:

Perform the following steps:

1. Select a RADIUS authentication policy.
2. Select a RADIUS accounting policy.

   Note:

   If you configure a ISA-WLAN GW group in Step 8, the default RADIUS policies must be a ISA RADIUS policy.

10

Click on the Caching tab and configure the parameters as required.

11

Click on the Interfaces tab and click Create or select an existing network Interface and click Properties. The RADIUS Proxy Interface form opens.

Perform the following steps:

1. Click Select and choose an interface from the Select Interface form.
2. Select an interface and click OK. The RADIUS Proxy Server (Create|Edit) form reappears.

12

Click on the Attribute Matching/Users tab to configure RADIUS attribute matching type and matching entries for RADIUS policies.

Perform the following steps:

1. On the Type tab, configure the Type and Vendor ID parameters.
2. Click on the Entry tab.
3. Click Create to create a new RADIUS attribute matching entry, or select an existing entry in the list and click Properties. The Attribute Matching Entry form appears.
4. Configure the Entry ID, Prefix Match String, and Suffix Match String parameters as required.
5. Select a RADIUS authentication policy and a RADIUS accounting policy.
6. Save your changes and close the form.

13

Save the changes and close the forms.

END OF STEPS

79.33 To configure a local DHCPv4 server on a VPRN site

79.33.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the RADIUS/DHCP/Diameter tab, then on the Local DHCP Servers tab.

5

On the DHCPv4 tab, click Create or select a server and click Properties. The Local DHCP Server, Site (Create|Edit).

6

Configure the required general parameters.

7

Configure the required parameters in the Client Pool panel.

8

Select a local user database in the Local User Database panel.

9

To configure DHCP failover support:

1. Click on the Failover tab.
2. Configure the Administrative State parameter.
3. Configure the parameters in the Client Lead Time panel.
4. Configure the parameters in the Startup Wait Time panel.
5. Configure the parameters in the Partner Down Delay panel.
6. To configure a failover DHCP server peer, click on the Peer tab.
7. Click Create. The Failover Peer, Site (Create) form opens.
8. Select a failover MC peer site.

9. Configure the Sync Tag parameter. The Sync Tag setting must be the same on both failover servers.
10. To configure node redundancy for the failover peer, click Node Redundancy. The Manage Node Redundancy form opens, pre-configured for MC peer groups.
11. Select an MC peer group, or click Create to configure a new MC peer group. See 40.4 “To configure an MC peer group” (p. 1330) . The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add failover peer sync tags.
12. Save the changes and close the forms.

10

To create a pool of one or more subnets on the local DHCP server:

1. Click on the IP Address Pools tab and click Create. The IP Address Pool form opens.
2. Configure the required general parameters.
3. Configure the parameters in the Minimum Lease Time panel.
4. Configure the parameters in the Maximum Lease Time panel.
5. Configure the parameters in the Offer Time panel.
6. Configure the parameters in the Pool Notifications (Minimum Free) panel.
7. Save the changes and close the form.

11

To add a subnet to the IP address pool:

1. Click on the Subnets tab click Create. The Subnet, Site (Create) form opens.
2. Configure the required general parameters.
3. Configure the required parameters in the Pool Notifications (Minimum Free) panel.
4. Click on the Address Ranges tab and click Create. The Subnet Address Range form opens.
5. Click Create. The Subnet Address Range, Site (Create) form opens.
6. Configure the parameters.

   You must exclude static IP addresses from the subnet address range because static IP addresses are dedicated.
7. Save the changes and close the form.
8. Click on the Options tab and click Create. The Subnet Option, Site (Create) form opens.
9. Configure the Option parameter and configure the required parameters.

   The Number parameter is configurable when the Option parameter is set to Custom Option.

   The Value parameter is configurable when the Type parameter is set to ASCII String or Hex String.

   The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
10. Save the changes and close the forms.

12

To configure IP address pool options, click on the Options tab and perform the following:

1. Click Create. The IP Address Pool Option, Site (Create) form opens
2. Configure the Option parameter and configure the required parameters.

   The Number parameter is configurable when the Option parameter is set to Custom Option.

   The Value parameter is configurable when the Type parameter is set to ASCII String or Hex String.

   The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
3. Save the changes and close the IP Address Pool Option and IP Address Pool forms.

13

To add a sticky lease to an IP address pool, perform the following:

1. Click on an IP address pool entry and click Create Sticky Lease. The Sticky Lease Action (Create) form opens.
2. Configure the required parameters.

   The Host Name parameter value must be unique. The MAC Address can be duplicated across multiple sticky lease objects, provided that each has a different corresponding Circuit ID value. Similarly, the Circuit ID value can be duplicated across multiple sticky lease objects, provided that each has a different corresponding MAC Address.

14

Save the changes and close the forms.

END OF STEPS

79.34 To configure a local DHCPv6 server on a VPRN site

Note: The tabs and parameters that are configurable vary depending on the NE.

79.34.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the RADIUS/DHCP/Diameter tab, then on the Local DHCP Servers tab.

5

On the DHCPv6 tab, click Create or select a server and click Properties. The Local DHCPv6 Server, Site (Create|Edit).

6

Configure the required general parameters.

7

Configure the required parameters in the Client Pool panel.

8

Configure the required parameters in the Lease Hold Time panel.

9

Configure the required parameters in the DHCP Unique Identifier (DUID) panel.

10

To configure DHCPv6 failover support:

1. Click on the Failover tab.
2. Configure the Administrative State parameter.
3. Configure the parameters in the Client Lead Time panel.
4. Configure the parameters in the Startup Wait Time panel.
5. Configure the parameters in the Partner Down Delay panel.
6. To configure a failover DHCPv6 server peer, click on the Peer tab.
7. Click Create. The Failover Peer, Site (Create) form opens.
8. Select a failover MC peer site.
9. Configure the Sync Tag parameter. The Sync Tag setting must be the same on both failover servers.
10. To configure node redundancy for the failover peer, click Node Redundancy. The Manage Node Redundancy form opens, pre-configured for MC peer groups.
11. Select an MC peer group, or click Create to configure a new MC peer group. See 40.4 “To configure an MC peer group” (p. 1330) . The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add failover peer sync tags.

12. Save the changes and close the forms.

11

To configure an IPv6 address pool:

1. Click on the IPv6 Address Pools tab click Create. The IPv6 Address Pool, Site (Create) form opens.
2. Configure the required general parameters.

12

To configure IPv6 address pool failover:

1. Click on the Failover tab, then the General tab.
2. Configure the Administrative State parameter.
3. Configure the parameters in the Client Lead Time panel.
4. Configure the parameters in the Startup Wait Time panel.
5. Configure the parameters in the Partner Down Delay panel.

13

To configure IPv6 address pool failover peer:

1. To configure a failover peer, click on the Peer tab.
2. Click Create. The Failover Peer, Site (Create) form opens.
3. Select a failover MC peer site.
4. Configure the Sync Tag parameter. The Sync Tag setting must be the same on both failover servers.
5. To configure node redundancy for the failover peer, click Node Redundancy. The Manage Node Redundancy form opens, pre-configured for MC peer groups.
6. Select an MC peer group, or click Create to configure a new MC peer group. See 40.4 “To configure an MC peer group” (p. 1330) . The Node Redundancy button allows you to access the MC sync group configuration within the MC peer group in order to add failover peer sync tags.
7. Save the changes and close the form.

14

To add prefixes to the IP address pool:

1. Click on the Prefixes tab.
2. Click Create. The Prefix, Site (Create) form opens.
3. Configure the required general parameters.
4. Configure the parameters in the Preferred Life Time panel.
5. Configure the parameters in the Valid Life Time panel.

6. Configure the parameters in the Renew Timer panel.
7. Configure the parameters in the Rebind Timer panel.
8. Click on the Options tab.
9. Click Create. The Prefix Option, Site (Create) form opens.
10. Configure the Option parameter and configure the required parameters.

    The Number parameter is configurable when the Option parameter is set to Custom Option.

    The Value parameter is configurable when the Type parameter is set to ASCII String, Hex String, or Domain.

    The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
11. Save your changes and close the form.
12. Click on the Minimum Free Thresholds tab.
13. Click Create or select an existing minimum free threshold entry and click Properties. The Prefix Minimum Free Threshold (Create|Edit) form opens.
14. Configure the required parameters:

    The Minimum Threshold and Minimum Number parameters are mutually exclusive. They cannot both be configured at the same time.
15. Save your changes and close the forms.

15

To configure IPv6 address pool options:

1. Click on the Options tab.
2. Click Create. The IP Address Pool Option, Site (Create) form opens.
3. Configure the Option parameter and configure the required parameters.

    The Number parameter is configurable when the Option parameter is set to Custom Option.

    The Value parameter is configurable when the Type parameter is set to ASCII String, Hex String, or Domain.

    The IP Address 1, IP Address 2, IP Address 3, and IP Address 4 parameters are configurable when the Type parameter is set to IP Address.
4. Save the changes and close the form.

16

To configure IPv6 address pool minimum free thresholds:

1. Click on the Minimum Free Thresholds tab.
2. Click Create or select an existing minimum free threshold entry and click Properties. The Address Pool Minimum Free Threshold (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the form.

17

Save the changes and close the forms.

END OF STEPS

79.35 To clear DHCP leases from a VPRN site local DHCP server

79.35.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the RADIUS/DHCP/Diameter tab, then on the Local DHCP Servers tab.

5

On the DHCPv4 tab, select a server and click Properties. The Local DHCP Server, Site (Edit) form opens.

6

Click on the Leases tab.

7

Click on the Regular tab to clear DHCP leases, and then do the following:

1. Select the DHCP lease you want to clear and click Clear DHCP Lease.
2. If you want to clear all of the leases on the local DHCPv4 server, click Clear All DHCP Leases.

8

Click on the Sticky tab to clear sticky leases, and then do the following:

1. Select the sticky lease you want to clear and click Clear Sticky Lease.
2. If you want to clear all sticky leases with a specific host name prefix, click Clear Sticky Lease With Host Name Prefix.

3. Specify a host name prefix and click Ok.

9

Confirm the deletion(s) and close the forms.

END OF STEPS

79.36 To configure TCP MSS adjustment on a VPRN site

Note: Perform this procedure to specify an ISA-BB group and a maximum segment size on a VPRN site that is to be part of a TCP MSS adjustment configuration. For more information, see 27.1.15 “TCP MSS adjustment for PPPoE sessions” (p. 822) and 27.1.16 “Workflow to configure TCP MSS adjustment” (p. 823).

79.36.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

On the General tab, select an ISA-BB group and configure the TCP Maximum Segment Size parameter.

5

Save your changes and close the forms.

END OF STEPS

79.37 To configure a group interface on a VPRN

79.37.1 Prerequisites

The 7450 ESS in mixed mode, 7750 SR supports the configuration of a group interface in a VPRN.

79.37.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces, right-click on a subscriber interface and choose Create VPRN Group Interface, or expand the subscriber interface, right-click on a group interface and choose Properties. The VPRN Group Interface (Create|Edit) form opens.

4

Configure the required general parameters.

The LNS parameter defines the type of group interface (regular or LNS). This parameter is set at creation time and cannot be modified. Regular group interfaces cannot configure LNS attributes and an LNS group interface does not allow PPPoE configuration or SAPs.

The Unicast RPF configuration is intended to enhance security for a MAC-only anti-spoofing configuration on the WLAN GW for soft GRE group interfaces; see 79.66 “To configure a WLAN GW for a VPRN group interface” (p. 2631) .

5

To configure cflowd sampling:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click Apply. The Cflowd sampling object appears in the Cflowd Sampling panel.
3. Save your changes.

6

Select a redundant interface.

7

Select a diameter application policy.

8

Select a diameter authentication policy.

9

Configure the required parameters in the IPv4/IPv6 panel. To configure IPv6:

1. Enable the IPv6 Allowed parameter and configure the required parameters.
2. Click on the IPv6 Advertisement tab and configure the required parameters.

   If the No Default Router check box is enabled, the Router Lifetime parameter is not configurable.
3. Configure the parameters in the Prefix Options panel.

   If the either of the Infinite check boxes is enabled, the associated parameter is set to its maximum value.
4. Configure the parameters in the DNS Options panel.

   If the Infinite check box is enabled, the RDNSS Lifetime parameter is set to -1.

10

To configure IPv6 router solicitation on the group interface:

1. Ensure the IPv6 Allowed parameter is enabled on the General tab.
2. Click on the IPv6 Router Solicit tab.
3. Configure the Administrative State parameter.
4. Select a local user database.
5. Configure the parameters in the Minimum Authentication Time panel.
6. Configure the parameters in the Inactivity Timer panel.

   If you enable the Infinite check box, the Inactivity Timer parameters are disabled.

11

To configure DHCPv6 on the group interface:

1. Ensure the IPv6 Allowed parameter is enabled on the General tab.
2. Click on the DHCPv6 tab and configure the required parameters.

   The Interface ID String parameter is configurable only if the Interface ID Option parameter is set to String.

   If the PD Managed Route parameter is enabled, the PD Managed Route Next Hop parameter is configurable.
3. Select a local user database.
4. Select a python policy.
5. Select a DHCPv6 filter.
6. Click on the Proxy Server tab and configure the required parameters.

   The ID Type and Value parameters are configurable only if the Server ID Type parameter is set to Enterprise.
7. Click on the Relay tab and configure the required parameters.

8. Configure the Server 1 through Server 8 parameters.
9. Select the interface name for each DHCPv6 server in the Zone Index panel.

   If you have entered a Unicast address, then the Interface Name parameter is not required.

12

To configure WPP on the group interface:

1. Ensure the Enable WPP parameter is enabled on the General tab.
2. Click on the WPP tab.
3. Configure the Administrative State parameter.
4. Select a WPP site in the WPP Site panel.
5. Configure the Portal Name parameter.
6. Select an Initial Subscriber Profile.
7. Select an Initial SLA Profile.
8. Select an Initial Application Profile.
9. Configure the Restore Default Profile On Disconnection parameter.
10. Select a Local User DB.
11. Configure the Enable Triggered Hosts parameter.
12. Disable the Default check box and configure the Lease Time (Days, Hours, Minutes, and Seconds) parameters, if required.

13

To configure anti-spoofing:

1. Click on the Anti-Spoofing tab.
2. Configure the ARP Populate parameter.

14

To configure subscriber management:

1. Click on the Subscriber Management tab.
2. Configure the required parameters.

   The parameters on the SHCV panel are configurable when the SHCV Enabled parameter is configured.
3. On the Subscriber Host Connectivity Verification panel, select an IPv4/IPv6 policy, IPv4 policy, and IPv6 policy, as required.

15

To configure ICMP for the group interface:

1. Click on the ICMP tab.

2. Configure the required parameters.

16

To configure ARP for the group interface:

1. Click on the ARP tab.
2. Configure the Timeout (seconds) parameter.
3. Click on the Proxy ARP tab.
4. Configure the required parameters.

17

To configure IPv4 DHCP relay for the group interface:

1. Click on the DHCP tab.
2. Configure the required general parameters.

   When the Relay Unicast Message parameter is set to Renew or Release Update Source IP, the GI address parameter can be configured as any local configured address in the same routing instance as the GI address for the DHCP relayed messages. If the Relay Unicast Message parameter is set to None, the GI address is restricted to the IP address configured on the subscriber interface.
3. Select a subscriber authentication policy in the Subscriber Authentication panel.
4. Select a local user database in the Local User Database panel.
5. Select a filter policy in the DHCP Filter panel.
6. Click on the Server tab and configure the required parameters.

   The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.
7. Click on the Client Applications tab and configure the Client Applications parameter.

18

To configure IPoE sessions for the group interface:

1. Click on the IPoE Session tab.
2. Configure the required parameters.

   The Stateless Redundancy parameter is configurable only on a regular (not soft GRE) group interface.

   Certain conditions and restrictions apply to an IPoE session on a soft GRE group interface:
   - RADIUS Session Timeout is set to Backwards Compatible by default
   - The IPoE session policy is set to a default policy (with SAP MAC address as key); cannot be changed
   - SAP Session Limit is set to 131071; cannot be changed
   - Administrative Sate is set to Enabled; cannot be changed.

3. Select an IPoE session policy.
4. Select a local user database.

19

To configure neighbor discovery on the group interface:

Neighbor discovery is configurable when the IPv6 Allowed parameter is enabled on the group interface.

1. Click on the Neighbor Discovery tab.
2. Configure the Maximum Number of Neighbors Learned and Duplicate Address Detection parameters.

20

To configure PPP for the group interface:

1. Click on the PPP tab.
2. Configure the Description and Administrative State parameters in the PPPoE panel.
3. Select a PPPoE policy in the PPPoE panel.
4. Select a PPPoE Local User DB in the PPPoE panel.
5. Select a python policy in the PPPoE panel.
6. Configure the required parameters in the PPPoE panel.
7. Configure the Description and Administrative State parameters in the PPP panel.
8. Select a PPP policy in the PPP panel.
9. Select a PPP Local User DB in the PPP panel.
10. Configure the Session Limit parameter in the PPP panel.

21

To configure IPoE linkage for the group interface:

1. Click on the IPoE Linkage tab.
2. Configure the required parameters.

22

To configure the ARP host for the group interface:

1. Click on the ARP Host Configuration tab.
2. Configure the required parameters.

23

To configure local address assignment for the group interface:

1. Click on the Local Address Assignment tab.

2. Configure the required parameters.
3. In the IPv4 Options panel, configure the PPP-v4 parameter and select a local DHCP server.
4. In the IPv6 Options panel, configure the IPoE SLAAC, IPoE WAN, and PPP SLAAC parameters and select a local DHCPv6 server.

24

To configure GTP for the group interface:

1. Ensure the Specific Type parameter is set to GTP on the General tab.
2. Click on the GTP tab.
3. Configure the Administrative State parameter.
4. Select a forward path extension.

25

To configure BRG for the group interface:

1. Click on the BRG Configuration tab.
2. Configure the Administrative State parameter.
3. Select a default BRG profile.
4. Enable the Authenticate BRG Only parameter, if only authenticated BRGs are permitted on the group interface.

26

To configure LNS for the group interface:

After you create an LNS group interface, you must configure the L2TP tunnel group profile or tunnel profile to terminate sessions for the LNS group interface that you just created; see 28.89 “To configure L2TP on a routing instance” (p. 989) . You can also configure the termination of sessions on a group interface using a RADIUS server.

1. Ensure that the LNS parameter is enabled on the General tab.
2. Click on the LNS tab.
3. Configure the Description parameter.
4. Select a subscriber profile in the Default Subscriber Profile panel.
5. Select an SLA profile in the Default SLA Profile panel.
6. Select a subscriber identification policy in the Subscriber Identification Policy panel.
7. Select an application profile in the Default Application Profile panel.
8. Configure the Default Subscriber Identification String parameter.

27

To configure SRRP for the group interface:

1. Click on the SRRP tab.

2. Configure the parameters in the Routing panel.
3. Click Create in the SRRP Instance panel. The SRRP Instance (Create) form opens.
4. Configure the required parameters in the SRRP Instance Information panel.
5. Select an operational group in the Operational Group panel.
6. Configure the Priority Step parameter.
7. Click on the Behavior tab.
8. Configure the required general parameters.
9. Select the SAP you need to use for the in-band messaging between the sites in the Message Path panel.
10. Select a policy pointer in the Policy Pointer 1 and Policy Pointer 2 panels.
11. Save the changes and close the form.

28

Save the changes and close the forms.

END OF STEPS

79.38 To create a bonding group interface on a VPRN

79.38.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces, right-click on a subscriber interface and choose Create VPRN Group Interface. The VPRN Group Interface (Create) form opens.

4

Configure the Name and Description parameters.

5

Select a redundant interface.

6

Select a diameter application policy.

7

Select a diameter authentication policy.

8

To configure bonding for the group interface:

1. Click on the Bonding tab.
2. On the General tab, configure the required parameters.
3. Select an FPE with subscriber management extensions enabled.
4. Click on the Connections tab.
5. Click Create. The Bonding Connection Object (Create) form opens.
6. Configure the required parameters, setting the Connection ID parameter to 1.
7. Save you changes and close the form.
8. Create a second bonding connection object, setting the Connection ID parameter to 2.

9

On the VPRN Group Interface, General tab set the Specific Type parameter to Bonding.

10

Save your changes and close the forms.

END OF STEPS

79.39 To configure a SAP on a VPRN group interface

79.39.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces→*subscriber_interface*, right-click on a group interface and choose Properties. The VPRN Group Interface (Edit) form opens.

4

Click on the Service Access Points tab.

5

Click Create. The VPRN Service Access Point (Create) form opens.

6

Configure the required general parameters.

7

Select a host lockout policy in the Host Lockout panel.

8

Configure a port:

1. Click on the Port tab.
2. Select a port for the L3 access interface from the Select Terminating Port form.

   Note:

   The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.
3. Configure the required parameters.

   The Auto-Assign ID parameter is configurable if the port uses dot1q encapsulation. When the parameter is enabled, the NFM-P automatically configures the Outer Encapsulation Value parameter using the lowest unassigned value.

   You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter on the User Preferences form. Choose Application→ User Preferences from the main menu.

   The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with QinQ encapsulation.

   The Outer Encapsulation Value (VPI) and Inner Encapsulation Value (VCI) parameters are configurable only for ATM ports.

9

To assign ingress and egress QoS policies to the SAP:

Items such as policies, schedulers, and filters can be applied later to multiple service components at once. Choose and right-click the components in the service navigation tree, choose Properties, and configure the parameters on the appropriate tab.

1. Click on the QoS tab.
2. Configure the required parameters.

   The Ingress Match QinQ Dot1P and Egress Mark QinQ Top Bits Only parameters are configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.
3. Select a policy in the Ingress Policy and Egress policy panels.

   Note:

If you select an ingress or egress policy which has a forwarding class mapped to an ingress or egress queue group, you must ensure that the port you selected in Step 8 has the access ingress or egress queue group with the same name created on it.

See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

4. Select an HS secondary shaper in the HS Overrides panel, if required.

10

To specify that an aggregation scheduler policy is not applied to the interface:

1. Set the Aggregation parameter to Off.

   Note:

   The Aggregation parameter is not configurable if the port you selected in Step 8 is an HSMDA port.

2. Configure the required parameters.

   The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

   The Frame-Based Accounting parameter is not configurable if the port you selected is an HSMDA port.

   You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.

3. Select an ingress scheduler in the Ingress Scheduler panel.
4. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
5. If the port you selected in Step 8 is an HSMDA port, save the changes and close the forms.
6. Select an egress scheduler in the Egress Scheduler panel.
7. Select an egress policer control policy in the Egress Policer Control Policy panel.

11

To specify that an aggregation scheduler policy is applied to the interface:

1. Set the Aggregation parameter to On.

   Note:

   You cannot specify an access scheduler policy if the port you selected in Step 8 is an HSMDA port.

2. Select an aggregation scheduler in the Aggregation Scheduler panel.

12

To assign ingress and egress ACL filters to the SAP:

1. Click on the ACL tab.

2. Select an ingress ACL filter in the Ingress Filter panel.
3. Select an egress ACL filter in the Egress Filter panel.

13

To assign an accounting policy to the SAP:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy.

14

To assign a virtual port to the SAP:

1. Click on the Virtual Port Name tab.
2. Configure the required parameters.

15

To configure anti-spoofing for the SAP, click on the Anti-Spoofing tab. Configure the Anti-Spoofing parameter and static hosts, as described in 74.24 “To create a static host for residential subscriber management on a SAP” (p. 2045) .

16

To assign a DoS protection policy or DDoS protection policy to the SAP:

Note: A default DoS protection policy is automatically assigned to the SAP.

1. Click on the Security tab.
2. Select a DoS protection policy in the NE DoS Protection panel.
3. Select an NE DDoS protection policy.
4. Configure the required parameters.

17

To define the levels of the ETH-CFM PDUs that are discarded on ingress into the SAP or to configure LMM or ETH-LMM frame loss measurement statistics collection:

1. Click on the OAM tab, then on the ETH-CFM tab.
2. Configure the parameters in the Squelch Ingress Level panel.

   Levels must be assigned contiguously from Level 0. If you select a level greater than 0, then all levels lower than the one you chose will automatically be selected.
3. Configure the required parameters in the LMM Session Stats Collection panel.

18

To configure residential subscriber management on the SAP:

1. Click on the Subscriber Management tab, then on the IGMP Host Tracking tab.
2. Select the import policy used to filter IGMP packets.
3. Configure the required parameters.
4. You can click on the Host Tracking Info tab to view a list of hosts that are being tracked on this service access point.
5. Click on the Profiles tab.
6. Configure the required general parameters.
7. Select a default subscriber profile, subscriber authentication policy, and default application profile in the Policies panel.
8. Configure the Profiled Traffic Only parameter in the Single Subscriber Configuration panel.
9. Select a non-subscriber subscriber profile for the SAP in the Non-Subscriber Traffic Subscriber Profile panel.
10. Select a non-subscriber traffic SLA profile for the SAP in the Non-Subscriber Traffic SLA Profile panel.
11. Select a non-subscriber traffic application profile for the SAP in the Non-Subscriber Traffic Application Profile panel.
12. To view active hosts for the subscriber instance, click on the Subscriber Hosts tab.

19

To configure a default host in cases where the associated DHCP filter policy is configured with the Bypass Host Creation action:

1. Click on the Default Hosts tab.
2. Depending on the required addressing format, click on either the IPv4 or IPv6 tab.
3. Click Create. The (IPv4|IPv6) Default Host (Create) form opens.
4. In the Default Host Configuration panel, Select an (IPv4|IPv6) address to associate with the subscriber interface SAP.

   The address must be one of the addresses on the subscriber interface that the SAP belongs to, and not another default host.
5. Configure the Next Hop (IPv4|IPv6) Address parameter.

   The next hop address can be duplicated on the same SAP, but not on another SAP default host on the same routing instance.
6. Save your changes and close the form.

   Default host configuration settings cannot be changed once the default host is created.

20

To assign an ANCP policy to the SAP:

1. Click on the ANCP Static Map tab.
2. Click Create. The ANCP Static Map (Create) form opens.
3. Configure the ANCP String parameter.
4. Select an ANCP policy in the ANCP Policy panel.
5. Save the changes and close the forms.

21

To configure ATM on the SAP:

1. Click on the ATM tab.
2. Configure the required parameters.
3. Select an ingress ATM policy in the Ingress ATM Policy panel.
4. Select an egress ATM policy in the Egress ATM Policy panel.

22

Save the changes and close the forms.

END OF STEPS

79.40 To configure LAG per-link hashing on a VPRN group interface SAP

79.40.1 Purpose

You can configure weighted per-link hashing on a SAP on a VPRN group interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member.

79.40.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces→*subscriber_interface*→Service Access Points, right-click on a SAP and choose Properties. The VPRN Service Access Point (Edit) form opens.

4

Click on the LAG Per Link Hash tab.

5

Configure the Class and Weight parameters.

6

Save the changes and close the forms.

END OF STEPS

79.41 To add an IGMP interface to a VPRN

79.41.1 Prerequisites

Before an IGMP interface can be configured, the following prerequisites must be met:

- The VPRN routing instance must be configured for IGMP; see 79.24 “To configure IGMP on a VPRN routing instance” (p. 2560) .
- The VPRN routing instance must be configured with an L3 interface; see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) .

79.41.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Routing Instance→Routing Instance→Protocols→*IGMP instance*.

4

Right-click on the IGMP instance and choose Create IGMP Interface. The IGMP Interface, Routing Instance (Create) form opens.

5

Configure the required parameters.

6

To assign an L3 access interface to the IGMP interface:

1. Click Select in the Interface panel and choose a L3 access interface
2. Configure the required parameters.

7

Click on the Behavior tab and configure the required parameters.

8

To configure multicast CAC, click on the Multicast CAC tab, select a multicast CAC policy and configure the required parameters.

9

To add a static multicast group or source:

1. Click on the Static Group/Source tab.
2. Click Create. The Static Grp Src, Interface routing instance (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

79.42 To add an IGMP group interface to a VPRN

79.42.1 Prerequisites

Before an IGMP group interface can be configured, the following prerequisites must be met:

- The VPRN routing instance must be configured for IGMP; see 79.24 “To configure IGMP on a VPRN routing instance” (p. 2560) .
- A VPRN group interface must be configured; see 79.37 “To configure a group interface on a VPRN” (p. 2586) .

79.42.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Routing Instance→Routing Instance→Protocols→*IGMP instance*.

4

Configure the required parameters.

5

Right-click on the IGMP instance and choose Create IGMP Group Interface. The IGMP Group Interface, Routing Instance (Create) form opens.

6

To assign a VPRN group interface to the IGMP group interface:

1. Click Select in the Interface panel and choose a VPRN group interface.
2. Configure the required parameters.

7

Click on the Behavior tab and configure the required parameters.

8

To configure multicast CAC, click on the Multicast CAC tab, select a multicast CAC policy and configure the required parameters.

9

Save the changes and close the forms.

END OF STEPS

79.43 To add an IP mirror interface to a VPRN

79.43.1 Purpose

Perform this procedure to configure an IP mirror interface in a VPRN service. This is a spoke terminated interface used to receive mirrored packets from a remote source.

79.43.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance, right-click on IP Mirror Interfaces and choose Create IP Mirror Interface. The IP Mirror Interface (Create) form opens.

4

Configure the required parameters.

The Name value for an IP mirror interface must be unique in VPRN site. There cannot be another L3 access interface, subscriber interface, or group interface with the same name on the VPRN site.

5

Save the changes and close the forms.

END OF STEPS

79.44 To configure a network interface on a VPRN site

Note: You cannot configure a network interface on a VPRN service if there are any pre-existing interfaces configured on the service. Once a network interface is configured on a VPRN service, no other interface types can be configured on the service.

79.44.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

Click on the Sites tab, choose a VPRN site, and click Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab and configure the Carrier Carrier VPN parameter.

The Carrier Carrier VPN parameter is configurable only if there are no other interfaces configured on the site.

5

Close the form.

6

On the VPRN Service (Edit) form, click on the Interfaces tab, then on the Network Interfaces tab and click Create. The Network Interface (Create) form opens.

7

Configure the required general parameters.

8

Configure the parameters in the Unicast RPF panel.

9

Configure the parameters in the NTP panel.

10

To configure Cflowd sampling:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click Apply. The Cflowd sampling object appears in the Cflowd Sampling panel.
3. To associate a service template with the Cflowd sampling object, select the newly created Cflowd sampling object and click Properties. The Cflowd Sampling (Edit) form opens.
4. Save your changes.

11

Configure the port.

1. Click on the Port tab.
2. Choose a terminating port for the interface in the Terminating Port panel and configure the required parameters.

   Note:

   If the Loopback Enabled parameter is selected on the General tab, the Port port tab does not appear.

   If you Clear or change the port associated with an existing Network Interface, the Network Policy associated to the Network interface remains but all the redirect information is cleared. This includes the Egress Queue Group Template Policy and its Egress Queue Group Instance ID, as well as the Ingress Queue Group Template Policy and its Ingress Queue Group Instance ID. These will subsequently need to be reconfigured if required.

12

Click on the Policies tab to select the network policy, the queue group template policies for the ingress forwarding plane redirect and egress port redirect, and the ingress and egress IP and IPv6 ACL policies.

Network policies are used to determine QoS settings based on the packet DSCP bits on the ingress and egress of the network.

Only global Ingress Queue Group Template Policies that have already been defined can be chosen. Queue group template policies allow SAP or IP interface forwarding classes to be redirected from the typical queue mapping to a shared queue.

Only local Egress Queue Group Template Policies which have a Network Egress Queue Group Instance defined on the associated port can be chosen.

Note: If you select a network policy with a forwarding class mapped to a queue group queue ID, you must ensure that the mapping queue group queue ID is in the selected Queue Group Template Policy.
If you need to add or remove a Port Redirect Group (queue or policer) to or from any of the Ingress or Egress Forwarding Classes in the global copy of the Network Policy, you must first define or remove the queue or policer in the global copy of the Egress Queue Group Template Policy.
You must ensure that the port you selected has a network egress queue group with the same name as the Queue Group Template policy created on it.
If the Network Policy you selected has Port Redirect Groups on its Ingress or Egress Forwarding Classes, you must specify both an Ingress or Egress Queue Group Template Policy and Instance ID. The Port Redirect Groups must already be defined within the Ingress or Egress Queue Group Template Policy you select. The default policy does not contain any redirects.
Queue Group Template policies are not applicable to L3 interfaces associated with HSMDA ports.
See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

13

To assign a non-default DoS protection policy or DDoS protection policy to the interface, click on the Security tab and select a DOS policy and NE DDos protection policy in the NE Protection panel.

14

To configure BFD for the interface, click on the BFD tab and configure the required parameters.

Note: You cannot configure BFD for the interface if BFD is disabled. See Chapter 28, “Routing protocol configuration” for information about enabling and disabling BFD for routing protocols.

15

To view local and remote session peers that are managed by the NFM-P, click on the BFD Session tab, choose a session and click Properties. View the following and close the form:

- BFD status
- protocol used
- local address
- remote address
- operational status and statistics

16

To configure ARP, click on the ARP tab and configure the Timeout (seconds) parameter.

1. Click Create to statically associate an IP or MAC address to the interface, and configure the required parameters.
2. Close the form.

17

Click on the Addresses tab to assign IP addresses to the interface:

1. Click Create. The IP Address, routing instance (Create) form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

18

Click on the ICMP tab and configure the required ICMP parameters.

19

Save the changes and close the forms.

END OF STEPS

79.45 To add a PIM interface to a VPRN

79.45.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Routing Instance→Routing Instance→Protocols→*PIM instance*, right-click on an instance and choose Create PIM Interface. The PIM Interface, Routing Instance (Create) form opens.

4

To assign an L3 access interface to the PIM interface:

1. Select an L3 access interface in the Interface panel.
2. Configure the required parameters.

5

Click on the Behavior tab and configure the required parameters.

The Sticky DR Priority parameter is configurable when the Sticky DR parameter is enabled.

6

To configure multicast CAC, click on the Multicast CAC tab, select a multicast CAC policy and configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

79.46 To implement dual homing using SRRP

79.46.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

Create the redundant interface used for SRRP out-of-band messaging between the two NEs:

Note: The two sites that participate in the dual homing configuration do not have to be part of the same service.

Ensure that the pair of sites each contains a properly configured subscriber interface and SAPs underneath the group interface that are participating in the redundant configuration.

Ensure that all subscriber interface IP addresses have a gateway address configured on them.

1. On the service navigation tree, click on the site to which you need to add the redundant pair; expand the entries for that site.
2. Right-click on Redundant Interfaces for one site of the redundant pair, and choose Create Redundant Interface. The Redundant Interface (Create) form opens.
3. Configure the required parameters.
4. Click on the Addresses tab.
5. Click Create. The IP Address (Create) form opens.
6. To specify IP addresses for the redundant interface on the current and remote sites, configure the required parameters.

   The remote IP address must be on the same subnet as the redundant interface IP address of the current site. For example, if the IP address of the current site is 7.7.7.7, with a prefix length of 24, then the redundant interface IP address of the remote site must be 7.7.7.d, where d is a value from 0 to 255, excluding 7.
7. Save the changes and close the forms.

4

Create an SDP spoke binding between the current and remote sites. The Source Interface is the Redundant Interface you created in Step 3 and the Tunnel Termination Site is the remote site. The Return Tunnel must come from the remote site. See 79.68 “To configure a VPRN spoke SDP binding” (p. 2635) for more information.

5

Assign the Redundant Interface to the Group Interface for the current site.

1. On the service navigation tree under the Subscriber Interface, right-click on the Group Interface and choose Properties. The VPRN Group Interface (Edit) form opens.
2. Select a redundant interface you created in Step 3 in the Redundant Interface panel.
3. Save the changes and close the form.

6

Create an SRRP Instance for the current site.

1. On the service navigation tree under Group Interfaces, right-click on the SRRP Instances item for the current site, and choose Create SRRP Instance. The SRRP Instance (Create) form opens.
2. Configure the required parameters.

   The SRRP ID value must be the same for the current and remote sites.
3. Click on the Behavior tab and configure the general parameters.

4. Select the SAP you need to use for the in-band messaging between the sites in the Port field.
5. Configure the Policy Pointers for the SRRP Instance.
6. Save the changes and close the form.

7

Click Turn Up to activate the SRRP instance.

8

Repeat Step 3 to Step 7 for the remote site.

Note: When you repeat Step 3 to Step 7 for the remote site, that site becomes the current site and the previously configured site is the remote site.
After the two sites have been properly set up, you can examine the SRRP peer associations at any time by right-clicking an SRRP Instance in the service navigation tree. This opens the SRRP Instance - Edit form, which contains a read-only field called SRRP Peer. The Site ID, Service ID, and Operational State of the associated peer appear in this field.
You can also examine the state of an SRRP Instance by checking the Operational Flags field. The flags indicate specific problems that might occur with the SRRP Instance, as follows:

- Duplicate Subscriber IF Address: one of the local subscriber IP addresses is the same as a subscriber IP address on the remote node.
- Redundant Interface Mismatch: the local SRRP instance and remote SRRP instance have mismatched redundant interfaces.
- SAP Mismatch: the local SRRP instance is backing a different set of SAPs than the peer.
- Subnet Mismatch: one of the subnets that SRRP is backing up does not have a match with the peer.
- Dual Master: both SRRP instances are master at the same time.
- SAP Tag Mismatch: the local SRRP instance is backing a set of SAPs with different remote and local tags.
- SRRP ID Mismatch: the peer has a different SRRP instance ID backing the same subnet.

9

Save the changes and close the forms.

END OF STEPS

79.47 To configure a subscriber interface on a VPRN

79.47.1 Prerequisites

The 7450 ESS in mixed mode, 7750 SR supports the configuration of a subscriber interface in a VPRN.

79.47.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance, right-click on Subscriber Interfaces and choose Create VPRN Subscriber Interface, or expand Subscriber Interfaces, right-click on a subscriber interface and choose Properties. The VPRN Subscriber Interface (Create|Edit) form opens.

4

Configure the required general parameters.

The Tunnel Fault Notification parameter is configurable on interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept, in order to receive the fault notification from the tunnel facility MEP.

The Name value for a subscriber interface must be unique in the context on the VPRN Service Site. That is, there cannot be another L3 access interface, subscriber interface, or group interface with the same name on the same VPRN site.

You must configure the Default Primary DNS Server Address parameter before you can configure the Default Secondary DNS Server Address parameter.

The WAN Mode parameter is configurable only on subscriber interface creation. The parameter is read-only on existing interfaces. The WAN Mode parameter is validated on retail/wholesale subscriber interfaces, in which case the parameter value should be equal across interfaces.

5

Configure the required parameters in the Interface Class panel.

If you selected a value of Unnumbered for the Class parameter, configure the Unnumbered Type parameter, and configure the Interface Name parameter or the IP Address parameter.

6

Configure the required parameters in the DHCP Server Synchronization panel.

7

In the WLAN GW panel, set the Redundancy/Pool Manager parameter to Enabled if you need to configure WLAN GW redundancy or an IPv6 pool manager. The WLAN GW Redundancy and WLAN GW IPv6 Pool Manager tabs appear.

8

Click on the WLAN GW Redundancy tab to configure WLAN GW redundancy.

9

Click on the WLAN GW IPv6 Pool Manager tab to configure DHCPv6 functionality for the subscriber interface.

In order to support WLAN GW IPv6 address pools, the subscriber interface must have the IPv6 and WLAN Redundancy/Pool Manager parameters enabled on the General configuration tab, and must be configured with a soft GRE group interface.

1. Select a WLAN GW group.
2. Configure the required parameters.

   The DHCPv6 Lease Query parameter must be enabled before the DHCPv6 Lease Query Max Retry parameter can be configured.

   The SLAAC and IA-NA Administrative State parameters must be set to Out Of Service before their related parameters can be configured.

10

To configure IPv6 forwarding on the subscriber interface:

1. Enable the IPv6 Allowed parameter.
2. Configure the IPv6 parameters.

   You must remove the check mark from the Default check box to access the Default Primary IPv6 DNS Server Address and Default Secondary IPv6 DNS Server Address parameters.
3. Click on the IPv6 Subscriber Prefixes tab.
4. Choose a subscriber prefix and click Properties, or click Create to create a new subscriber prefix. The Subscriber Prefix (Edit|Create) form opens.
5. Configure the parameters.
6. Save the changes and close the form.

11

To create a subscriber interface for a wholesale and retail VPRN:

Note: When you configure the routing instance for the VPRN, you can set the Type parameter to Subscriber Split Horizon to enable the forwarding service and forwarding subscriber information. The subscriber split horizon VPRN controls the flow of traffic for

wholesale subscriber applications. See 79.5 “To create a VPRN service” (p. 2534) for more information on how to configure this parameter.

You cannot create a group interface under a forwarding subscriber interface.

1. Select a forwarding service to assign to the subscriber interface in the Forwarding Service panel.
2. Configure the Private Retail Subnets parameter.
3. Select a forwarding subscriber interface in the Forwarding Subscriber Interface panel.
4. Configure the Export Host Routes parameter, if required.

   The Export Host Routes parameter is only available on retail subscriber interfaces when the Allow Unmatched Subnets parameter is enabled, or when the Interface Class parameter is set to Unnumbered.

12

To create one or more IP addresses for the subscriber interface that are inherited by the SAPs in the group interfaces that are child objects of the subscriber interface:

1. Click on the Addresses tab.
2. Click Create. The IP Address (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

13

To configure IPv4 DHCP for the interface:

1. Click on the DHCP tab.
2. If you specified a forwarding service in Step 11 , the General tab is configurable. Configure the parameters that define the forwarding service information. Otherwise go to 4 .
3. If you specified a forwarding service in Step 11 and depending on the type and release of the device that you are configuring, the Subscriber Authentication Policy panel is present. Select a Subscriber Authentication policy. Otherwise, go to 4 .
4. Configure the required parameters.
5. Click on the Server tab and configure the required parameters.

   The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.
6. Click on the Client Applications tab and configure the Client Applications parameter. You can enable either or both of the PPPoE or DHCP choices.

14

To configure PPPoE for the interface:

Note: PPPoE is only configurable on a VPRN subscriber interface if the interface is a retailer interface.

1. Click on the PPPoE tab.
2. Configure the required parameters.

15

If the subscriber interface is part of a wholesale/retail configuration, click on the IPoE Linking tab and configure the required parameters.

The IPoE Linking tab is only configurable if a forwarding service is specified for the subscriber interface (see Step 11).

16

To configure the ARP host for the interface:

Note: The ARP host is only configurable on a VPRN subscriber interface when the interface is a retailer interface.

1. Click on the ARP Host Configuration tab.
2. Configure the required parameters.

17

Save the changes and close the forms.

END OF STEPS

79.48 To force a WLAN GW switchover to standby management on a VPRN subscriber interface

79.48.1 Purpose

Perform this procedure to force a WLAN gateway service to switch to a standby system for the UE associated with a subscriber interface.

Note: The Force Switchover command is only available on an active subscriber interface.

79.48.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service form opens.

3

On the service navigation tree, expand Sites→Routing Instance→Subscriber Interfaces, right-click on the subscriber interface object on which you want to perform a WLAN GW switchover and choose Properties. The VPRN Subscriber Interface (Edit) form opens.

4

Click on the WLAN GW Redundancy tab.

5

Click Force Switchover.

6

Close the forms.

END OF STEPS

79.49 To add a TMS interface to a VPRN

79.49.1 Prerequisites

The 7750 SR-7 and 7750 SR-12 devices support the configuration of a TMS interface in a VPRN.

The following conditions must be met before you can configure a TMS interface.

- You must have one VPRN configured. The TMS interface is configured on this on-ramp VPRN.
- Optionally, you can have two addition VPRNs configured, one for the management VPRN and one for the off-ramp VPRN.
- Each VPRN must be associated within the same site on the same NE.
- An XP-IOM-3 is required.
- The ISA-TMS daughter card must be installed.

79.49.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service Subscriber (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance, right-click on TMS Interfaces and choose Create TMS Interface. The TMS Interface (Create) form opens.

4

Configure the required parameters.

The Name value for a TMS interface must be unique in the context of the VPRN Service Site. That is, there cannot be another L3 access interface, subscriber interface, group interface or TMS interface with the same name on the same VPRN site.

5

To assign an off-ramp VPRN to the TMS interface, select a VPRN in the Off-Ramp VPRN panel.

6

Set the off-ramp ingress routing context. Configure the Routing Instance parameter.

7

To assign a management VPRN to the TMS interface, select a VPRN in the Management VPRN panel.

8

To associate an ISA-TMS daughter card with the TMS interface, select an ISA-TMS in the TMS Info panel and configure the ISA-TMS Authentication parameter.

Note: An ISA-TMS card can only be assigned to one TMS interface.

9

To assign an IP address to the TMS interface:

1. Click on the Addresses tab.
2. Click Create. The IP Address (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

10

Save the changes and close the forms.

END OF STEPS

79.50 To configure VRF import and export policies on a VPRN site

79.50.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Configure the VRF import and export policies.

1. Click on the VRF Import Policies tab or on the VRF Export Policies tab.
2. Select the required VRF import or export policies 1 through 15.

   Note:

   Any RT values that are reserved in the RT reservation table for this service are automatically populated in the VRF import or export policy. See 79.63 “To reserve route targets for specific VPRN services” (p. 2628) for more information about how to reserve an RT for a specific VPRN service.

   If you enabled the Enable RT/RD Auto Assignment parameter when you created the VPRN service, the NFM-P automatically creates and assigns a VRF import or export policy to the VPRN site. Note 2 does not apply if the condition in Note 1 is met. See 79.51 “To configure a policy to reserve an RT and RD range for VPRN services” (p. 2616) for information about how to create a VPRN RT/RD Auto Assignment range policy.

5

Save the changes and close the forms.

End of steps

79.51 To configure a policy to reserve an RT and RD range for VPRN services

79.51.1 Purpose

Configure a range policy to reserve route targets and route distinguishers.

When you create the associated VPRN service, you must enable the Enable RT/RD Auto Assignment parameter on the General tab. You can only configure this parameter during VPRN service creation. See 79.5 “To create a VPRN service” (p. 2534) .

When you create the associated hub, spoke, or mesh VPRN site or sites, the NFM-P automatically generates the RT and RD based on the range policy. The NFM-P creates the routing policies and community for the sites, and automatically releases and distributes the policies and community to the NEs. The NFM-P assigns the VRF import and export policies to the VPRN site or sites, with the RT values that you are configuring in this range policy, and assigns the auto-generated RD to the VPRN site.

79.51.2 Steps

1

Choose Tools→Network Resources→VPRN RT/RD Auto Assignment from the NFM-P main menu. The VPRN RT/RD Auto Assignment (Edit) form opens.

2

Configure the Route Distinguisher Type parameter in the Route Distinguisher Range panel and configure the required parameters.

3

Configure the Route Target Format and Extended Community Type parameters in the Route Target Range panel and configure the required parameters.

4

Close the forms.

END OF STEPS

79.52 To automatically assign RT policies and RD configuration to VPRN sites

79.52.1 Purpose

The NFM-P can automatically assign route target policies and route distinguisher configuration to the VPRN sites from a pre-configured range policy. See 79.51 “To configure a policy to reserve an RT and RD range for VPRN services” (p. 2616) for more information about how to create a policy for RT and RD ranges.

79.52.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

Enable the Enable RT/RD Auto Assignment parameter in the VPRN Reservation panel.

4

Save and close the forms.

END OF STEPS

79.53 To configure a network ingress filter policy on a VPRN site

79.53.1 Purpose

Perform this procedure to assign an ingress IP and/or IPv6 template filter policy that will be applied to unicast traffic arriving on all network interfaces for the VPRN service site. The filter policy is applicable to both automatically created and explicitly defined spokes, and is supported for inter-AS and intra-AS network ports and for any label type. Network chassis mode D is required.

79.53.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

In the Network Ingress Filters panel:

1. Select the required ACL IP Filter policy.
2. Select the required ACL IPv6 Filter policy.

   You can specify both an IP and an IPv6 policy here, if required.

5

Save the changes and close the forms.

END OF STEPS

79.54 To configure ingress QoS policies on a VPRN site

79.54.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

In the Network Ingress QoS panel:

1. Select the required Network policy.
2. Select the required Queue Group policy.
3. Configure the Forwarding Plane Queue Group Instance ID parameter.

   Note:

   If the selected Network policy contains a policer, that same policer must also exist in the selected Ingress Queue Group template policy.

   The Network policy, Ingress Queue Group policy, and Instance ID must all be configured concurrently.

   All VPRN sites that employ the selected Network policy are listed in the VPRN Sites tab of the policy's properties form.

5

Save the changes and close the forms.

END OF STEPS

79.55 To configure a system-wide alternate source IP address for GRE encapsulation

79.55.1 Purpose

Use this procedure to configure an alternate source IPv4 address for the GRE tunnels in VPRN services. When an alternate address is configured, the system address is no longer used in VPRN services that auto-bind to the GRE tunnel.

79.55.2 Steps

1

On the equipment tree, right-click on an NE and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

In the GRE Encapsulation panel, click Select and select an interface in the Select Interface form.

4

Click OK and close the forms.

END OF STEPS

79.56 To bind a VPRN site to service tunnels

79.56.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

To configure an auto-bind tunnel, click on the SDP Binding tab, and then on the Auto-Bind Tunnel tab, and configure the required parameters.

When the Transport parameter is set to MPLS:LDP, you can choose segment routing as a tunnel option by selecting the SR-ISIS or SR-OSPF parameters. The following prerequisites apply:

- Segment routing must be configured; see 28.11 “Segment routing” (p. 894)
- You must configure segment routing on the ISIS or OSPF interfaces; see 28.63 “To configure an IS-IS interface” (p. 958) and 28.69 “To add a Layer 3 interface to an OSPF router” (p. 966)

5

Configure the Transport parameter:

a. To explicitly specify service tunnels and circuits for the service, set the Transport parameter to None.

b. To automatically bind the service to MPLS service tunnels utilizing MPLS, set the Transport parameter to LDP.

Note: To use MPLS as the transport type, you must bind LSPs to service tunnels during service tunnel configuration; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190)

c. To automatically bind the service to MPLS service tunnels using RSVP-LSP, set the Transport parameter to RSVP-LSP.

Note: The RSVP-LSP option is available only on the 7450 ESS in mixed mode, the 7705 SAR, 7750 SR, and 7950 XRS.

d. To automatically bind the service to MPLS service tunnels using RSVP or LDP, set the Transport parameter to RSVP or LDP.

This choice provides the ability to simultaneously support both tunnel options, in networks that have a mixture of LDP and RSVP-TE in place. NFM-P always tries to resolve the VPN route by using RSVP-LSP tunnels first (lowest metric). If no RSVP-LSP service tunnels are available, then tunnels configured for LDP are used. If RSVP-LSP tunnels subsequently become available again, then the route resolution automatically returns to RSVP-LSP.

Note: The RSVP or LDP option is available only on the a 7450 ESS in mixed mode, the 7705 SAR, 7750 SR, or 7950 XRS.

e. To automatically bind the service to GRE service tunnels set the Transport parameter to GRE.

6

Save the changes and close the forms.

END OF STEPS

79.57 To configure static routes on a VPRN site

79.57.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab, then on the Static Routes tab.

5

Define a static route that the PE VRF is to exchange with the CE:

1. Click Create. The Static Route (Create) form opens.
2. Configure the required general parameters.
3. Configure the parameters in the Destination panel.

   The IP Address parameter is configurable when the Type parameter is set to an option other than Black Hole.
4. Select an interface for the static route in the Destination panel.
5. Select an IPsec tunnel for the static route in the Destination panel.
6. Configure the parameters in the Other panel.
7. Configure the parameters in the CPE Check panel.

   The Target IP Address, Interval (seconds), Drop Count, Log, and Packet Size parameters are configurable only when the Enable CPE Check parameter is enabled.
8. Select a prefix list name in the Prefix List panel and configure the Prefix List Flag parameter. The Prefix List Flag parameter is configurable only when you choose a prefix list.

   Note:

   You cannot specify a Prefix List if BFD Enabled or Enable CPE Check parameters are enabled for the static route.
9. Configure the parameters in the Source/Destination Class Index panel.
10. Save the changes and close the form.

6

Save the changes and close the forms.

END OF STEPS

79.58 To configure route aggregates on a VPRN site

79.58.1 Purpose

Route aggregation allows you to group a number of routes with common IP prefixes into a single entry in the routing table. This reduces the number of routes that need to be advertised by the hosting router and reduces the number of routes in the routing tables of downstream routers.

79.58.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab, then the Route Aggregation tab and click Create. The Aggregation (Create) form opens.

5

Configure the required parameters.

The IP Address Prefix you enter must be an IPv4 address in the form x.x.x.0, where the last integer is the host bits and must have a value of 0.

The Aggregator AS and Aggregator IP Address parameters are configurable when the Aggregator parameter value is set to True. The Indirect Address parameter is configurable when the Next Hop Type parameter is set to Indirect.

6

Save the changes and close the forms.

END OF STEPS

79.59 To enable a tunnel facility MEP on the VPRN site

79.59.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click the OAM tab, then the ETH-CFM tab.

5

Set the Tunnel Fault Notification parameter to Accept.

The Tunnel Fault Notification parameter is configurable on interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation.

6

Save the changes and close the forms.

END OF STEPS

79.60 To enable the automatic selection of a route distinguisher on a VPRN site

79.60.1 Purpose

Since a route distinguisher (RD) must be unique on each PE in the network, you can allocate either a RD that you manually select or an NE-selected route distinguisher for each service. When you configure an auto-RD on a VPRN site, a Type-1 RD is automatically allocated by the NE based on the community range that you configure.

79.60.2 Steps

1

Before you configure a site for auto-RD selection, you must:

1. Enable BGP on the routing instance of the NE; see Procedure ; 28.29 “To enable BGP on a routing instance” (p. 916)
2. Enable the automatic selection of an RD on the NE and specify the community range. See 12.11 “To enable the automatic selection of an RD on an NE” (p. 349) .

2

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

3

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

4

On the navigation tree, right-click on VPRN Service and choose Create VPRN Site, or right-click on a site and choose Properties. The VPRN Site (Create|Edit) form opens.

5

Click on the Routing tab, then on the General tab.

6

Set the Route Distinguisher Type parameter to Type 1.

7

In the Type 1 panel, set the Auto Route Distinguisher parameter to true.

The Operational RD is displayed after you apply the changes to the site and to the service.

8

Save the changes and close the forms.

END OF STEPS

79.61 To add a Global Route Table to a VPRN site

79.61.1 Purpose

Packets within a VRF are able to perform a parallel lookup against a Global Route Table (GRT), as well as within the local VRF. A successful routing table match found in the local VRF is typically preferred over any match found in the GRT. However, a static route can be used to allow for specific prefixes covered by the static route to fail the lookup in the local VRF table, thus resulting in the guaranteed use of a route from the GRT.

The GRT is populated by defining export policies for each participating VPRN service, and the maximum number of routes that are exported from a specific VRF can be limited.

Note: The addition of a GRT to a VPRN service is supported on the 7450 ESS in mixed mode, 7705 SAR and 7750 SR.

79.61.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance, right-click on a site and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Routing tab and configure the Enable GRT Lookup parameter.

Enable the Allow Local Management parameter to enable SNMP/Telnet/SSH/FTP traffic to flow to the base dystem IP interface via the VPRN service.

5

Click on the GRT Export Policies tab and configure the required parameters.

6

Select up to five export policies.

7

Save the changes and close the forms.

END OF STEPS

79.62 To create an OSPF sham link between two VPRN sites

79.62.1 Purpose

Perform this procedure to create an intra-area OSPF sham link between two VPRN sites. See 79.1.10 “OSPF sham link support” (p. 2519) in “VPRN service management” (p. 2513) for more information.

79.62.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, choose the first site to participate in the sham link. If you need to create the site, perform 79.11 “To configure a VPRN site” (p. 2545) .

4

Choose an L3 access interface for the site. If you need to create the L3 access interface, see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) .

5

Configure OSPFv2 for the site. Perform 79.22 “To configure ISIS, L2TP, MLD, OSPFv2, OSPFv3, PIM, RIP, or WPP on a VPRN routing instance” (p. 2555) . Ensure that you select sub79.22 c in Step 4 .

6

Right-click on OSPFv2 under the Protocols icon in the service navigation tree and choose Create Sham Link. The ShamLink (Create) form opens.

7

Select the Remote Neighbor IP Address. This is the IP Address of the other site participating in the sham link.

8

Select and interface in the Interface panel.

9

Select the area ID in the Area panel.

10

Configure the Administrative State parameter.

11

Click on the Protocol Properties tab and configure the required parameters.

12

The Authentication tab is configurable, depending on the OSPF version. To configure authentication for the sham link, click on the Authentication tab and either:

a. Create an MD5-based authentication key, by setting the Authentication Type parameter to MD5-based and clicking Create. The Md5Key (Create) form opens. Configure the required parameters and save and close the form.

b. Create a simple password, setting the Authentication Type parameter to Simple Password and clicking Change Password to enter a password. The Password (Create) form opens. Configure the parameters and close the form.

13

To view OSPF configuration information on neighbor sites, click on the Virtual Neighbor tab.

14

Save the changes and close the form.

15

Repeat for the second site participating in the sham link.

16

Save the changes and close the forms.

END OF STEPS

79.63 To reserve route targets for specific VPRN services

79.63.1 Purpose

Reserve route targets for a specific VPRN so that no other VPRN services can use the import/export RT.

When you create the associated VPRN service with the customer name that you specify for the RT reservation in Step 3 of this procedure, and the name of the VPRN service matches the Service Name parameter that you specify for the RT reservation in Step 3 , the Route Target Reservation (Edit) form updates with the service ID and the SVC Mgr Service ID of the VPRN.

When you create the associated hub, spoke, or mesh VPRN site or sites, the NFM-P creates the routing policies and community for the sites, and automatically releases and distributes the policies and community to the NEs. The NFM-P assigns the VRF import and export policies to the VPRN site or sites, with the RT values that you are configuring in this RT reservation. The Reserved Route Targets (Edit) form, in Step 4 , updates with the automatically created VRF export and import policies for each hub, spoke, and mesh site.

See 79.50 “To configure VRF import and export policies on a VPRN site” (p. 2615) for more information about how to configure VRF import and export policies. See 27.1.12 “NE routing policies” (p. 821) in Chapter 27, “NE routing and forwarding” for more information about routing policies.

79.63.2 Steps

1

Choose Tools→Network Resources→VPRN RT Reservation from the NFM-P main menu. The VPRN RT Reservation form opens.

2

Click Create or choose a route reservation entry and click Properties. The Route Target Reservation (Create) form opens.

3

Configure the Customer Name and Service Name parameters. The NFM-P automatically creates the customer if it does not already exist. You must manually create the VPRN service and associated routing policies.

4

Add a route target reservation instance.

1. Click Create in the Route Targets panel. The Reserved Route Targets (Create) form opens.

2. Configure the Route Target 1 and Route Target 2 parameters.

5

Close the forms.

END OF STEPS

79.64 To configure a VXLAN termination on a VPRN

79.64.1 Purpose

Use this procedure to configure VXLAN tunnel termination on a VPRN site. The address terminating the VXLAN tunnels can be an IPv4 or IPv6 address. An FPE must be created and configured to terminate on the VPRN, with an SPD ID range configured and VXLAN termination enabled. See 12.41 “To create an FPE” (p. 374).

79.64.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Create a VXLAN tunnel termination:

1. Click on the Routing tab, then the VXLAN Tunnel Termination sub-tab.
2. Enter the tunnel termination IP address.
3. In the FPE panel, click Select and choose the FPE ID.
4. Click Apply and close the form.

5

Close the forms.

END OF STEPS

79.65 To configure WLAN GW functionality on a VPRN site

79.65.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Enable the Enable WLAN GW parameter.

5

Click on the WLAN GW tab and configure the APN Network Identifier parameter.

6

Configure the parameters in the Mobile Triggered Accounting panel.

7

Click on the Mobile GW Address Map tab to configure a mobile GW address map.

1. Click Create. The Mobile Gateway Address Map (Create) form opens.
2. Configure the Mobile Gateway Address Prefix and Prefix Length parameters.
3. Select a mobile gateway/peer profile policy.
4. Save your changes and close the form.

8

Click on the Distributed Subscriber Management tab and configure the TCP Maximum Segment Size Adjustment parameter.

9

Click on the Distributed Subscriber Management tab and configure the TCP Maximum Segment Size Adjustment parameter.

10

If the WLAN GW is intended for a home LAN extension configuration, click on the Cross Connect tab.

Configure the required parameters and select a WLAN GW group.

11

Save the changes and close the forms.

END OF STEPS

79.66 To configure a WLAN GW for a VPRN group interface

79.66.1 Prerequisites

Before you can configure a WLAN GW for a VPRN group interface, you must configure a VPRN group interface with the Specific Type parameter set to Soft GRE. See 79.37 “To configure a group interface on a VPRN” (p. 2586) .

79.66.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→*routing_instance*→Subscriber Interfaces→*subscriber_interface*.

4

Right-click on a WLAN GW group interface. The VPRN Group Interface (Edit) form opens.

5

Click on the WLAN GW tab and configure the required parameters.

The Interim Update parameter must be enabled if the Include Counters and Hold Down Time parameters are to be configured.

6

Select a WLAN GW group in the WLAN GW Group panel.

7

Select a tunnel router instance in the Tunnel Router Instance panel.

8

Select a default retail VPRN service in the Default Retail Service panel.

9

Configure egress QoS, if required.

1. In the Egress QoS panel, configure the required parameters.

   Do not configure the Aggregate Rate Limit parameter with a positive value if you specify a scheduler policy in 3 .
2. Select a QoS policy.
3. Select a Scheduler Policy.

10

If the WLAN GW is part of the home LAN extension configuration, configure the Maximum Number of Bridged Domains parameter on the LAN Extension panel.

This parameter can only be configured when the administrative state of the WLAN GW is Out of Service.

11

Select an authentication policy.

12

To configure an authentication hold time:

1. Enable the check box in the Authentication Hold Time panel.
2. Specify an authentication hold time in the Hours, Minutes, and Seconds fields.

13

Click on the SAP Parameters tab to configure SAP parameters, as required.

1. Configure the required parameters.

   If the Default ID Type parameter is set to Use String, configure the Default ID String parameter.

   If the WLAN GW is supporting managed routes, set the Anti-Spoofing parameter to Next Hop IP and MAC Address. To enhance security for MAC-only anti-spoofing, configure unicast RPF on the group interface; see 79.37 “To configure a group interface on a VPRN” (p. 2586) .
2. Select a subscriber profile.
3. Select an SLA profile.
4. Select an application profile.
5. Select a subscriber identification policy.

14

Click on the Gateway Addresses tab to configure gateway addresses for the WLAN GW.

1. Click Create or choose a gateway address entry and click Properties. The Gateway Address (Create|Edit) form opens.
2. Type a gateway IP address in the field and click OK.

   You can configure up to 10 gateway addresses (six IPv6 and four IPv4).

   One IPv6 address can be configured with the XConnect option.

15

Click on the VLAN Tag Ranges tab to configure VLAN tag ranges, as required.

1. Click Create or choose a VLAN tag range entry and click Properties. The VLAN Tag Range (Create|Edit) form opens.
2. Configure the required parameters.
3. Select a retail service.
4. Select a NAT policy.
5. Select an HTTP redirect policy.
6. Select an authentication policy.
7. Select a RADIUS proxy cache server.
8. If the VLAN tag range is part of an L2 wholesale-retail configuration, select a VPLS site in the L2 Retail panel, and configure the Administrative State and Description parameters as required.
9. Click on the DHCP tab and configure DHCP parameters, as required.
10. Click on the Distributed Subscriber Management tab to configure distributed subscriber management.
11. Configure the required parameters.
12. Select an accounting policy.
13. Select an ingress policer policy.
14. Select an egress policer policy.
15. Select an IP filter policy.
16. Select an application profile string, if required.
17. Click on the SLAAC tab to configure lifetimes for SLAAC-configured IPv6 hosts.

    The Administrative State parameter must be Up to configure parameters on the SLAAC tab.
18. Disable the Default check box and configure the lifetime parameters (in hours, minutes, and seconds) for any of the following categories:
    - Preferred Lifetime (Initial State)
    - Preferred Lifetime (Active State)
    - Valid Lifetime (Initial State)
    - Valid Lifetime (Active State)

19. Click on the DHCP6 tab to configure lifetimes for DHCPv6-configured IPv6 hosts.

    The Administrative State parameter must be Up to configure parameters on the DHCP6 tab.

20. Disable the Default check box and configure the lifetime parameters (in hours, minutes, and seconds) for any of the following categories:
    - Preferred Lifetime (Initial State)
    - Preferred Lifetime (Active State)
    - Valid Lifetime (Initial State)
    - Valid Lifetime (Active State)

21. Click on the BRG Configuration tab to configure BRG for the VLAN tag range.
    - Configure the Administrative State parameter.
    - Select a default BRG profile.
    - Enable the Authenticate BRG Only parameter, if only authenticated BRGs are permitted.

22. If the VLAN tag range is part of a LAN extension configuration, click on the Cross Connect tab.

    Configure the required parameters and select an ISA RADIUS policy.

    You must enable Authentication on DHCP on the VLAN Tag Range - General tab in order to set the Cross Connect configuration admin state to In Service

23. If the VLAN tag range is part of a home LAN extension configuration, click on the Home LAN Extension tab.

    Configure the required parameters.

    You must enable Authentication on DHCP on the VLAN Tag Range - General tab, and enable the Administrative State of the BRG on the VLAN Tag Range - BRG Configuration tab in order to enable the home LAN extension configuration.

24. Save the changes and close the form.

16

Click on the L2 AP tab to configure an L2 access point.

1. Click Create or select an existing L2 AP entry and click Properties. The WLAN GW L2 Access Point (Create|Edit) form opens.
2. Select a port.
3. Configure the required parameters.

17

Save the changes and close the forms.

END OF STEPS

79.67 To resync WLAN GW tunnels on a VPRN site

79.67.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the WLAN GW Tunnels tab and click Resync WLAN GW Tunnels.

5

Close the forms.

END OF STEPS

79.68 To configure a VPRN spoke SDP binding

79.68.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Right-click on Spoke SDP Bindings and choose Create Spoke SDP Binding or expand Spoke SDP Bindings, right-click on an spoke SDP binding and choose Properties. The Spoke SDP Binding (Create|Edit) form opens.

5

Select a source interface for the SDP binding in the Source Interface panel.

6

Configure the required parameters.

These parameters appear only when a source interface is assigned to the spoke SDP binding.

7

Configure the parameters in the Hash Label panel.

The Enable Hash Label and Enable Signal Capability parameters can only be configured for spoke-SDP bindings that are access interface terminated.

8

To specify a destination node for the spoke SDP binding:

a. If the destination NE is a managed NE, select an NE from a list of managed NEs.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

9

Configure the required parameters.

10

To specify a transport tunnel for the spoke SDP binding:

a. For the NFM-P to configure the transport tunnel automatically, enable the Auto-Select Transport Tunnel parameter and configure the Profile Name or the Tunnel Auto-Selection Transport Preference parameter.

b. To configure the transport tunnel manually, select a tunnel in the Tunnel Panel.

c. To configure an MPLS-TP transport tunnel manually, select an MPLS-TP tunnel in the Tunnel Panel.

11

Select an application profile for the spoke SDP binding.

12

To choose an AA transit policy types:

Note: You can only associate one AA transit policy type with a service object.
You can bind a transit policy to only one L3 access interface or spoke SDP binding per node.
The transit policy and the application profile must belong to the same application assurance group or partition.

a. Select an AA transit IP policy in the Transit IP Policy panel.

b. Select an AA transit prefix policy in the Transit Prefix Policy panel.

13

To configure custom object attributes for AA reporting:

1. Click on the NSP Analytics Parameters tab, then on the Reporting tab.
2. Click Create. The AA Reporting (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.
5. Click on the General tab.

14

To specify a transport tunnel for the Return SDP binding:

Note: You can create a return tunnel only between sites that are within the same service. If the sites are not in the same service, the Return tab does not appear.

a. Let the NFM-P configure the transport tunnel automatically.
   1. Click on the Return tab.
   2. Enable the Auto Select Return Transport Tunnel parameter.
   3. Configure the Profile Name or the Return Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually.
   1. Click on the Return tab.
   2. Select a service tunnel in the Return Tunnel panel.

15

To configure CPU protection:

1. Click on the Security tab.
2. Configure the IP Src Monitoring parameter.

16

To associate a MEP to the spoke SDP binding:

1. Click on the OAM tab, then the ETH-CFM tab.
2. Click Create in the MEPs panel. The MEP (Create) form opens.
3. Select a MEG.
4. Configure the required parameters.

   The CCM Padding Packet Size (Bytes) parameter cannot be configured when the CCM interval parameter is set to 10 ms or 100 ms.

17

If the MD for the MEP has a Name Type of none and the associated MEG has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable; click on the Y.1731 Tests tab and configure the required parameters.

The Eth Test Pattern parameter is configurable when the Eth Test Enabled parameter is enabled.

18

Click on the AIS tab and configure the required parameters.

The AIS Meg Level parameter is configurable when the AIS Enabled parameter is enabled.

19

Save the changes and close the forms.

END OF STEPS

79.69 To create an L2 SDP spoke termination on a VPRN service

79.69.1 Prerequisites

Ensure that a service and site have been created in the VPRN. To terminate an L2 service on a VPRN SDP spoke, you must identify the VC and an interface belonging to the VC. The interface must not have an associated port.

79.69.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites→Routing Instance→Spoke SDP Bindings, right-click on Spoke SDP Bindings and choose Create Spoke SDP Binding or expand Spoke SDP Bindings, right-click on an spoke SDP binding and choose Properties. The Spoke SDP Binding (Create|Edit) form opens.

4

Select a source interface for the SDP binding in the Source Interface panel.

If you are creating this spoke SDP binding for use with an IP mirror interface, select the required IP mirror interface as the source interface.

5

To specify a destination node for the spoke SDP binding:

a. If the destination NE is a managed NE, select an NE from a list of managed NEs.

b. If the destination NE is an unmanaged NE, specify the system ID for the Tunnel Termination Site parameter.

6

Configure the required parameters.

These parameters appear only when a source interface is assigned to the spoke SDP binding.

7

Configure the parameters in the Hash Label panel.

The Enable Hash Label and Enable Signal Capability parameters can only be configured for spoke-SDP bindings that are access interface terminated.

8

To specify a transport tunnel for the spoke SDP binding:

a. For the NFM-P to configure the transport tunnel automatically, enable the Auto-Select Transport Tunnel parameter and configure the Profile Name or the Tunnel Auto-Selection Transport Preference parameter.

b. To configure the transport tunnel manually, select a tunnel in the Tunnel Panel.

c. To configure an MPLS-TP transport tunnel manually, select an MPLS-TP tunnel in the Tunnel Panel.

9

Select an application profile for the spoke SDP binding.

10

To choose an AA transit policy types:

Note: You can only associate one AA transit policy type with a service object.
You can bind a transit policy to only one L3 access interface or spoke SDP binding per node.
The transit policy and the application profile must belong to the same application assurance group or partition.

a. Select an AA transit IP policy in the Transit IP Policy panel.

b. Select an AA transit prefix policy in the Transit Prefix Policy panel.

11

To configure custom object attributes for AA reporting:

1. Click on the NSP Analytics Parameters tab, then on the Reporting tab.
2. Click Create. The AA Reporting (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.
5. Click on the General tab.

12

Click on the QoS tab.

13

Configure the required parameters in the Forwarding Plane Redirect panel.

1. Select a network policy.
2. Select and Ingress Queue Group Template Policy.
3. Configure the Instance ID parameter.

14

Configure the required parameters in the Port Redirect panel.

1. Select a network policy.
2. Select and Egress Queue Group Template Policy.
3. Configure the Instance ID parameter.

15

Click on the States tab and configure the Administrative State parameter.

16

To assign ingress and egress ACL filters to the spoke SDP binding:

1. Click on the ACL tab.
2. Select an ingress ACL filter in the Ingress Filter panel.

   If you are configuring IP mirroring, you can optionally select an ingress IP filter. This is a packet mirroring option which specifies that packets matching the IP filter are mirrored to the mirror destination.
3. Select an egress ACL filter in the Egress Filter panel.

17

If you are configuring IP mirroring, save the changes and close the forms.

18

To assign an accounting policy to the spoke SDP binding:

1. Click on the Accounting tab.
2. Configure the Collect Accounting Statistics parameter.
3. Select an accounting policy.

19

Save the changes and close the forms.

END OF STEPS

79.70 To configure an MPLS-TP static pseudowire on a VPRN spoke SDP binding

79.70.1 Purpose

Perform this procedure to create an MPLS-TP static pseudowire on the spoke SDP binding. An MPLS-TP service tunnel must be used in the SDP binding, and the Control Word parameter for pseudowire OAM must be set to Preferred. See 79.69 “To create an L2 SDP spoke termination on a VPRN service” (p. 2638) .

79.70.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service tree, expand Sites→Routing Instance→Spoke SDP Bindings, right-click on the spoke SDP binding you need to configure and choose Properties. The Spoke SDP Binding (Edit) form opens.

4

Click on the Control Channel tab.

5

Configure the required parameters.

6

Configure the static PW:

1. Click on the Static PW tab.
2. Click Create. The PW Path ID (Create) form opens.
3. Configure the Path AGI parameter.
4. Configure the parameters in the Source Attachment Individual Identifier panel.
5. Configure the parameters in the Target Attachment Individual Identifier panel.
6. Click OK. The PW Path ID (Create) form closes.

7

Save the changes and close the forms.

END OF STEPS

79.71 To configure BFD on a VPRN spoke SDP binding

79.71.1 Purpose

BFD is used over the VCCV control channel for PW fault detection. BFD carried over a PW associated channel enables the monitoring of the PW between the terminating PEs, regardless of whether the service spans multiple hops. This allows faults that are local to individual PWs to be detected.

79.71.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and configure BFD on the spoke SDP binding:

1. On the General tab, enable the Enable BFD parameter.
2. Choose a BFD template. To create a BFD template, see 28.25 “To configure a BFD template policy” (p. 911) .

Note:

You must assign a BFD template to the spoke SDP binding if you enable BFD.

3. In the Failure Action dropdown, select the action to be performed on BFD failure.
4. Configure the Up-Timer parameter.

   Note:

   The Up-Timer parameter is applicable only when the value of Failure Action is set to down.

5

Save and close the forms.

End of steps

79.72 To clear BFD sessions and statistics on a VPRN spoke SDP binding

79.72.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to view the BFD session information.

4

Clear BFD sessions or statistics on the spoke SDP binding:

1. Click on the BFD tab, then on the BFD Session tab.
2. Click Clear All to clear all BFD sessions.
3. Click Clear All Statistics to clear all BFD statistics.

5

Close the forms.

End of steps

79.73 To view the BFD session status on a VPRN SDP spoke binding

79.73.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Spoke SDP Bindings and click on the spoke SDP binding on which you want to configure BFD.

4

Click on the BFD tab and view the status of the BFD session on the spoke SDP binding:

1. Click on the BFD Session tab.
2. Choose a BFD session and click Properties. The BFD Session (View) form opens.

5

Close the forms.

END OF STEPS

79.74 To run an OAM validation test for a VPRN service

79.74.1 Prerequisites

An OAM validator test suite must be created for the tested entity. See Chapter 90, “OAM diagnostic tests” for more information about how to create an OAM validator test suite.

Note: As an alternative, you can also run an OAM validation test on the service by performing a One Time Validation. This is a mostly automated procedure and is described in 90.55 “To run a one-time validation test on a service” (p. 3084) .

79.74.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VPRN Service (Edit) form opens.

3

Click Validate. If the Validate button is not visible, click on the More Actions button and choose Validate. If an OAM validator test suite is not associated to the service, a dialog box appears. Perform the following:

1. Click OK to associate the service with an existing OAM validator test suite. The Choose Validator Test Suite form opens.
2. Select an OAM validator test suite.

4

View the State Cause indicators. When the validation test fails, a check mark appears beside the OAM Validation Failure indicator.

5

Click on the Tests tab, then on the Tested Entity Result tab.

6

Choose an entry and click Properties. The Tested Entity Result (Edit) form opens.

7

Click on the Results tab to display the validation test results.

8

If you need to compare two test results from the same type of test, choose the two test results and click Compare; the Difference form opens.

9

Compare the test results.

10

Close the forms.

END OF STEPS

79.75 To configure OAM components on a VPRN site

79.75.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the OAM tab, then on the ETH-CFM tab.

5

Configure the Tunnel Fault Notification parameter in the Facility MEPs panel.

This parameter must be set to Accept, in order to receive the fault notification from a tunnel facility MEP. Tunnel Fault Notification is configurable on sites where the device has ports configured in access or hybrid mode with QinQ encapsulation.

6

Click the OAM tab, then the Configuration tab.

7

Configure the required Test Generation Options parameters.

8

Click the OAM tab, then the TWAMP tab.

9

To add a TWAMP reflector to this site, perform 79.76 “To create a TWAMP Light reflector on a VPRN site” (p. 2647) .

10

Save the changes and close the forms.

END OF STEPS

79.76 To create a TWAMP Light reflector on a VPRN site

79.76.1 Prerequisites

A TWAMP Light reflector is required to conduct a TWAMP Light Test Session. See Chapter 92, “Performance Monitoring tests” for more information.

79.76.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

In the service navigation tree, expand Sites and right-click on the routing instance icon in the service navigation tree on which you need to configure a TWAMP Light reflector. The Site (Edit) form opens.

4

Click the OAM tab, then the TWAMP tab.

5

Configure the required parameters.

6

Click Create in the Reflector Prefixes panel. The Prefix TWAMP Light Reflector (Create) form opens.

7

Configure the required parameters.

Prefixes are added to the reflector to determine which PM Sessions can target the reflector. You should specify the prefix address and length if you require masking. Only those TWAMP Light test sessions with a Source IP matching a prefix will be valid. A maximum of 50 prefixes can be added to a reflector. TWAMP Light sessions without a matching prefix will cause an alarm to be generated.

8

Click the TWAMP Light Sessions tab.

9

Click Add to create a TWAMP Light Test Session. See 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) for more information.

10

Click the PM Sessions tab.

11

Click Add to create a PM Session. See 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) for more information.

12

Save the changes and close the forms.

13

From the equipment view, right-click on the managed NE for which you want to configure a TWAMP Light Reflector and choose Properties. The Network Element (Edit) form opens.

14

Click on the Globals tab, then on the OAM tab.

15

Configure the TWAMP Light Reflector Inactivity Timer (seconds) parameter in the TWAMP-Light Reflector panel.

16

Save the changes and close the forms.

END OF STEPS

79.77 To view the last cleared BFD statistics and sessions on a VPRN site

79.77.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, click on the site on which you want to view the last cleared BFD statistics and sessions.

4

Click on the Clear Status tab and view the five last cleared BFD statistics and the five last cleared BFD sessions.

5

Close the forms.

END OF STEPS

79.78 To view VPRN services that use an auto-assigned RT and RD or reservation table

79.78.1 Steps

1

Choose Tools→Network Resources→Network Resources from the NFM-P main menu. The Network Resources form opens.

2

Click on the Route Distinguishers and Route Targets tab.

3

Click Search to view the services and sites that use an auto-assigned RT and RD or RT from a reservation table.

4

Close the forms.

END OF STEPS

79.79 To view DHCPv6 leases

79.79.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, click on the site where you need to view DHCPv6 leases; expand the entries for that site.

4

Choose routing instance from the navigation tree. The Site (Edit) form opens.

5

Click on the RADIUS/DHCP/Diameter tab. The Local DHCP Servers tab is displayed.

6

Click on the DHCP V6 tab. The list of associated DHCPv6 servers is displayed.

7

Choose a DHCPv6 server and click Properties. The Local DHCPv6 Server (Edit) form opens.

8

Click Show Leases. If the Show Leases button is not visible, click on the More Actions button and choose Show Leases. The DHCPv6 Server Show Leases form opens.

9

Perform one of the following:

a. To view the entire list of leases that are active on the DHCPv6 server, click OK.

b. To view the lease information for a specific prefix, enter the IP address in the Prefix parameter for the particular DHCPv6 server and click OK.

Note: If the prefix does not exist in the list of leases, a message displays.

c. To view the detailed lease information for a specific prefix, enter the IP address in the Prefix parameter for the particular DHCPv6 server, select the Detail check box, and click OK. The form displays the CLI details for the specific prefix.

10

Close the forms.

END OF STEPS

79.80 To view DHCPv6 log events

79.80.1 Purpose

This procedure describes how to view lease not owner and pool unknown log events for local DHCPv6 servers.

79.80.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, click on the site where you need to view DHCPv6 log events; expand the entries for that site.

4

Choose routing instance from the service navigation tree. The Site (Edit) form opens.

5

Click on the RADIUS/DHCP/Diameter tab. The Local DHCP Servers tab is displayed.

6

Click on the DHCP V6 tab. The list of associated DHCPv6 servers is displayed.

7

Choose a DHCPv6 server and click Properties. The Local DHCPv6 Server (Edit) form opens.

8

Click on the Logs tab. The Lease Not Owner Log and Pool Unknown Log tabs display.

9

Click on the tab for the type of log that you wish to display. Configure the filter criteria and click Search. A list of NE SHCV event logs appears.

10

Select an event logs and click Properties. The Log form opens.

11

View the log entry.

12

Close the forms.

END OF STEPS

79.81 To view the service topology map associated with a VPRN service

79.81.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Topology View. A Topology View dialog box appears.

3

Click Yes to proceed. The Service Topology - map opens.

See Chapter 4, “Topology map management” for more information about service topology views.

END OF STEPS

VPRN L3 access interfaces procedures

79.82 Workflow to configure VPRN L3 access interfaces

79.82.1 Overview

The following workflow lists the high-level steps required to configure L3 access interfaces on a VPRN site.

79.82.2 Stages

1

Configure an L3 access interface on a VPRN site; see 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) .

2

Configure an override source VPRN L3 access interface; see 79.84 “To configure an override source VPRN L3 access interface on a VPRN site” (p. 2658) .

3

Bind an application profile to a VPRN L3 access interface; see 79.85 “To bind an application profile to a VPRN L3 access interface” (p. 2659) .

4

Configure LAG per-link hashing on a VPRN L3 access interface; see 79.86 “To configure LAG per-link hashing on a VPRN L3 access interface” (p. 2660) .

5

Configure load balancing on a VPRN L3 access interface; see 79.87 “To configure load balancing on a VPRN L3 access interface” (p. 2661) .

6

Configure custom object attributes for AA reporting on a VPRN L3 access interface; see 79.88 “To configure custom object attributes for AA reporting on a VPRN L3 access interface” (p. 2661) .

7

Assign ingress and egress QoS policies to a VPRN L3 access interface; see 79.89 “To assign ingress and egress QoS policies to a VPRN L3 access interface” (p. 2662) or 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665).

8

Configure scheduling on a VPRN L3 access interface; see 79.91 “To configure scheduling on a VPRN L3 access interface” (p. 2666) .

9

Assign ingress and egress ACL filters to a VPRN L3 access interface; see 79.92 “To assign ingress and egress ACL filters to a VPRN L3 access interface” (p. 2668) .

10

Assign a virtual port to a VPRN L3 access interface; see 79.93 “To assign a virtual port to a VPRN L3 access interface” (p. 2669).

11

Assign an accounting policy to a VPRN L3 access interface; see 79.94 “To assign an accounting policy to a VPRN L3 access interface” (p. 2670) .

12

Assign an accounting template policy to a VPRN L3 access interface, tunnel interface, or subscriber group interface; see 79.95 “To assign an accounting template policy to a VPRN interface” (p. 2670) .

13

Associate a security zone policy with a VPRN L3 access interface on a 7705 SAR; see 79.96 “To associate a security zone policy with a VPRN L3 access interface on a 7705 SAR” (p. 2672) .

14

Assign a time of day suite to a VPRN L3 access interface; see 79.97 “To assign a time of day suite to a VPRN L3 access interface” (p. 2672) .

15

Bind a VPRN L3 access interface to a VPLS site or VPLS I-site; see 79.98 “To bind a VPRN L3 access interface to a VPLS site or VPLS I-site” (p. 2673) .

16

Associate a local DHCPv4 or DHCPv6 server with a VPRN L3 access interface; see 79.99 “To associate a local DHCPv4 or DHCPv6 server with a VPRN L3 access interface” (p. 2674) .

17

Assign an NE DoS protection policy or NE DDoS protection policy to a VPRN L3 access interface; see 79.100 “To assign an NE DoS or DDoS protection policy to a VPRN L3 access interface” (p. 2675) .

18

Configure residential subscriber management for a VPRN L3 access interface; see 79.101 “To configure residential subscriber management for a VPRN L3 access interface” (p. 2676) .

19

Assign an IP address to a VPRN L3 access interface; see 79.102 “To assign an IP address to a VPRN L3 access interface” (p. 2677) .

20

Configure IPv4 ICMP for a VPRN L3 access interface; see 79.103 “To configure IPv4 ICMP for a VPRN L3 access interface” (p. 2678) .

21

Configure IPv6 ICMP on a VPRN L3 access interface; see 79.104 “To configure IPv6 ICMP on a VPRN L3 access interface” (p. 2678) .

22

Configure BFD for a VPRN L3 access interface; see 79.106 “To configure BFD for a VPRN L3 access interface” (p. 2680) .

23

Configure ARP for a VPRN L3 access interface; see 79.107 “To configure ARP for a VPRN L3 access interface” (p. 2682) .

24

Configure neighbor discovery on a VPRN L3 access interface; see 79.108 “To configure neighbor discovery on a VPRN L3 access interface” (p. 2682) .

25

Configure IPv4 DHCP for a VPRN L3 access interface; see 79.109 “To configure IPv4 DHCP for a VPRN L3 access interface” (p. 2683) .

26

Create a VRRP instance on a VPRN L3 access interface for a virtual router; see 79.110 “To create a VRRP instance on a VPRN L3 access interface for a virtual router” (p. 2685) .

27

Configure anti-spoofing filters for a VPRN L3 access interface; see 79.111 “To configure anti-spoofing filters for a VPRN L3 access interface” (p. 2686) .

28

Configure router advertisement on a VPRN L3 access interface; see 79.112 “To configure router advertisement on a VPRN L3 access interface” (p. 2687) .

29

Assign an ANCP policy to a VPRN L3 access interface; see 79.113 “To assign an ANCP policy to a VPRN L3 access interface” (p. 2688) .

30

Specify queue or meter overrides on a VPRN L3 access interface; see 79.114 “To specify QoS policy overrides on a VPRN L3 access interface” (p. 2689) .

31

Configure IPv6 DHCP on a VPRN L3 access interface; see 79.115 “To configure DHCPv6 on a VPRN L3 access interface” (p. 2690) .

32

Associate a Multi-Chassis shunting profile to a VPRN L3 access interface; see 79.116 “To associate a Multi-Chassis shunting profile to a VPRN L3 access interface” (p. 2691) .

79.83 To configure an L3 access interface on a VPRN site

79.83.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Right-click on L3 Access Interfaces and choose Create VPRN L3 Access Interface or expand L3 Access Interfaces, right-click on an access interface and choose Properties. The VPRN L3 Access Interface (Create|Edit) form opens.

5

Configure the required parameters.

If you enable the Loopback Enabled parameter, you cannot associate a port with the L3 access interface.

The Admin Link Local Address parameters are only configurable when the IPv6 Allowed parameter is enabled.

The Unnumbered Type parameter is configurable when the Class parameter is set to Unnumbered.

The Unnumbered IP Address parameter is configurable when the Unnumbered Type parameter is set to IP Address.

The Unnumbered Interface Name parameter is configurable when the Unnumbered Type parameter is set to Name.

IPv6–specific parameters are only configurable when the IPv6 Allowed parameter is enabled.

The Tunnel Fault Notification parameter is configurable on interfaces where the device has ports configured in access or hybrid mode with QinQ encapsulation.

If you are configuring a tunnel facility MEP, Tunnel Fault Notification must be set to Accept, in order to receive the fault notification from the tunnel facility MEP.

6

Configure the required parameters in the Frame Size Constraints panel.

7

Configure the required parameters in the Unicast RPF panel.

8

Select a host lockout policy in the Host Lockout panel.

9

Configure the required parameters in the PTP HW panel.

You can only enable the PTP HW Assist parameter after the port is associated to the interface, and an IP address is configured on the interface.

10

To configure Cflowd sampling:

1. Click Create in the Cflowd Sampling panel. The Cflowd Sampling (Create) form opens.
2. Configure the required parameters and click Apply. The Cflowd sampling object appears in the Cflowd Sampling panel.
3. Save your changes.

11

Configure the required parameters in the Redundant Next Hop Addresses panel.

12

Select a monitored group in the Operational Group panel.

13

Associate a port with the L3 access interface:

1. Click on the Port tab.

   Note:

If the Loopback Enabled parameter in Step 6 is enabled, you cannot associate a port with the L3 access interface. Go to Step 14 .

2. Select a port in the Terminating Port panel.

   Note:

   Only ports in access or hybrid mode, or PW ports are listed. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. The port is listed when you click Search.

   The form only lists PW ports that are bound to a service tunnel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for information about how to create a service tunnel with a PW port binding.

   You can select a bundle in the Terminating Port panel.

3. Configure the required parameters.

   The Auto-Assign ID parameter is configurable if the port uses dot1q encapsulation. When the parameter is enabled, the NFM-P automatically configures the Outer Encapsulation Value parameter using the lowest unassigned value.

   You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter on the User Preferences form. Choose Application→ User Preferences from the main menu.

   The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with QinQ encapsulation.

   The Outer Encapsulation Value (VPI) and Inner Encapsulation Value (VCI) parameters are configurable only for ATM ports.

4. Configure the required parameters in the Properties panel.
5. Select an NE DoS protection policy or an NE DDoS protection policy in the Security panel and configure the required parameters.

14

Save the changes and close the forms.

End of steps

79.84 To configure an override source VPRN L3 access interface on a VPRN site

79.84.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand the Sites icon, right-click on a routing instance and choose Properties. The VPRN Site (Edit) form opens.

4

Click on the Source Addresses tab and click Create. The Source Address (Create) form opens.

5

Configure the Source IP Application parameter.

6

Set the Source Address Termination parameter to Interface Index.

7

Select a source address VPRN interface.

8

Save the changes and close the forms.

END OF STEPS

79.85 To bind an application profile to a VPRN L3 access interface

79.85.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→L3 Access Interfaces, right-click on an L3 access interface and choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Application Assurance tab.

5

Select an application profile string.

Only local profiles that already exist on the NE can be selected.

6

To choose one of the following AA transit policy types:

Note: You can only associate one AA transit policy type with a service object.
To bind a transit policy to an L3 access interface, a port must already exist on the interface.
You can bind a transit policy to only one L3 access interface or spoke SDP binding per NE.
The transit policy and the application profile must belong to the same application assurance group or partition.

a. Select an AA transit IP policy in the Transit IP Policy panel.

b. Select an AA transit prefix policy in the Transit Prefix Policy panel.

7

Select an AA redundant protocol in the AA Redundant Protocol panel.

8

Save the changes and close the forms.

END OF STEPS

79.86 To configure LAG per-link hashing on a VPRN L3 access interface

79.86.1 Prerequisites

You can configure weighted per-link hashing on a VPRN L3 access interface if the terminating port has LAG per-link hashing enabled. The interface must be a LAG member.

79.86.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Sites→Routing Instance→L3 Access Interfaces, right-click on the L3 access interface and choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the LAG Per Link Hash tab.

5

Configure the Class and Weight parameters.

6

Save the changes and close the forms.

END OF STEPS

79.87 To configure load balancing on a VPRN L3 access interface

79.87.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose an VPRN and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Configure the L4 Load Balance, TEID Load Balancing, Egress IP Load Balancing, IPv6 Flow Label Load Balancing, and SPI Load Balancing parameters, as required.

5

Save the changes and close the forms.

END OF STEPS

79.88 To configure custom object attributes for AA reporting on a VPRN L3 access interface

79.88.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

On the General tab, choose an Application Profile for the application Group and Partition.

5

Click on the Application Assurance tab, then the NSP Analytics Parameters sub-tab.

6

Click on the Reporting tab and click Create. The AA Reporting (Create) form opens.

7

Configure the required parameters.

8

Save the changes and close the forms.

END OF STEPS

79.89 To assign ingress and egress QoS policies to a VPRN L3 access interface

79.89.1 Before you begin

The available panels and parameters vary depending on the NE, chassis type, and release.

If you are assigning QoS policies to an access interface on a 7210 SAS NE, see 79.90 “To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site” (p. 2665).

79.89.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the QoS tab and configure the Ingress Match QinQ Dot1P parameter.

The QoS tab is configurable only if a port is assigned to the interface.

The Ingress Match QinQ Dot1P parameter is configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.

5

Select an ingress QoS policy in the Ingress Policy panel.

Note: If you select an ingress policy which has a forwarding class mapped to an ingress queue group, you must ensure that the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) has the access ingress queue group with the same name created on it.
See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.

6

Select an ingress queue group template policy in the Forwarding Plane Redirect panel.

7

Configure the Instance ID and the Egress Mark QinQ Top Bits Only parameters.

The Egress Mark QinQ Top Bits Only parameter is configurable only when the encapsulation type of the port is BCP dot1q, dot1q, or QinQ.

8

Configure the required parameters in the Aggregate Rate Limit panel.

9

Select an egress policy in the Egress Policy panel.

Note: If you select an egress policy which has a forwarding class mapped to an egress queue group, you must ensure that the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) has the access egress queue group with the same name created on it.
See 16.24 “To configure Ethernet ports” (p. 599) for more information about how to configure Ethernet ports. See Chapter 49, “Policies overview” for more information about queue group template policies.
Queue Group Template policies are not applicable to L3 interfaces associated with HSMDA ports.

10

Select an Egress Queue Group Template Policy in the Port Redirect panel.

Note: Selecting an Egress Queue Group Template Policy here permits the redirection of Ethernet traffic packets to a queue ID specified in the egress port queue group of the SAP. The following properties and restrictions apply:

- If an Egress Queue Group Template Policy is specified here, the policy must have port redirection enabled.
- You cannot use policy-based redirection with the queue group when the SAP has SAP-based redirection enabled.
- Port access egress redirection is only supported on Ethernet/LAG ports. It is not supported on SAPs bound on non-Ethernet, Eth-tunnel, or CCAG ports.
- Supported ports include access, hybrid, and HSMDA.
- Queue groups can be applied to SAPs that incorporate LAGs. The LAGs can include port members from just a single card or from multiple cards.
- If you edit a LAG incorporated by the SAP, you cannot remove the last LAG member if a queue group reference exists to the containing SAP.
- You cannot add a secondary LAG member that has a queue group mismatch with primary LAG member.

11

If you are configuring an L3 access interface for a 7705 SAR, or if the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) is not an HSMDA port, then save the changes and close the forms.

Otherwise, configure the Packet Byte Offset (bytes) parameter. You must enable the associated Override parameter.

12

Select a WRR policy in the Egress HSMDA Override panel.

13

Select an HSMDA egress secondary shaper policy in the Egress HSMDA Override panel.

14

In the Shaper Group panel, select a Shaper Group for the access ingress port or access egress port.

15

In the IXR Specific panel, select an Egress Remark policy, Egress VLAN QoS policy, and a Shared Policer policy, as required. See 50.82 “To configure a 7250 SROS Remarking policy” (p. 1634), 50.54 “To configure a 7250 SROS VLAN QoS policy” (p. 1594), and 50.101 “To configure a shared policer policy” (p. 1661).

16

Select an HS secondary shaper in the HS Overrides panel, if required.

17

Save the changes and close the forms.

END OF STEPS

79.90 To assign ingress and egress QoS policies to a VPRN L3 access interface on a 7210 SAS site

Note: The available parameters and policies vary depending on the device type and chassis variant. The configurations that are supported on the site NE are shown on the form.

79.90.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the QoS tab and expand the 7210/7250 Specific panel.

5

Select a SAP Access Ingress policy in the Ingress Policy panel.

When you assign an ingress policy, the Color Mode parameter setting for the meters in the policy must coincide with the Enable DEI parameter setting on the physical port. When Enable DEI is selected, the Color Mode for meters must be set to Color Aware. When Enable DEI is not selected, the Color Mode for meters must be set to Color Blind. See 16.24 “To configure Ethernet ports” (p. 599) and 50.29 “To configure a 7210, 7250, and 1830 SAP Access Ingress policy” (p. 1544) .

To support H-metering, you must choose an ingress policy with all meter rate modes set to trTCM (RFC 4115).

For 7250 IXR sites, the selected SAP Access Ingress policy must contain an assigned 7250 Ingress CoS policy.

6

Select a SAS egress policy.

7

To enable table-based color-aware ingress classification, select the Enable Table Classification parameter. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

8

Select an Egress Remarking policy in the Egress Remark Policy panel.

9

Configure the required parameters in the Aggregate Rate Limit panel.

You can configure the Ingress Meter parameter only during SAP creation. The parameter must be set to true to support H-metering.

You can configure the Ingress Meter Rate (kbps) and Ingress Meter Burst parameters only after SAP creation.

You can configure the Egress Meter Rate and Egress Meter Burst parameters only when resources are allocated to the SAP Egress Aggregate Meter parameter in the system resource profile; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380). To allocate resources on the 7210 SAS-R, configure the Egress SAP Aggregate Meter parameter in the system resource profile policy assigned to the device; see 12.51 “To configure a system resource profile policy for the 7210 SAS-R/T/Mxp, 7210 SAS-S/Sx, or 7210 SAS-S/Sx VC” (p. 382).

You must also enable port-based scheduling on 7210 SAS-Mxp and 7210 SAS-R NEs; see 12.53 “To configure port-based scheduling on the 7210 SAS” (p. 384).

The Enable Egress Meter Stats parameter is available when a value is configured for the Egress Meter Rate parameter.

10

Configure the required parameters in the IXR Specific panel. Select an Egress Remark policy, Egress VLAN QoS policy, and a Shared Policer policy, as required.

11

Save the changes and close the forms.

END OF STEPS

79.91 To configure scheduling on a VPRN L3 access interface

79.91.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Schedulers tab and configure the required parameters.

Note: The Schedulers tab is configurable only if a port is assigned to the SAP in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) .

5

To configure scheduling on a 7210 SAS site, enable and configure the required parameters in the Egress Aggregate Rate Limit panel and go to Step 9 .

6

To configure scheduling on a 7705 SAR site:

Note: For the 7705 SAR, scheduler behavior is determined by the scheduler mode, which is 4-Priority by default. You can only configure the Egress and Ingress Aggregate Rate Limit parameters when the Scheduler Mode parameter is set to 16-Priority. You can set the Scheduler Mode to 16-Priority only when the port is on an MDA that supports 16-Priority. See the 7705 SAR documentation for more information.

If you change the Scheduler Mode parameter from 16-Priority to 4-Priority, the NFM-P automatically restores the default settings for the Egress Aggregate Rate Limit and Ingress Aggregate Rate Limit panels when you click on the Apply or OK button.

1. In the Egress Scheduler panel, configure the Scheduler Mode parameter.
2. In the Ingress Scheduler panel, configure the Scheduler Mode parameter.
3. If you set the Scheduler Mode parameter to 16-Priority in the Egress Scheduler panel or Ingress Scheduler panel, configure the parameters in the Egress Aggregate Rate Limit panel and Ingress Aggregate Rate Limit panel.

7

To specify that an aggregation scheduler policy is not applied to the interface:

1. Set the Aggregation parameter to Off.

   Note:

   The Aggregation parameter is not configurable if the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) is an HSMDA port.
2. Configure the required parameters.

The Aggregate Rate Limit (kbps), Frame-Based Accounting, and Limit Unused Bandwidth parameters are configurable only when there is no scheduler specified in the Egress Scheduler panel.

The Frame-Based Accounting parameter is not configurable if the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) is an HSMDA port.

You cannot specify an egress scheduler when the Aggregate Rate Limit (kbps) parameter is set to a value greater than zero.

3. Select an ingress scheduler in the Ingress Scheduler panel.
4. Select an ingress policer control policy in the Ingress Policer Control Policy panel.
5. If the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) is an HSMDA port, go to Step 8 .
6. Select an egress scheduler in the Egress Scheduler panel.
7. Select an egress policer control policy in the Egress Policer Control Policy panel.
8. Go to Step 9 .

8

To specify that an aggregation scheduler policy is applied to the interface:

Note: You cannot specify an access scheduler policy if the port you selected in 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656) is an HSMDA port. Go to Step 9 .

1. Set the Aggregation parameter to On.
2. Select an aggregation scheduler in the Aggregation Scheduler panel.

9

Save the changes and close the forms.

END OF STEPS

79.92 To assign ingress and egress ACL filters to a VPRN L3 access interface

79.92.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the ACL tab.

Note: When you configure ACL filters on a 7210 SAS NE, you must configure the system resource profile appropriately. See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) for more information.

5

Select the required ACL filter policies.

Note: Select the Ingress IP Filter Pair option to configure IP and IPv6 filters simultaneously to filter IP and IPv6 traffic respectively. Once the filters are paired, they can not be used separately or paired with another filter. This option is supported only for VPRN services on 7250 IXR NE, 22.2 R1 or later.

6

Save the changes and close the forms.

END OF STEPS

79.93 To assign a virtual port to a VPRN L3 access interface

79.93.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Virtual Port Name tab.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

79.94 To assign an accounting policy to a VPRN L3 access interface

79.94.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Accounting tab.

5

Select an accounting policy.

6

Configure the required parameters.

Collect Egress Queue Statistics parameter can be configured only during VPRN L3 access interface creation. See 79.83 “To configure an L3 access interface on a VPRN site” (p. 2656).

7

Save the changes and close the forms.

END OF STEPS

79.95 To assign an accounting template policy to a VPRN interface

79.95.1 Purpose

This procedure applies to VPRN L3 access interfaces, tunnel interfaces, or subscriber group interfaces.

For information about creating an accounting template policy, see 54.12 “To configure an accounting template policy” (p. 1757) .

79.95.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria. A list of services appears.

3

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

4

Expand the Sites icon in the service navigation tree.

5

Expand the Routing Instance icon in the service navigation tree on which you need to configure the accounting template policy.

6

Right-click on the L3 access interface, tunnel interface, or subscriber group interface you need to configure and choose Properties.

7

Click on the Policies tab.

8

Select a accounting template in the accounting template panel.

9

Save and close the form.

END OF STEPS

79.96 To associate a security zone policy with a VPRN L3 access interface on a 7705 SAR

79.96.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Zone tab.

5

Select a security zone policy.

You can create a security zone policy by clicking Create.

6

Configure the ByPass Zone Config parameter.

7

Save the changes and close the forms.

END OF STEPS

79.97 To assign a time of day suite to a VPRN L3 access interface

79.97.1 Prerequisites

You can assign a ToD suite to a L3 access interface only if accounting statistics collection is disabled on the interface. You must disable the Collect Accounting Statistics parameter in Step 4 of 79.94 “To assign an accounting policy to a VPRN L3 access interface” (p. 2670) .

79.97.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the TOD Suite tab.

5

Select a time of day suite.

Note: SapEgrQosPlcyStats and SapIngQosPlcyStats statistics are collected only if a Time Of Day Suite is applied on the SAP.

6

Save the changes and close the forms.

END OF STEPS

79.98 To bind a VPRN L3 access interface to a VPLS site or VPLS I-site

79.98.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Routed VPLS tab.

Note: The operational state of the IP interface binding will not be turned up until the Enable IP Interface Binding parameter is set to true.
You can create and manage a routed VPLS connector from the navigation tree on the Composite Service (Edit) form.

5

Enter a VPLS site name or select a VPLS site.

6

Configure EVPN Tunnel parameters as needed.

7

Expand the Ingress or Egress panel as needed.

8

To configure the filter policies:

1. Select an IPv4 filter in the IPv4 Filter panel.
2. Select an IPv6 filter in the IPv6 Filter panel.

9

To enable table-based color-aware ingress classification, select the Enable Table Classification For VPLS parameter, then select a 7210/7250 DSCP classification policy as the Routed Override QoS policy. See 50.23.2 “Table-based ingress classification on the 7210 SAS” (p. 1529).

10

To configure the egress QoS policies by selecting an QoS policy in the Egress QoS Policy panel.

11

Save the changes and close the forms.

END OF STEPS

79.99 To associate a local DHCPv4 or DHCPv6 server with a VPRN L3 access interface

79.99.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Local DHCP tab:

a. For local DHCPv4 servers, select a local DHCP server in the Local DHCP Server panel.

Note: You cannot associate a local DHCPv4 server to the L3 group Interface if the Administrative State parameter in the Local Proxy Service panel is up.

b. For local DHCPv6 servers, select a local DHCPv6 server in the Local DHCPv6 Server panel.

Note: To associate local DHCPv6 servers the IPv6 Allowed parameter must be enabled on the General tab.

5

Save the changes and close the forms.

END OF STEPS

79.100 To assign an NE DoS or DDoS protection policy to a VPRN L3 access interface

Note: See the procedure to configure an NE DoS protection policy in the *NSP System Administrator Guide* for more information about configuring and applying NE DoS protection policies.

79.100.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Port tab.

5

Select a DoS protection policy in the NE DoS Protection Policy panel.

6

Select a DDoS protection policy.

7

Save the changes and close the forms.

END OF STEPS

79.101 To configure residential subscriber management for a VPRN L3 access interface

79.101.1 Prerequisites

Residential subscriber management is supported on the 7450 ESS in mixed mode and 7750 SR.

79.101.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Subscriber Management tab.

5

Enable the SHCV Enabled parameter.

6

Configure the required parameters on the Subscriber Host Connectivity Verification panel.

7

Select an SHCV IPv4 policy.

8

Save the changes and close the forms.

END OF STEPS

79.102 To assign an IP address to a VPRN L3 access interface

79.102.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Addresses tab and click Create. The IP Address (Create) form opens.

5

Configure the required parameters.

The Broadcast Address Format parameter only appears if the IP Address parameter is set to an IPv4 address.

The parameters in the IPv6 panel only appear if the IP Address parameter is set to an IPv6 address.

6

Select a Track SRRP Instance, if required.

7

Save the changes and close the forms.

END OF STEPS

79.103 To configure IPv4 ICMP for a VPRN L3 access interface

79.103.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the ICMP tab and configure the required parameters.

5

Configure the required parameters in the Redirects panel.

6

Configure the required parameters in the Unreachables panel.

7

Configure the required parameters in the TTL Expired panel.

8

Save the changes and close the forms.

END OF STEPS

79.104 To configure IPv6 ICMP on a VPRN L3 access interface

79.104.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Ensure that the IPv6 Allowed parameter is enabled.

5

Click on the IPv6 ICMP tab and configure the required parameters.

6

Configure the required parameters in the Redirects panel.

7

Configure the required parameters in the Unreachables panel.

8

Configure the required parameters in the Packet Too Big panel.

9

Configure the required parameters in the Param panel.

10

Configure the required parameters in the Time Exceeded panel.

11

Save the changes and close the forms.

END OF STEPS

79.105 To assign an ICMP ping template to a VPRN L3 access interface

79.105.1 Before you begin

ICMP ping templates are used to populate values for ICMP ping tests when the tests are used on L3 interfaces to control the operational state. ICMP ping templates are supported for IPv4 only. See 90.51 “To configure an ICMP Ping template” (p. 3079) for information about configuring ICMP ping templates.

Note: Nokia recommends that you assign an NE DoS protection policy, configured with a protocol of ICMP-Ping-Check, to the affected interfaces; see 79.100 “To assign an NE DoS or DDoS protection policy to a VPRN L3 access interface” (p. 2675).

79.105.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service tree, expand Site→L3 Access Interfaces, right-click on the L3 access interface you need to modify, and choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the ICMP Ping Template tab.

5

Click Create, or choose the template in the list and click Properties. The Virtual Interface ICMP Template Config form opens.

6

Click Select to choose a template pointer for the template. The Select Template Pointer form opens.

7

Choose an ICMP Ping template from the list and click OK to close the Select Template Pointer form.

8

On the Virtual Interface ICMP Template Config form, configure the Admin State and Destination Address parameters for the ICMP Ping test.

9

Save the changes and close the forms.

END OF STEPS

79.106 To configure BFD for a VPRN L3 access interface

Note: If the IPv6 Allowed parameter on the General tab is enabled, you can configure BFD parameters for IPv6.

79.106.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the BFD tab.

5

In the IPv4 BFD Configuration panel, set the Admin Status parameter to Up and configure the required parameters.

6

In the IPv6 BFD Configuration panel, set the Admin Status parameter to Up and configure the required parameters.

7

To view local and remote session peers that are managed by the NFM-P, click on the BFD Session tab. A list of BFD current sessions on a router interface or an L3 interface appears.

8

Click on a session. The properties form for the session opens. View the following:

- BFD status
- protocol used
- local address
- remote address
- operational status and statistics

9

Save the changes and close the forms.

END OF STEPS

79.107 To configure ARP for a VPRN L3 access interface

79.107.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the ARP tab and configure the required parameters.

5

To configure static ARP:

1. Click Create. The Static ARP (Create) form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

6

Click on the Proxy ARP tab and configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

79.108 To configure neighbor discovery on a VPRN L3 access interface

79.108.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Set the IPv6 Allowed parameter to enabled.

5

Click on the Neighbor Discovery tab and click Create. The Neighbor Discovery (Create) form opens.

6

Configure the required parameters.

For the neighbor address to be communicated to the NE, you must configure an SNMP community on the site. See 79.17 “To configure an SNMP community on a VPRN site” (p. 2551) .

7

Save the changes and close the form.

8

Click on the Proxy ND tab and configure the required parameters.

Note: Do not leave an empty policy parameter between two configured policy parameters. For example, do not configure the Policy 1 and Policy 3 parameters and leave the Policy 2 parameter unconfigured, or the NFM-P reorders the policies and moves the policy specified for the Policy 3 parameter to the Policy 2 parameter.

9

Click on the Secure ND tab and configure the required parameters.

10

Save the changes and close the forms.

END OF STEPS

79.109 To configure IPv4 DHCP for a VPRN L3 access interface

79.109.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the DHCP tab and configure the required parameters.

The Lease Populate parameter is configurable when the Enable parameter is enabled.

5

Select a subscriber authentication policy in the Subscriber Authentication Policy panel.

6

Configure the required parameters in the Option 82 panel.

The Remote ID String parameter is configurable when the Remote ID is set to Remote ID String.

When the Relay Unicast Message parameter is set to Renew or Release Update Source IP, the GI address parameter can be configured as any local configured address in the same routing instance as the GI address for the DHCP relayed messages. If the Relay Unicast Message parameter is set to None, the GI address is restricted to the IP address configured on the subscriber interface.

7

Configure the parameters in the Vendor Specific Option panel.

8

Configure the parameters in the GI-Address panel.

9

Select a Python policy, if required.

10

Click on the Server tab and configure the required parameters in the DHCP Servers and Local Proxy Server panels.

The Number of Days, Number of Hours, Number of Minutes, Number of Seconds, and Lease Time Override parameters are configurable only when the Lease Time parameter is set to Specified Time Period.

11

Save the changes and close the forms.

END OF STEPS

79.110 To create a VRRP instance on a VPRN L3 access interface for a virtual router

79.110.1 Prerequisites

You must know the VRID for an existing virtual router and ensure that the interface is a member of the same subnet as the virtual router.

See Chapter 37, “VRRP” for configuration information about VRRP instances and virtual routers.

The following configurations are required for the operation of the IPv6 VRRP instance:

- Two tabs are available under the VRRP tab, one for IPv4 instances and the other for IPv6 instances. You can only create an IPv6 VRRP Instance if you enable the IPv6 Allowed parameter on the General tab of the VPRN L3 Access Interface (Edit) form.
- The Link Local Address on the parent interface must be set to preferred or to disable DAD, depending on your node release, and configured as one of the backup addresses (or same subnet) for the IPv6 VRRP instance. The Admin Link Local Address and related parameters on the General tab of the VPRN L3 Access Interface (Edit) form must be set accordingly.
- The IPv6 address on the parent interface must be set to disable DAD to be used as a backup address (on same subnet) for the IPv6 VRRP instance. The IP Address and Disable DAD parameters in Step 5 of 79.102 “To assign an IP address to a VPRN L3 access interface” (p. 2677) must be set accordingly.
- The Send Advertisement and Use Virtual MAC Address parameters must be enabled in Step 6 of 79.112 “To configure router advertisement on a VPRN L3 access interface” (p. 2687) for the router advertisement on the parent interface.

79.110.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the VRRP tab and click Create. The VRRP Instance (Create) form opens.

5

Configure the Virtual Router ID parameter.

6

Perform Step 3 to Step 15 of 37.4 “To create and configure a VRRP instance” (p. 1283) .

You can use the VR Instances tab to create, modify, and view VR instances.

7

Save the changes and close the forms.

END OF STEPS

79.111 To configure anti-spoofing filters for a VPRN L3 access interface

79.111.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Anti-Spoofing tab and configure the required parameters.

5

Configure the required general parameters.

The ARP Populate parameter is configurable when all of the IP addresses of the defined static hosts on the interface are in one of the subnets configured for the interface.

6

To configure static subscriber host entries, if subscriber entries are not available through DHCP lease management:

1. Click on the Static Hosts tab.

2. Click Create. The Access Interface Anti-Spoofing Static Host Display (Create) form opens.
3. Configure the required parameters.

   Specify at least one IP address or MAC address for each static host. The values specified for the Anti-Spoofing and ARP Populate parameters determine the type of address entry that is required for the static host. For example, when you set the Anti-Spoofing parameter to Source Ip Addr, you must specify at least the IP address for the static host.

   Note:

   You can configure a static host on a SAP only when no static ARP entries exist on the IP interface.

   When the ARP Populate parameter is enabled, the IP address of the new static host must be in one of the subnets that is configured for the interface in Procedure 79.102 “To assign an IP address to a VPRN L3 access interface” (p. 2677) .

7

Save the changes and close the forms.

END OF STEPS

79.112 To configure router advertisement on a VPRN L3 access interface

79.112.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Advertisement tab.

5

Click Create to add a router advertisement entry. The Router Advertisement (Create) form opens.

6

Configure the required parameters.

If you are configuring the L3 interface for an IPv6 VRRP instance, then the Send Advertisement and Use Virtual MAC Address parameters must both be enabled.

7

To configure RDNSS advertisement options:

1. Click on the DNS Options tab and click Create. The DNS Options form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

8

To configure router advertisement prefixes:

1. Click on the Prefix tab and click Create. The Router Advertisement Prefix form opens.
2. Configure the required parameters.
3. Save the changes and close the form.

9

Save the changes and close the forms.

END OF STEPS

79.113 To assign an ANCP policy to a VPRN L3 access interface

79.113.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the ANCP Static Map tab.

5

Click Create. The ANCP Static Map (Create) form opens.

6

Configure the ANCP String parameter.

7

Select an ANCP policy in the ANCP Policy panel.

8

Save the changes and close the forms.

END OF STEPS

79.114 To specify QoS policy overrides on a VPRN L3 access interface

79.114.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Click on the Override Policy Items tab.

Note: The Override Policy Items tab contains a number of tabs. However, the tabs that are displayed depend on the port type that you have chosen for this interface.

- If you configured a non-HSMDA port, then the Access Ingress Queues, Access Egress Queues, Ingress Policer, and Egress Policer tabs are active.
- If you configured an HSMDA port, then the Access Ingress Queues, Access Egress HSMDA Queues and Ingress Policer tabs are active.

Configure the policy overrides, as described in 50.97 “To configure QoS policy overrides on an L2 or L3 access interface” (p. 1654) .

To configure meter overrides on a 7210 SAS, see 50.98 “To configure QoS policy overrides on access ingress meters for the 7210 SAS” (p. 1657) .

To configure queue overrides on a 7210 SAS, see 50.99 “To configure QoS policy overrides on access ingress queues for a 7210 SAS-X” (p. 1659) .

5

Save the changes and close the forms.

END OF STEPS

79.115 To configure DHCPv6 on a VPRN L3 access interface

79.115.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface and choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Set the IPv6 Allowed parameter to enabled.

5

Click on the DHCPv6 tab and configure the required parameters on the DHCPv6 Relay — General tab.

6

Click on the Server tab and configure Server 1 through Server 8 parameters.

7

To configure the interface name for each DHCPv6 server that you configured, select a configured interface in the Zone Index panel.

8

Click on the DHCPv6-Prefix tab and click Create. The DhcpRelayV6PrefixDelegation (Create) form opens.

9

Configure the required parameters.

10

Save the changes and close the forms.

END OF STEPS

79.116 To associate a Multi-Chassis shunting profile to a VPRN L3 access interface

79.116.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a VPRN service and click Properties. The VPRN Service (Edit) form opens.

3

On the service navigation tree, expand Site→L3 Access Interfaces, right-click on an L3 access interface and choose Properties. The VPRN L3 Access Interface (Edit) form opens.

4

Select a profile in the Multi-Chassis Shunting Profile panel; see 27.31 “To configure a Multi-Chassis shunting profile on a base routing instance or VPRN routing instance” (p. 876) .

Note: Only the shunting profile created on a VPRN routing instance can be selected.

5

Save the changes and close the forms.

END OF STEPS

79.117 To start or stop the ignore SAP port state tool on a VPRN interface

79.117.1 Purpose

Use the ignore SAP port state tool to bypass the checking of the physical port operational state if it is down during operational checks on the NE. The tool is available for 7x50 NEs on L3 and subscriber interfaces of IES and VPRN services. The interface can be IPV4/IPV6 or dual stack. An SAP must be attached.

If this procedure is performed when the IP interface is operationally up the command will be accepted but will enter a pending state. It will not become active unless the port state becomes non-operational.

79.117.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The VPRN Service (Edit) form opens.

3

Perform one of the following:

On the service navigation tree, expand Site→Subscriber Interfaces→Group Interface and right-click on a group interface.

a. On the service navigation tree, expand Site→L3 Access Interfaces and right-click on an L3 access interface.

b. On the service navigation tree, expand Site→Subscriber Interfaces→Group Interface and right-click on a group interface.

4

Choose Start Ignore SAP Port State or Stop Ignore SAP Port State.

END OF STEPS

80 SPB service management

80.1 Overview

80.1.1 Purpose

Shortest Path Bridging (SPB) is supported on OS 6900 and OS 10K NEs, Release 7.3.2 and later. SPB uses the Provider Backbone Bridge (PBB) network model to encapsulate (using IEEE 802.1ah headers) and tunnel customer traffic through the network backbone. The shortest path trees upon which the PBB network infrastructure operates are determined using a version of the Intermediate System-to-Intermediate System (IS-IS) link state protocol that supports TLV extensions for SPB (ISISSPB).

80.1.2 Contents

SPB service management procedures

80.2 Workflow to create SPB services (OmniSwitch)

80.2.1 Overview

The following workflow describes the sequence of high-level tasks required to create an SPB service on OmniSwitch devices.

80.2.2 Stages

1

Create an OmniSwitch Backbone VLAN service. See 80.3 “To create an OmniSwitch Backbone VLAN service” (p. 2695) for more information.

2

Configure an SPB control instance protocol. See 80.4 “To configure an SPB control instance protocol” (p. 2696) for more information.

3

Configure an SPB network interface. See 80.5 “To configure an SPB network interface” (p. 2696) for more information.

4

Configure an OmniSwitch Ethernet service L2 profile. See 80.6 “To configure an OmniSwitch Ethernet service L2 profile” (p. 2697) for more information.

5

Configure an SPB access interface. See 80.7 “To configure an SPB access interface” (p. 2698) for more information.

6

Create an SPB service. See 80.8 “To create an SPB service” (p. 2699) for more information.

7

Associate an access interface with an SPB service. See 80.9 “To associate an access interface with an SPB service” (p. 2700) for more information.

80.3 To create an OmniSwitch Backbone VLAN service

80.3.1 Steps

1

Choose Create→Service→VLAN from the NFM-P main menu. The VLAN Service (Create) form opens.

2

Choose a customer to associate with the VLAN.

3

Configure the required parameters.

The Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Click on the VLAN tab and set the Application parameter to Backbone VLAN.

5

Select a group to associate with the VLAN in the Group panel.

Note: All network elements within a backbone infrastructure must belong to the same VLAN group.

6

On the service tree, right-click on VLAN Service and choose Create VLAN Site. The Select Network Elements - VLAN Service form opens.

7

Choose one or more sites and click OK. The form closes.

8

Save your changes and close the form.

END OF STEPS

80.4 To configure an SPB control instance protocol

80.4.1 Steps

1

On the equipment tree, right-click on a Backbone VLAN site and choose Properties. The Network Element (Edit) form opens.

2

Click on the Backbone VLAN Parameters tab.

3

Configure the ECT Algorithm and Tandem Multicast Mode parameters.

Note: The ECT Algorithm parameter must be configured with the same value for all Backbone VLAN sites.

4

Click on the Globals tab.

5

Click on the Service tab and select an SPB control instance.

6

Click OK. The SPB Control Instance (Edit) form opens.

7

Select the required Backbone VLAN service in the Service panel.

8

Save your changes and close the forms.

END OF STEPS

80.5 To configure an SPB network interface

80.5.1 Steps

1

On the equipment tree, right-click on a Backbone VLAN site and choose Properties. The Network Element (Edit) form opens.

2

Click on the Network Interfaces tab.

3

Click Create. The Network Interface (Create) form opens.

4

Select a port in the Interface panel.

5

Close the form.

6

Click on the General tab and set the Administrative State parameter to Up.

7

Save your changes and close the form.

END OF STEPS

80.6 To configure an OmniSwitch Ethernet service L2 profile

80.6.1 Steps

1

Choose Policies→Ethernet→AOS Ethernet Service from the NFM-P main menu. The Manage Ethernet Service Policies form opens.

2

Click Create and choose Create L2 Profile. The Layer-2 Profile, Global Policy (Create) form opens.

3

Configure the required parameters.

4

Click OK. The form closes.

5

To create more L2 profiles repeat Step 2 to Step 4 .

6

Close the form.

End of steps

80.7 To configure an SPB access interface

80.7.1 Steps

1

Perform one of the following:

a. To configure an OmniSwitch LAG as an SPB access interface, continue to Step 2 .

b. To configure an OmniSwitch port as an SPB access interface, go to Step 6 .

2

Create an OmniSwitch LAG as described in 13.19 “To configure an OmniSwitch LAG” (p. 440) .

3

On the equipment tree, expand Device→LAG.

4

Right-click on the LAG object and choose Properties. The LAG (Edit) form opens.

5

Go to Step 9 .

6

On the equipment tree, right-click on an OmniSwitch object and choose Properties. The Network Element (Edit) form opens.

7

On the equipment tree, expand Shelf→Port.

8

Click on the required port. The Physical Port (Edit) form opens.

9

Set the Mode parameter to Access and click Apply to confirm.

10

Set the SPB Service Mode parameter to Enabled.

11

Configure the Description and Vlan Xlation Mode parameters in the Service Access Info panel.

12

Click Apply to confirm.

13

Click on the Policies tab and select a UNI profile in the UNI Profile panel.

Note: If no UNI profiles exist, perform 61.2 “To configure an OmniSwitch Ethernet service UNI profile” (p. 1815) to create one.

14

Select an L2 profile in the L2 Profile panel.

Note: If no L2 profiles exist, perform 80.6 “To configure an OmniSwitch Ethernet service L2 profile” (p. 2697) to create one.

15

Close the form.

End of steps

80.8 To create an SPB service

80.8.1 Steps

1

Choose Create→Service→SPB from the NFM-P main menu. The SPB Service (Create) form opens.

2

Select a customer to associate with the SPB.

3

Configure the required parameters.

The Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

The SVC Mgr Service ID parameter is configurable when the Auto-Assign ID parameter is disabled.

4

Select a Backbone VLAN service to associate with the SPB in the SPB panel.

5

Configure the ISID parameter.

6

On the equipment tree, right-click on Sites and choose Create SPB Site. The Select Network Elements - SPB Service form opens.

7

Choose one or more sites and click OK. The form closes.

8

Save your changes and close the form.

END OF STEPS

80.9 To associate an access interface with an SPB service

80.9.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the SPB service to be associated with an interface and click Properties. The SPB - *Service Name* (Edit) form opens.

3

On the equipment tree, click on the site to which you want to associate the new access interface; expand the entries for that site.

4

Right-click on L2 Access Interfaces and choose Create SPB L2 Access Interface. The SPB L2 Access Interface (Create) form opens.

5

Configure the required parameters.

6

Click on the Port tab and select an SPB L2 access interface.

Note: The form lists only ports and LAGs that have been configured as an SPB access interface. See 80.7 “To configure an SPB access interface” (p. 2698) for more information.

7

Configure the required parameters.

8

Save your changes and close the forms.

END OF STEPS

81 PW routing and dynamic MS-PW service management

PW routing and dynamic MS-PW service management

81.1 Overview

81.1.1 General information

Service providers require a solution for interconnecting customer sites that span different domains, such as MAN-WAN or Inter-AS (in the same Service Provider or Inter-Provider). The NFM-P supports two methods for interconnecting customers sites:

- PW routing
- MS-PW routing

81.1.2 PW routing

The NFM-P supports inter-domain services for VLL through the use of VLL spoke switching, which allows creation of a VLL service by cross-connecting two spoke SDPs. However, this requires statically configuring the PW switching points at the gateway S-PEs between domains. See 76.2.2 “VLL spoke switching” (p. 2104) in Chapter 76, “VLL service management” for more information.

Static PW routing is supported on the 7210 SAS, 7450 ESS, 7750 SR, and 7950 XRS. Support for 7210 SAS NEs varies depending on the chassis type and release; see the NE documentation for information.

81.1.3 MS-PW routing

Multi-segment pseudowire routing provides solutions for inter-domain services by using dynamic MS-PW routing and signaling, in which the switching points are automatically instantiated in the Switching-Provider Edge (S-PE) NEs. The path of the MS-PW is dynamically signaled end-to-end by T-LDP, using PW routing information stored in the S-PEs by MP-BGP. Per-PW configuration is only required at the endpoints of the MS-PW in the Termination-Provider Edge (T-PE) NEs.

Dynamic MS-PWs are characterized by the following:

- They are supported for VLL Epipe services
- Dynamic and static routes, as well as explicit paths for MS-PWs are supported
- Dynamic MS-PWs may be established across LDP or RFC 3107 labeled BGP SDPs
- Dynamic MS-PWs may be used as a part of a set of PWs for PW redundancy, including MC-LAG. Diverse routes for the active and standby MS-PWs can be configured by using explicit paths, or dynamically by using a BGP route distinguisher.

- Diverse routes for the active and standby MS-PWs can be achieved by using explicit paths or dynamically using a BGP route distinguisher.

Dynamic MS-PWs are supported on the following NEs:

- 7210 SAS
- 7450 ESS and 7750 SR chassis types that support Epipe service creation
- 7950 XRS

Support for 7210 SAS NEs varies depending on the chassis type and release; see the NE documentation for information.

81.1.4 Dynamic MS-PW services

Spoke SDP FEC configuration will bind a service to an existing Service Distribution Point (SDP), using a dynamic MS-PW. When using dynamic MS-PWs, the particular SDP to bind to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under the Spoke SDP FEC. The selected SDP will terminate on the first hop S-PE of the MS-PW. Therefore, an SDP should be defined that can reach the first hop of the MS-PW. The NE creates a spoke SDP binding to associate an SDP with a service. If the required SDP does not exist, the NE creates one based on the parameters specified in the spoke SDP FEC, for example, the SAII/TAII, Path, or PW template. The Creation Mode attribute for such an SDP indicates Multi-Segment PW.

This differs from the regular spoke SDP binding creation in that it creates a spoke SDP binding using a PW with the FEC128. However, the Spoke SDP FEC enables PWs with other FEC types to be used.

Spoke SDP FEC configuration is currently only applicable to an Epipe site. It can be configured under the Epipe service or at a service endpoint of an Epipe site, with or without ICB. The creation under an Epipe service is demonstrated in 81.4 “To configure a dynamic MS-PW service using PW routing” (p. 2710) .

81.2 Workflow to configure PW routing and dynamic MS-PW services

81.2.1 Purpose

The following workflow lists the high-level steps required to configure PW routing and dynamic MS-PW services. Enabling MS-PW support requires configuration in the following areas:

- network commissioning
- SR NE pre-provisioning
- dynamic MS-PW service
- optional PW routing and MS-Service tasks

81.2.2 Stages

Network commissioning

1

To use PW routing to establish dynamic pseudo-wires from source to destination, without using static routes and a specified path:

a. Enable BGP on the routing instance of the required node. See 28.29 “To enable BGP on a routing instance” (p. 916) for more information.

b. Configure the BGP Family parameter to include the options MS-PW and IPv4. This must be done on the BGP site, associated BGP groups, and BGP peers.

See the following procedures:

- BGP Site configuration: See 28.31 “To configure global-level BGP” (p. 918)
- BGP Peer Group configuration: See 28.32 “To configure peer-group-level BGP” (p. 922)
- BGP Peers configuration: See 28.33 “To configure peer-level BGP” (p. 927)

2

If more than one equal cost route is required for the PW routing, configure the following:

a. In the routing instance for the required NE, set the Maximum Number of Equal Cost Routes parameter to a value greater than 1. See 27.2 “To configure a routing instance or a VRF instance” (p. 826) .

b. For the required BGP Site, set the Multi Path parameter to a value greater than 1. See 28.31 “To configure global-level BGP” (p. 918) .

3

Configure a routing policy statement with the Family parameter having the MS-PW option enabled. This will then be exported under the BGP site. See 54.5 “To configure a routing policy statement” (p. 1745) .

4

If path diversity is required for primary/standby MS-PWs, perform the following:

a. Configure a routing policy statement having multiple statement entries, each associated with a different community, from a different neighbour, with family including MS-PW and a local preference. The communities configured here are used further on in the procedure when configuring a MS-PW local prefix. See 54.5 “To configure a routing policy statement” (p. 1745) for more information.

b. Import the above policy under the BGP site. See 28.31 “To configure global-level BGP” (p. 918) .

5

If there are ASBRs, set Next Hop Self to True under each peer group. For example, set one for the inner-AS and one for the external-AS. Without this setting, the SDPs created by MS-PW routing only reach to the BGP interface's IP address, but not the system ID. See 28.32 “To configure peer-group-level BGP” (p. 922) . Refer also to Chapter 79, “VPRN service management” for information on ASBRs.

6

If all T-PEs and S-PEs are within the same AS, then the S-PEs must be configured as route reflectors (RR). Perform the following:

a. Set the Cluster ID parameter on the BGP site or BGP peer group that includes the T-PEs as BGP peers on the S-PE. See Step 3 in 28.31 “To configure global-level BGP” (p. 918) . Two T-PEs can not be configured as BGP peers to each other. They must be configured as BGP peers of the RR.

b. Configure a routing policy statement with Advertise Next Hop Self set to true and Action set to Accept, and then export this policy under the BGP site. See 54.5 “To configure a routing policy statement” (p. 1745) .

This completes the summary of essential network commissioning requirements.

SR NE pre-provisioning

7

Configure PW routing. See 81.3 “To configure PW routing on an NE” (p. 2708).

Configure a dynamic MS-PW service

8

Configure a dynamic MS-PW service using PW routing. See 81.4 “To configure a dynamic MS-PW service using PW routing” (p. 2710) .

Optional PW routing and MS-Service tasks

9

As required, perform one or more of the following:

a. On-demand resync of the MS-PW routing table for a specified NE is supported and can be accessed from a number of places in the NFM-P. See 81.5 “To display MS-PW routing tables” (p. 2712) for information on displaying MS-PW routing tables.

b. Conduct various spoke SDP FEC operations from the endpoint of an Epipe site. See 81.6 “To perform spoke SDP FEC operations from an Epipe site endpoint” (p. 2712) .

c. Discover and manage MS-PW switching sites. See 81.7 “To discover and manage MS-PW switching sites” (p. 2713) .

d. Run an OAM VCCV Ping or Trace validation test for the MS-PW service. See 81.8 “To

conduct MS-PW routing OAM tests” (p. 2713) .

PW routing and dynamic MS-PW service management procedures

81.3 To configure PW routing on an NE

81.3.1 Steps

1

Right-click on an NE in the navigation tree and choose Properties. The Network Element (Edit) form opens.

2

Click on the Globals tab, then on the Service tab.

3

Click on the PW Routing tab.

4

Configure the general parameters.

The SPE Address must be unique in the network. This parameter must be configured on all T-PEs and S-PEs that are involved in this PW routing.

The Number of BGP Routes, Number of Static Routes, Number of Local Routes, and Number of Host Routes attributes are for information only and can be resync-ed from the NE.

5

To create one or more local prefixes:

1. Click on the Local Prefixes tab.
2. Click Create. The PW Routing NE Local Prefix (Create) form opens.
3. Configure the Local Prefix parameter.

   You must configure this parameter to enable the MS-PW routing configuration on a T-PE node. This Local Prefix is also used in other portions of the MS-PW configuration, such as setting static routes on the remote T-PE or intermediate S-PEs, and setting Source Attachment Individual Identifier (SAII) and Target Attachment Individual Identifier (TAII) addresses on a spoke SDP FEC on the T-PE.
4. Click Create in the BGP RD panel. The PW Routing NE BGP RD (Create) form opens.

   A route distinguisher must be configured to enable this local prefix to be advertised by MP-BGP to the network.
5. Configure the Route Distinguisher parameter.
6. Configure the BGP Community parameter.

The BGP Community can only be selected from the list. Members were created as part of the routing policy configuration in 54.5 “To configure a routing policy statement” (p. 1745) and importing the routing policy under the BGP site in 28.31 “To configure global-level BGP” (p. 918) .

7. Save the changes and close the forms.

   You can create up to four route distinguishers per local prefix. You can create up to 16 local prefixes per NE. The local prefixes you create must be unique in the network.

   At this point the NE can start to advertise the defined local prefix to its BGP peers.

6

To configure one or more static routes:

1. Click on the Static Routes tab.
2. Click Create. The PW Routing NE Static Route Configuration (Create) form opens.

   A static route can be configured on both the T-PE and S-PE, based on your requirements.
3. Configure the Target T-PE’s PW Address parameter to specify the local prefix defined in the remote T-PE. It can be selected from a managed T-PE or S-PE, or directly entered.
4. Configure the Next Hop parameter to specify the system IP address of the next-hop S-PE or T-PE. It can be selected from a managed T-PE or directly entered.
5. Save the changes and close the form.

   At this point the NE can start to advertise the defined static route to its BGP peers.

7

To configure an explicit path:

1. Click on the Configured Paths tab. All configured paths are listed.

   The path should be configured on the T-PEs, which is then used in the spoke SDP FEC to establish a pseudo-wire through the specified path.
2. Click Create. The PW Routing NE Path Configuration (Create) form opens.
3. Specify a Path Name.
4. Click Create in the Hops panel. The PW Routing Hop (Create) form opens.
5. Configure the Hop ID parameter.
6. Configure the Hop Address parameter.

   This specifies the system IP address of the next-hop S-PE or T-PE (where the T-LDP session to a given S-PE terminates). It can be entered directly or chosen from the managed S-PE or T-PE by clicking the Select button.
7. Save the changes and close the forms.

8

Save the changes and close the forms.

END OF STEPS

81.4 To configure a dynamic MS-PW service using PW routing

81.4.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Configure the filter criteria for Service→VLL (VLL)→Epipe Service→(Epipe). A list of services appears at the bottom of the Manage Services form.

3

Choose the required VLL Epipe service that contains the T-PE site and click Properties. The Epipe Service (Edit) form opens.

4

Click on the Spoke SDP FECs tab.

5

Click Create. The Spoke SDP FEC (Create) form opens.

Only one spoke can be configured on the terminating site, including both regular spoke SDP bindings and spoke SDP FECs.

6

Configure the required parameters.

The Auto Config parameter enables automatic endpoint configuration. This allows the configuration of a spoke SDP endpoint without specifying the TAII associated with that spoke SDP. In this mode, the far end T-PE actively initiates MS-PW signaling and will send the initial label mapping message using T-LDP, while the T-PE for which Auto Config is specified will act as the passive T-PE. The Auto Config parameter cannot be enabled if the Signaling Type is set to Master or if the TAII Address is specified.

If you are going to use SAII and TAII addresses instead of Auto Config, then you must decide which method will be used for the PW routing.

- For dynamic PW routing, the local prefix associated with the SAII/TAII address should be configured with a route distinguisher, with an optional community (see 81.3 “To configure PW routing on an NE” (p. 2708)). The SAII and TAII addresses must be unique in the network.
- For static routing, the local prefix should be associated with static routes configured under PW routing (see Step 5 and Step 6 in 81.3 “To configure PW routing on an NE” (p. 2708)). A spoke SDP FEC cannot be turned up if no local prefix is configured under PW routing.
- When using a specified path, the Path parameter must be configured. The associated Select button allows you to select a path that will be used for this spoke SDP and which is

configured under PW Routing. If no path is configured, then each next hop of the MS-PW used by the spoke SDP will be chosen locally at each S-PE and T-PE, using dynamic PW routing.

7

If redundancy is required, select an Endpoint.

8

Configure the required parameters.

9

To configure a return spoke SDP FEC:

1. Click on the Return Spoke SDP FEC tab.

   This tab is displayed only when the TAII Address configured in Step 6 was set by selection for the current spoke SDP FEC.

2. Configure the PW ID parameter for the return spoke SDP FEC and click Apply.

   When a valid PW ID for the return spoke SDP FEC is entered, additional parameters are displayed. If the PW ID is not valid, then the return spoke SDP FEC is not created.

3. Configure the required parameters.

4. Select an Endpoint.

   When the return spoke SDP FEC is created, it will have corresponding attributes such as the SAII and TAII Address values and AC ID cross-matching with the current spoke SDP FEC.

10

Save the changes and close the forms.

When a spoke SDP FEC is successfully created, it will be listed under the Spoke SDP FECs tab for this Epipe site.

The Spoke SDP Associated tab lists the spoke SDP bindings created by using this Spoke SDP FEC for MS-PW routing. All such spoke SDP bindings have their Creation Mode attribute set as Multi-Segment PW. The related SDP FEC Auto-bind tab shows the associated SAII or TAII defined in the Spoke SDP FEC.

The Faults tab displays any mis-configuration alarms associated with this spoke SDP FEC.

When the configuration for both T-PEs is complete, the NEs can start negotiating and establishing the spoke SDP bindings from T-PE to S-PE and from S-PE to T-PE.

END OF STEPS

81.5 To display MS-PW routing tables

81.5.1 Steps

1

Right-click on an applicable NE from one of the following locations and choose Show MS-PW Routes→BGP|Service:

- Navigation tree
- Physical Topology map
- Service Tunnel map

The Service Routing rtr BGP or Service Routing rtr Service form opens. The requested routing table information is presented, along with a detailed status/error report.

2

Close the Service Routing rtr report form.

END OF STEPS

81.6 To perform spoke SDP FEC operations from an Epipe site endpoint

81.6.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the required VLL Epipe service that contains the primary and backup spoke SDP FECs and click Properties. The Epipe Service (Edit) form opens.

3

In the service navigation tree, expand Site→Spoke SDP FECs, right-click on the required spoke SDP FEC and choose:

- Force Switchover
- Clear Forced Switchover
- Turn Up
- Shut Down
- Delete
- Resync
- Properties

The Forced Switchover and Clear Forced Switchover options are only applicable where the Epipe site has been configured for redundancy, as detailed in 81.4 “To configure a dynamic

MS-PW service using PW routing” (p. 2710) . The switchover applies to the primary and backup spoke SDP FECs configured on the Epipe site.

4

Close the forms.

END OF STEPS

81.7 To discover and manage MS-PW switching sites

81.7.1 Purpose

With PW routing configured on T-PEs and S-PEs, and spoke SDP FECs configured on the T-PEs, the NE will establish the pseudo-wires between T-PEs. This includes creating switching sites between two T-PEs. A switching site for MS-PW is a read-only site (nothing can be modified) and it cannot be deleted. It is controlled by the NE.

Use the following method to discover such sites.

Note: The switching sites and spoke SDP bindings created by MS-PW routing are tagged with their Creation Mode attribute indicating Multi-Segment PW (as opposed to Manual for non-MS-PW entities).

When a switching site is discovered, the associated spoke SDP bindings under it are also discovered and re-synced. The NFM-P will create a composite service and connector(s) if there are matching spoke SDP bindings found in other Epipe services (based on cross-matching of the spoke SDP ingress/egress labels). The related service is added to the newly-created composite service.

81.7.2 Steps

1

Conduct an RCA Audit to discover switching sites. This is done within the Epipe service that contains the T-PEs having MS-PW configured. The audit will find all MS-PW switching sites associated with T-PEs in this service and provide management solutions.

See 95.7 “To configure an RCA audit policy” (p. 3238) for details on configuring an RCA audit.

2

Perform 95.8 “To perform an RCA audit of a service or multiple services” (p. 3240) .

END OF STEPS

81.8 To conduct MS-PW routing OAM tests

81.8.1 Purpose

Perform this procedure to test the MS-PW routing configuration using OAM VCCV Ping and Trace tests.

81.8.2 Steps

1

To conduct a VCCV Ping test, see 90.8 “To create and run a VCCV ping OAM diagnostic test from the STM” (p. 3008) .

The test can be created either from the Service Test Manager or from a configured Epipe service, Epipe site, or a Spoke SDP FEC.

If you choose a spoke SDP binding created by MS-PW routing, the VCCV ping is specifically for MS-PW, and the SAII and TAII parameters will be auto-filled. This allows you to conduct a MS-PW VCCV Ping on an S-PE.

If the test is created from a configured Spoke SDP FEC, the test object will be auto-filled with the configured spoke SDP FEC.

Alternatively, if the test is created from the Service Test Manager, an Epipe service, or an Epipe site (both the termination site and the switching site), then you must select a spoke SDP FEC (only for T-PE) or a spoke SDP binding of type FEC 129 created by the MS-PW routing.

2

To conduct a VCCV Trace test, see 90.9 “To create and run VCCV trace OAM diagnostic test from the STM” (p. 3009) or 90.10 “To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM” (p. 3010) .

The test can be created either from the Service Test Manager or from a configured Epipe service, Epipe site, or a Spoke SDP FEC.

If the test is created from a configured Spoke SDP FEC, the test object will be auto-filled with the configured spoke SDP FEC.

Alternatively, if the test is created from the Service Test Manager, an Epipe service, or an Epipe site (both the termination site and the switching site), then you must select a spoke SDP FEC (only for T-PE) or a spoke SDP binding of type FEC 129 created by the MS-PW routing.

Note: When configuring either a VCCV Ping or VCCV Trace test for dynamic MS-PW from an S-PE, the Reply Type must be set to IP.
If the selected spoke SDP binding is created by MS-PW routing, you can only conduct a VCCV ping for MS-PW FEC129, not for FEC128.
If a PW template is used in the Spoke SDP FEC configuration, it must have its Enable Control Word parameter enabled, since both VCCV Ping and Trace require Control Word support.

3

To create a Test Suite, see 89.12 “To create an STM test suite” (p. 2951) .

A Test suite can be generated based on the selected policy for a specified VLL Epipe service.

If the selected policy has its Entity Type set to VLL Service, then the VCCV Ping and VCCV Trace tests that are generated will adhere to the following:

- No tests are generated for spoke SDP bindings created by MS-PW routing.

- Both VCCV Ping and VCCV Trace tests for MS-PW will be generated if there are spoke SDP FECs in the specified Epipe service.

If the selected policy has its Entity Type set to Service Connector, then there is also no test generated for the spoke SDP bindings created by MS-PW.

END OF STEPS

82 Network Group Encryption

82.1 Network Group Encryption overview

82.1.1

Network Group Encryption (NGE) is a mechanism for the end-to-end encryption of MPLS- or GRE-based traffic at the service level that does not require meshes of IPsec tunnels at the network layer. NGE is supported on 7705 SAR, 7705 SAR-Hm, VSR, and 7750 SR devices. With 7750 SR devices, only WLAN GW interface encryption is supported.

See the device documentation for information about the chassis variants that support NGE, and for detailed information about NGE operation on an NE.

NGE provides the following main levels of encryption to secure an IP/MPLS network:

- SDP encryption: GRE or MPLS
- VPRN encryption — MP-BGP-based VPRN-level encryption
- WLAN gateway interface encryption
- Router interface and Ethernet port encryption— Layer 3 and Layer 2 user plane and control plane encryption

  L2 encryption is supported for 7705 SAR NEs only
- PW template encryption for L2 service PW templates:
  - the template must be using auto-GRE
  - the template cannot be in use by more than 100 services
  - all sites in use by the template must support NGE

L3 routing interface encryption and L2 Ethernet port encryption are supported in an NGE domain; see 82.4 “NGE domains” (p. 2721).

NFM-P NGE management minimizes network downtime in the event of a catastrophic failure such as a natural disaster, and maintains network security functions and critical network traffic transmission during events such as unexpected NE reboots and link disruptions.

The NFM-P generates the keys for the participating NGE NEs. An NFM-P operator assigns encryption and authentication keys to a key group. The operator then associates NGE supported objects for inbound and outbound traffic with the key group, as required, and the NFM-P distributes the key group to each NE that hosts an associated NGE supported object. The keys in a key group are randomly generated by the NFM-P using the FIPS 140-2 standard.

NGE version support

The current NSP release only supports NGE version 2. To use NGE, all 7705 SAR NEs must be running Release 8.0 R4 or later, and all VSR and 7705 SAR-Hm NEs must be running Release 15.0 R4 or later.

82.1.2 NFM-P management of NGE

The NFM-P has a comprehensive suite of NGE functions that include the following:

- network-wide encryption configuration and management
- selective key distribution only to service sites associated with a key group
- automatic NGE configuration of sites added to an encrypted service
- key synchronization among participating NEs
- coordinated key updates without service degradation
- fault tolerance using NFM-P and NE redundancy functions
- alarm management
- statistics collection

The NFM-P uses SNMP to deploy general NGE attributes to NEs, and SSH2 sessions to configure the key values. You can use an existing SSH2 user account on each NE, or, to facilitate the tracking of key value configuration activity, you can use the UserNGE account. The NFM-P creates the account on each participating NGE NE and uses the account only for creating and updating key values. The NFM-P user activity log records all NGE configuration activity.

In a key group, an NFM-P operator specifies the service objects that are to be encrypted using the key group, then initiates the encryption. The NFM-P then deploys the key group to each NE associated with the service objects; for example, the two NEs associated with an SDP, or the sites in an MP-BGP VPRN service.

The NFM-P subsequently configures the key group for outbound traffic, then the key group for inbound traffic, on each NE. The NGE key group associations are displayed on the properties form of each associated service object.

The NFM-P ensures that the keys on all NEs in a key group are synchronized.

If a service object associated with a key group is deleted by the NFM-P or through a CLI, the NFM-P removes the object association from the key group.

A connectivity loss between the NFM-P and participating NEs does not affect the existing encrypted services.

82.1.3 Prerequisites for NGE management

To use NGE, all nodes must be configured to provide reachability to the NFM-P.

All NGE nodes must be managed by the same NFM-P.

Some potential scenarios to allow for a node to be reachable include:

- Out of band management – The node uses an out of band network to allow for NFM-P traffic. When the node comes back up, the NFM-P is reached using the out-band network and downloads the key group keys using this network.
- In-band using GRT (RI encryption disabled) – The node can establish routing and control plane functions when it comes back online. Based on established routing, the NFM-P re-discovers the node, and sends SNMP and SSH messages to the node to update the required key groups.
- In-band using GRT (RI encryption enabled) – IP exception ACLs are preconfigured in the

network to allow SNMP traffic through RI encryption enabled interfaces. Note that when the node that was off-line comes back up, it will not be able to send or receive any control plane packets as the RI interface NGE keys are old. The node will need static routes for the SNMP/SSH traffic to and from the NFM-P to allow these packets to be routed. If the network is over a L2/L3 service provider, then the IP exception ACLs are required on the off-line node itself and the gateway node. If the node is part of a private IP/MPLS network the IP exception ACLs are required on the expected bordering nodes to the node that goes off-line.

The default ACLs allow reachability to the NFM-P to then update key groups with new keys to bring up RI encryption and other services that were using old keys. The default ACLs do not need to be used for normal NFM-P management traffic after the NFM-P updates key groups.

- In-band using management VPRN service (preferred for 7705 SAR-Hm scenarios) – A default manually configured GRE-SDP with reachability to a gateway node and a default in-band management VPRN are configured, using the default GRE-SDP to reach the NFM-P. This VPRN could also be NGE encrypted using a key group that remains static (does not get rekeyed) for double encrypting the NFM-P traffic. When the node comes back on-line, this default VPRN service is available for the NFM-P to reach the node and update key groups.

82.2 Migration of CLI managed NGE to NFM-P managed NGE

82.2.1 Migration overview

If you have been using CLI for NGE management before setting up the NFM-P, you can migrate the NGE to the NFM-P after the NFM-P is installed and operational.

When the NFM-P is performing NGE migration from CLI to NFM-P managed NGE, the NFM-P tracks the key groups and encryption keys from the CLI NGE configurations and auto downloads key groups and encryption keys as necessary for all the services requiring encryption. In the end, the cleanup scheduled task removes un-used key groups and keys from CLI.

After the migration, do not make any NGE changes via CLI, including for troubleshooting purposes. If this occurs, the NFM-P generates a mismatch alarm and, if there is service impact, the NFM-P reverts the changes from CLI.

82.2.2 Migration scenarios

Migration from CLI NGE to NFM-P NGE is available in several scenarios. The process the user must perform in the NFM-P after node discovery differs slightly depending on the scenario.

See 82.11 “Workflow for migration of NGE management from CLI to NFM-P” (p. 2728) for details and prerequisites.

The following table shows the available scenarios and how the process differs in each.

Scenario	Configuration process in NFM-P after the node is discovered
Key group is the same on CLI and NFM-P VPRN and PW Template are configured on CLI but not in the NFM-P	Add the VPRN and the PW Template in the NFM-P, and encrypt
Key group and services are configured in CLI but not in the NFM-P	Create the key group, add the services, and encrypt

Scenario	Configuration process in NFM-P after the node is discovered
Key group and services are configured the same in both CLI and NFM-P	No further action is needed: the node will be discovered directly into the key group and services.
Key group and services are configured in CLI and NFM-P CLI key group is different from the key group in NFM-P, for example, the CLI key group is number 77 and the NFM-P key group is 60.	No further action is needed: the NFM-P will create key group 60 on the CLI and move the services to key group 60 on the NFM-P. If the migration leaves key group 77 unused on the CLI, the NGE cleanup task will remove it.

82.2.3 NSP alarms

If there is a difference between CLI and NFM-P NGE configuration, alarms will be generated as follows:

- The NSP raises a ConfigurationUnknown alarm if configuration is present on the CLI that is not present on the NFM-P

  ConfigurationUnknown alarms for global encryption label, key group, VPRN, and PW templates will be suppressed for first time node discovery. However, these alarms will be raised when there is subsequent addition of the key group or service association from CLI after the node has been discovered. Or if node resync is done after the node is already discovered by NFM-P and before the migration from CLI to NGE is completed.

  ConfigurationUnknown alarms for SDP, WLAN GW, L2 and L3 interfaces configured via CLI will still be present upon node discovery.
- The NSP raises a ConfigurationMismatch alarm if the NGE configuration on an NE is modified using CLI.

During migration from CLI, the alarms are suppressed.

82.3 Configuration

82.3.1 Configuration overview

You can configure NGE using an NFM-P GUI or OSS client.

82.3.2 Global encryption label

NFM-P NGE management requires a global encryption label that is used as a common NGE identifier by all participating NEs in the managed network. A global encryption label must be configured before any services are encrypted. A global encryption label is intended to be set once, for permanent use, and cannot be modified. The global encryption label can be deleted only if no key groups exist in the NFM-P and no local key groups exist on NEs.

If the deployment of a global encryption label fails, the NFM-P and NE labels do not match, or a global encryption label is detected during device discovery, an alarm is raised.

Note: An attempt to create a static MPLS ingress label is blocked if the label has the same value as the NGE group encryption label.

An attempt to create the group encryption label is blocked if the label has the same value as a static MPLS ingress label in the network.

If the NFM-P discovers an NE that supports NGE, and the NE has a static MPLS ingress label that matches the group encryption label, an alarm is raised.

82.3.3 Key groups

NGE deployment to one or more NEs requires a key group that contains the NGE keys. A key group defines the algorithms that the NFM-P uses to generate the encryption and authorization keys. A key group also contains a list of the current Security Associations (SAs) between the key group and the service objects that use the key group.

After you create a key group, you cannot modify the encryption and authentication algorithms; if such changes are required, you must create a new key group and delete the previous key group.

After the initial key group deployment, you can use a scheduled task for the regular and automatic replacement of the keys in the key group; see 82.8 “Key updates” (p. 2724) .

Keys are always encrypted when stored in the database; transfer of keys to the network elements is over a secure and encrypted connection.

Note: You cannot delete a key group if any SDPs, service objects, NGE domains, L3 router interfaces, or L2 Ethernet ports are associated with the key group.

82.3.4 NGE cleanup scheduled task

The NGE cleanup task is auto-created during NFM-P system initialization. The NGE cleanup task removes unused SPIs and key groups from the NE.

Key groups that are only present on CLI, but not on NFM-P, will be removed by the NGE cleanup schedule if there is no service association in the CLI.

When encryption is disabled on the last service using the key group of a site, the cleanup task will remove the key group from the site.

The NGE cleanup task runs every night at midnight by default. The task must not be deleted, however, it can be executed manually at any time.

Choose Manage→Network Group Encryption from the NFM-P main menu and choose Cleanup Scheduled Task from the drop-down menu to open the NGE Cleanup Scheduled Task.

82.4 NGE domains

82.4.1 NGE domain overview

An NGE domain is a group of L3 router interfaces, L2 Ethernet ports, or both, enabled for NGE. The domain is configured on a key group.

An NGE domain has the following components:

- general: the key group used by the NGE domain
- the list of gateway sites and gateway interfaces
  A gateway interface is required in order to add an unmanaged NE to a domain; see “NGE discovery” (p. 2722).
- the list of sites in the domain

A domain site is an NE that contains interfaces or ports that are participating in the domain. An NE can have interfaces or ports in more than one domain.

- the list of L3 router interfaces in the domain
- the list of L2 Ethernet ports in the domain
- the NGE encryption type (L3, L2, or both)

When NGE is configured, an interface can receive unencrypted packets or NGE encrypted packets from any configured key group on the NE, but no other type of IPsec formatted packet is allowed. If an IPsec packet is received on an NGE-enabled interface, it will not pass NGE authentication and will be dropped. Therefore, IPsec packets are prohibited from existing within the NGE domain without first being converted to an NGE packet. This delineates the boundary of the NGE domain and other IPsec services.

The maximum number of domains per network is 64. The domains can be in the same key group or distributed across multiple key groups.

82.4.2 NGE domain sites

When two or more sites are added to the domain, the NFM-P performs an auto-population: L3 router interfaces are added if their addresses are in the same subnet as interfaces already in the domain, and L2 Ethernet ports are added based on the router interfaces. If the L2 Ethernet port associated to the L3 router interface has a LAG, all member ports will be added to the domain.

Cellular interfaces on 7705 SAR-Hm NEs are populated when the site is added to the domain, regardless of subnet.

Ports and interfaces can also be added, encrypted, and deleted manually.

A port or interface cannot be added to more than one domain.

The following domain site encryption statuses are supported:

- partial: encryption enabled on some of the interfaces or ports of the same site
- Layer 3 encrypted: encryption enabled on all interfaces of the same site, but not on L2 ports
- Layer 2 and Layer 3 encrypted: encryption enabled on all interfaces and ports of the same site

NGE discovery

The NFM-P can add undiscovered NEs directly into an NGE domain. ACL IP exception filter policies are created automatically by the NFM-P and applied to a gateway interface to allow for NE discovery. A device discovery rule must be available; see 9.10 “Workflow for device discovery” (p. 288).

The NGE discovery process creates inbound and outbound ACL IP exception filters on the gateway NE. If the NSP server is restarted or switched to a secondary server during the NGE discovery process, the exception filters will become invalid and NGE discovery will fail. The NGE discovery process must be initiated again after the switchover or restart is completed. The NGE discovery process will be recovered from the failure state and the process will continue.

It is possible to delete a gateway interface during the NGE discovery process. The NGE discovery will continue but the ACL IP exception filter entries will not be removed automatically by the NFM-P after encryption has been enabled. You must manually delete the exception filter entries to resume

traffic between the NFM-P and the encrypted NEs. The exception filter to be deleted will have the IP address of the site most recently managed using NGE discovery as its Destination IP or Source IP. See 51.7 “To configure an ACL IP exception filter policy” (p. 1683).

82.4.3 L3 router interface encryption

Encryption for L3 interfaces in the NGE domain are configured with an inbound and an outbound key group: outbound packets are encrypted using the interface key group; inbound packets must be encrypted using the interface key group keys.

A security zone and NGE cannot be applied to the same interface.

L3 encryption exemption

It may be necessary for L3 packets to enter the NGE domain in clear text. To allow this, the NFM-P applies an ACL IP exception filter. See 51.1 “Filter policies” (p. 1663) for more information about filter policies.

82.4.4 L2 Ethernet port encryption

L2 encryption is configured on the Ethernet port. When L2 encryption is configured, all configured IS-IS and LLDP packets are encrypted using NGE.

L2 NGE encryption is enabled by configuring an outbound and inbound key group on the Ethernet port, similar to assigning key groups for router interface encryption, as described in 82.4.3 “L3 router interface encryption ” (p. 2723).

L2 Ethernet port encryption is not supported for VSR or 7705 SAR-Hm NEs.

82.5 PW template encryption

82.5.1 PW template encryption overview

PW templates allow you to automatically create SDP bindings. If the PW template is configured to use GRE delivery, there is no SDP configuration required, therefore no option to configure NGE on the services that make use of the SDP.

With PW template encryption, you can configure encryption on the PW template, providing encryption for the L2 services that will use the SDP. With PW template encryption, the NFM-P adds the NGE inbound and outbound encryption key-groups to a PW template.

This type of encryption can be useful for 7705 SAR-Hm networks that use the auto-bind function of the PW template for GRE-MPLS transport over the cellular interface.

82.6 Encryption for offline nodes

82.6.1 Encryption for offline nodes overview

A node is in offline mode if the NFM-P has lost connectivity to the node at the time of specific NGE action (rekey, encrypt, or disable encryption). A node is not in offline node if it responds to SSH and SNMP requests. Nodes in offline mode are skipped when enabling or disabling encryption.

The Administrative Status parameter of an SDP, service, or interface reflects the last encryption action on the node, and the Deployment Status parameter shows the current status of encryption.

For example, if encryption is enabled when the node is offline, the Administrative Status parameter will be updated to Encryption Enabled; however, since the encryption operation skipped the node, the Deployment Status remains Not Deployed. If encryption is disabled when the node is offline, the Administrative Status is updated to Encryption Disabled, but the Deployment Status remains Encrypting.

When the node is back online, the NFM-P will automatically apply the action based on the Administrative Status.

82.7 WLAN GW encryption

82.7.1 WLAN GW encryption overview

WLAN GW encryption allows the WLAN GW to terminate and originate WLAN AP traffic to and from a 7705 SAR-Hm or 7750 SR NE that is NGE encrypted. This provides a simple encryption approach using a single NGE key group on the WLAN GW interface to provide security to many WLAN APs.

The WLAN GW group interface is configured with the same key group for inbound and outbound.

When a WLAN group interface is added to or removed from a key group, all associated GRE-SDPs will be added or removed from the SDP tab of the WLAN GW automatically. When a new GRE-SDP with a far-end address the same as the WLAN GW address is added to the network, this GRE-SDP will be added to the SDP list under the WLAN GW in the key group automatically, with the same encryption status as the WLAN GW.

82.8 Key updates

82.8.1 Key updates overview

For increased security, Nokia recommends frequent replacement of the keys in a key group, which is called a rekeying operation. For each key group, you can configure a rekeying scheduled task that defines how often the NFM-P generates and deploys a new key set to each NE associated with the key group. The rekeying mechanism ensures that there is no service degradation during rekeying.

The rekeying scheduled task associated with the key group must be deleted before the key group can be deleted.

A rekeying scheduled task cannot use a schedule in which a delay is configured. However, an inter-nodal wait time can be configured as part of the rekey schedule. If an inter-nodal wait time is configured, the NFM-P will wait a specified amount of time after rekeying each NE.

Note: The execution of a rekeying scheduled task is skipped if a manual encryption operation using the same key group is in progress.

You can modify the wait time of a scheduled task. If a change to another parameter of a scheduled task is required, you must create a new scheduled task and delete the current scheduled task.

A rekeying operation has the following phases:

- Phase 1: deploy new key
- Phase 2: set new key as active outbound
- Phase 3: delete old key

If a rekeying operation fails, the NFM-P runs a cleanup operation the next time a rekey is triggered. The NFM-P checks and compares where the process failed, implements correction and inactive key cleanup, and continues with the rekey process.

After a rekeying operation, the NFM-P verifies that each key is correctly set by comparing the CRC checksums of the local and NFM-P key values. If the verification is delayed or unable to complete, an alarm is raised.

If a rekeying operation is unable to complete before the next rekeying operation is to begin, for example, when a large number of NEs are rekeyed using a schedule of high frequency, the NFM-P attempts the rekeying during the next scheduled task run.

You can view the results of rekeying scheduled task runs, which include the old and new CRC checksum values, from the properties form of the task; see 82.19 “To view rekeying results and statistics” (p. 2739).

The NFM-P also raises alarms for the following rekeying faults:

- failure to create an SA in a key group
- failure to delete an existing SA in a key group

82.8.2 Key updates with offline nodes

A node is in offline mode if the NFM-P has lost connectivity to the node at the time of specific NGE action (rekey, encrypt, or disable encryption). A node is not in offline mode if it responds to SSH and SNMP requests.

In some networks, it might be desirable to allow key groups to be rekeyed even though some of nodes in the key group are offline. For example, in a 7705 SAR-Hm network with vehicle-mounted deployments, it may not be possible to rekey the entire key group at once since some nodes (vehicles) will be periodically offline.

By default, the rekey procedure is halted if any node in the key group is offline.

Forced rekey

The Force Re-key option allows for the online nodes to be rekeyed when some nodes are offline. The offline nodes are skipped and online nodes are rekeyed. If rekey sites in a key group contain any site that belongs to a cellular domain, Force Re-key is skipped with a warning message in the server log.

When a node that has been skipped during a rekey operation comes back online, the node is rekeyed to apply the new key.

To allow nodes to learn new keys that were applied while the node was offline, the nodes require configurations that provide reachability to the NFM-P. The configuration cannot rely on any key groups that may be rekeyed while the node is offline. When the node comes back up, the node reaches the NFM-P, the NFM-P downloads the correct NGE keys for the key group to the node, and any services that were using the keys come back up.

Progressive rekey

A progressive rekey operation provides an alternative to forced rekeying if some nodes in the key group are offline. When offline nodes are encountered during the rekey operation, the NFM-P will wait for them to come back online, then proceed with the rekey operation. The use of progressive rekey prevents the need for the in-band configurations that provide reachability to the NFM-P for offline nodes if Force Re-key is in use.

The NFM-P provides the following information during the progressive rekey process:

- Current phase (Key Deployment, Key Activation, or Cleanup)
- Number of nodes remaining for key deployment
- Number of offline nodes pending rekey

If a progressive rekey operation is hung due to a node that remains offline, you can perform a force rekey operation for both the Key Deployment and Key Activation phases. While waiting in Key Deployment phase, you can also stop and fail the rekey process. When nodes come back online that did not receive the key-group keys, the NFM-P detects key groups that are not in sync and automatically updates the key group as needed.

You can add or remove sites from a key group during a progressive rekey operation. If a site is added, the new site is included in the rekey operation. If a site is deleted, the site information is deleted from the node.

82.9 NGE statistics

82.9.1 NGE statistics overview

You can collect NGE statistics for a key group on demand, and schedule NGE statistics collection using a MIB entry policy. The statistics are displayed on the Statistics tab of the Site properties form in a key group. See Chapter 9, “Device discovery” for information about configuring MIB entry policies.

NGE statistics can also be viewed and collected from the Statistics tab of the following object properties forms:

- card slot
- daughter card slot
- VPRN spoke-SDP binding
- network router interface
- Ethernet port

82.9.2 Rekeying operation statistics

The NFM-P monitors rekeying operations, and records statistics about the duration of rekeying operations for troubleshooting and assurance purposes. The statistics describe the duration of key replacement activities.

If the statistics reveal that rekeying activity is taking an increasing or excessively long time, investigation may be required to identify the cause of the latency. See 82.19 “To view rekeying results and statistics” (p. 2739) for information about viewing the rekeying operation statistics.

82.10 Workflow for NGE management using NFM-P

82.10.1 Stages

The following is the sequence of high-level actions required to manage NGE.

1

Enable SSH2 for the secure key transfers.

1. Ensure that SSH2 is enabled on each NE that is to participate in NGE; see 9.15 “To verify that SSH2 is enabled on a device” (p. 300).
2. Enable SSH2 host key persistence on devices that support host key persistence; see 9.16 “To enable SSH host key persistence on a device” (p. 301).
3. Create a mediation policy that specifies SSH2 as the CLI protocol; see 9.17 “To configure device mediation” (p. 301).

   If you specify SNMPv3 in the mediation policy, ensure that the associated SNMPv3 user has console access enabled.
4. Apply the SSH2 mediation policy as the Security Mediation Policy in each discovery rule associated with an NE that is to participate in NGE; see 9.23 “To configure a discovery rule” (p. 310).

2

Configure the global encryption label; see 82.12 “To create the NGE global encryption label” (p. 2730).

3

Create a key group to specify the security algorithms, encrypt objects, and create a rekeying scheduled task; see 82.13 “To create an NGE key group” (p. 2730).

4

As required, add objects to a key group, and encrypt the objects; see 82.14 “To add an object to a key group” (p. 2732).

5

As required, add NGE domains to a key group; see 82.15 “To create an NGE domain on a key group” (p. 2733).

6

As required, add managed sites to NGE domains, and apply encryption; see 82.16 “To configure an NGE domain” (p. 2734).

7

As required, add unmanaged sites to NGE domains, and apply encryption; see 82.17 “To add unmanaged sites to an NGE domain” (p. 2737).

8

If required, manually execute a rekeying scheduled task; see 82.18 “To manually execute a rekeying scheduled task” (p. 2739).

9

View the results of one or more rekeying operations; see 82.19 “To view rekeying results and statistics” (p. 2739).

10

Remove the NGE security from one or more objects; see 82.20 “To disable encryption on an object” (p. 2741).

82.11 Workflow for migration of NGE management from CLI to NFM-P

82.11.1 Stages

The following is the sequence of high-level actions required to migrate NGE management from CLI to NFM-P. Some steps may not be needed depending on the CLI and NFM-P configuration of the services and key group; see 82.2.2 “Migration scenarios” (p. 2719).

1

Configure the global encryption label; see 82.12 “To create the NGE global encryption label” (p. 2730).

2

Perform the prerequisite steps. These steps must be completed before starting to discover CLI managed NGE nodes because migration is automatic when discovery starts.

Perform the following:

1. Ensure the Global Encrypt label used on the CLI managed nodes is identical to that configured in the NFM-P.
2. Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

   Choose Key Group (NetworkGroupEncryption) from the drop-down menu and click Search to display the list of key groups.

3. Verify that both encryption and authentication algorithms of key groups on CLI managed nodes are identical to those configured on the NFM-P.
4. Disable re-keying for all key groups currently managed by the NFM-P if re-keying is enabled.
5. Choose Manage→Network Group Encryption from the NFM-P main menu and choose Cleanup Scheduled Task from the drop-down menu. Select the NGE Scheduled Task and click Shut Down to disable the task.
6. Check the list of Active Outbound Security Associations in the key group list against the CLI allocated SPIs to verify that there are no SPI conflicts on the CLI managed nodes.
7. If SPI or algorithm conflicts are found, resolve them in CLI.

3

Discover the nodes in the NFM-P; see 9.10 “Workflow for device discovery” (p. 288).

4

If needed, create a key group with the same ID as configured in the CLI; see 82.13 “To create an NGE key group” (p. 2730).

5

If needed, add the services to the key group and enable encryption:

1. Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

   Choose Key Group (NetworkGroupEncryption) from the drop-down menu and click Search to display the list of key groups.
2. Select a key group and click Properties. The Key Group (Edit) form opens.
3. Click on the Encryption tab, then the sub-tab for the type of service you need to add.
4. Click Add and use the form that opens to choose one or more objects.
5. Add any other services that are configured to the key group.
6. Click Encrypt Services to enable encryption.

6

Perform the following post-requisite steps:

1. Choose Manage→Network Group Encryption from the NFM-P main menu and choose Cleanup Scheduled Task from the drop-down menu. Select the NGE Cleanup Scheduled Task.
2. Click Turn Up to re-enable the NGE cleanup scheduled task, and Execute to run it.

   Nokia recommends running the NGE Cleanup Scheduled Task before re-keying is enabled. This will remove unused key groups and remove unused SPIs from the key group, allowing room for the new keys.
3. Enable any re-keying schedules you disabled in Stage 2.

82.12 To create the NGE global encryption label

82.12.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Click Create→Group Encryption Label. The Group Encryption Label (Create) form opens.

3

Configure the Group Encryption Label parameter.

4

Click OK. The Group Encryption Label (Create) form closes.

5

Close the Manage Network Group Encryption form.

END OF STEPS

82.13 To create an NGE key group

82.13.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Click Create→Key Group. The Key Group (Create) form opens.

3

Configure method of rekey operations: check the Force Rekey or Progressive Rekey check box, or leave both boxes unchecked to choose (default) rekey.

4

Configure the Encryption Algorithm and Authentication Algorithm parameters.

5

Click Apply. The NFM-P generates the encryption and authentication keys; the key values are displayed in the Security Associations panel.

6

Click on the Encryption tab.

7

Add one or more service objects to the key group.

1. Click on the sub-tab for the type of object that you need to add.

   Note the following:
   - The terminating site of an SDP must support NGE and be managed by the NFM-P.
   - A PW template must have the Use GRE Delivery parameter set to True.
2. Click Add and use the form that opens to choose one or more objects.

8

To apply the NGE encryption keys, perform one of the following: to one or more objects, select the objects and click Encrypt.

- To apply the NGE encryption keys to SDPs or VPRNs, click Encrypt Services.
- To apply the NGE encryption keys to other objects, select the objects and click Encrypt.

9

Click Yes to confirm the action. The NFM-P deploys the key group and keys to the participating NEs. The indicators in the Execution Status panel on the General tab display the status of the operation.

If the deployment is successful, the indicators appear as follows:

- Execution State—Encryption
- Last Execution Status—Success

10

To create a rekeying scheduled task, perform the following steps.

For maximum security, Nokia recommends that you create a rekeying scheduled task, which ensures that the keys in a key group are updated regularly.

1. Click on the Rekey Schedule tab.
2. Click Create. The Rekey Schedule, Key Group (Create) form opens.
3. Configure the parameters.
4. In the Schedule panel, click Select and use the form that opens to choose or create a schedule.

   You cannot use a schedule that has a Frequency value of Per Second or Per Minute.
5. Click OK to save your changes and close the form. A rekeying scheduled task is created.

11

Click OK to save your changes and close the Key Group (Create) form.

12

Close the Manage Network Group Encryption form.

END OF STEPS

82.14 To add an object to a key group

82.14.1 Purpose

Perform this procedure to add an NGE compatible object to an existing key group.

82.14.2 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Encryption tab.

5

Click on the sub-tab for the type of object that you need to add.

The far-end site of an SDP must support NGE and be managed by the NFM-P.

6

Click Add and use the form that opens to choose one or more objects.

7

To apply the NGE encryption keys to the objects, click Encrypt Services.

8

Click Yes to confirm the action. The NFM-P deploys the key group and keys to the participating NEs. The indicators in the Execution Status panel on the General tab display the status of the operation.

If the deployment is successful, the indicators appear as follows:

- Execution State — Enable Encryption

- Last Execution Status — Success

9

Close the Key Group (Edit) form.

10

Close the Manage Network Group Encryption form.

END OF STEPS

82.15 To create an NGE domain on a key group

82.15.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Encryption tab.

5

Click on the NGE Domains sub-tab.

6

Click Create. The NGE Domain (Create) form opens.

7

Configure the parameters in the General tab.

8

Click OK.

9

Save your changes and close the forms.

END OF STEPS

82.16 To configure an NGE domain

82.16.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Encryption tab.

5

Click on the NGE Domains sub-tab.

6

Choose an NGE domain and click Properties. The NGE Domain (Edit) form opens.

7

Add one or more managed sites to the domain:

1. Click on the Domain Sites tab.
2. Click Add→Add Managed Sites. The Select Managed Sites window opens.
3. Click Search. The available NEs are listed.
4. Select an NE and click OK.

   The site is added to the domain.

When two or more sites are added to the domain, the L3 Router Interfaces in the same subnet will be populated in the L3 Router Interfaces tab. The Ethernet ports associated with interfaces in the L3 Router Interfaces tab will be populated in the L2 Ethernet Ports tab.

You can add unmanaged sites to a domain; see 82.17 “To add unmanaged sites to an NGE domain” (p. 2737).

8

If needed, add interfaces or ports to the domain manually. When an L3 Router Interface is added, the Ethernet ports associated with the interface will be populated in the L2 Ethernet Ports tab.

To add an interface or port manually:

1. Click on the L3 Router Interfaces or L2 Ethernet Ports tab.
2. Click Add.
3. Select the interfaces or ports and click OK.

The objects are added to the domain.

9

If needed, you can bind a manually created ACL IP Exception filter to an L3 interface for L3 encryption exemption.

See 51.7 “To configure an ACL IP exception filter policy” (p. 1683) to create the policy.

To bind an ACL IP Exception Filter to a gateway L3 interface:

1. Click on the L3 Router Interfaces.
2. Select an interface and click Properties. The Key Group Routing Interface Binding (Edit) form opens.
3. Click on the Select button for the Inbound IP Exception or Outbound IP Exception fields and select a policy.
4. Save your changes and close the Key Group Routing Interface Binding (Edit) form.

10

To apply the NGE encryption keys to one or more objects:

- To encrypt interfaces:
  1. Select a domain and click Properties.
  2. Click on the L3 Router Interfaces tab.
  3. Select one or more interfaces and click Encrypt.
- To encrypt ports:
  1. Select a domain and click Properties.
  2. Click on the L2 Ethernet Ports tab.
  3. Select one or more ports and click Encrypt.
- To encrypt all ports or interfaces in a domain, select a domain and click Encrypt→L3 Router Interfaces or Encrypt→L2 Ethernet Ports.

11

To disable encryption on one or more objects:

- To disable encryption on an interface:
  1. Select a domain and click Properties.
  2. Click on the L3 Router Interfaces tab.
  3. Select one or more interfaces and click Disable Encryption.
- To disable encryption on a port:
  1. Select a domain and click Properties.
  2. Click on the L2 Ethernet Ports tab.
  3. Select one or more ports and click Disable Encryption.
- To disable encryption on all ports or interfaces in a domain, select a domain and click Disable Encryption→L3 Router Interfaces or Disable Encryption→L2 Ethernet Ports.

12

To remove a port or interface from the domain:

1. Disable encryption on the ports or interfaces you need to delete; see Step 11.
2. From the NGE Domains tab, select a domain and click Properties. The NGE Domain (Edit) form opens.
3. Click on the L3 Router Interfaces or L2 Ethernet Ports tab.
4. Select an interface or port and click Delete.
5. Click OK to close the NGE Domain (Edit) form.

13

To remove a site from the domain:

1. Disable encryption on all ports and interfaces on the site; see Step 11.
2. From the NGE Domains tab, select a domain and click Properties. The NGE Domain (Edit) form opens.
3. Click on the Domain Sites tab.
4. Select a site and click Delete.
5. Click OK to close the NGE Domain (Edit) form.

14

To delete a domain:

1. Disable encryption on the all ports and interfaces in the domain; see Step 11.
2. From the NGE Domains tab, select a domain and click Delete.

15

Close the forms.

END OF STEPS

82.17 To add unmanaged sites to an NGE domain

82.17.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Encryption tab, then on the NGE Domains sub-tab.

5

Choose an NGE domain and click Properties. The NGE Domain (Edit) form opens.

6

Configure at least one site and interface in the domain as a gateway. At least one gateway must be present before an unmanaged site can be added to the domain.

To configure a gateway:

1. Click on the Gateways tab, then on the Sites sub-tab.
2. Click Add and select a site.
3. Click OK to add the site to the list of gateway sites.
4. Click on the Interfaces sub-tab.
5. Click Add and select an L3 interface.
6. Click OK to add the interface to the list.

7

Add one or more unmanaged sites to the domain.

To add an unmanaged site:

1. Click on the Domain Sites sub-tab.
2. Click Add→Add Unmanaged Sites. The NGE Domain Site (Create) form opens.
3. Configure the parameters.

   If you need to discover a VSR which is managed by a VSR-a, configure the Managed by VSR-a check box and enter the VSR-a IP address. The VSR will be discovered as part of discovering the VSR-a.
4. Click OK.

The unmanaged site is added to the list of sites in the domain.

Repeat this step as required to add other sites.

8

Verify that the Unmanaged Site check box is enabled for all unmanaged sites.

Select the unmanaged sites and click NGE Discovery. The IP addresses of the unmanaged sites are added to the discovery rule specified in Step 7.

When NGE discovery is triggered, the NFM-P inserts the IP address of the unmanaged NE to the specified discovery rule. ACL IP exception filters will then be created on the gateway interface, in preparation for successful encryption on the interface of the newly discovered NE.

9

Wait for the NGE discovery process to complete an NE resynchronization. When the resynchronization is complete, the discovered sites become managed sites in the NGE domain and the Unmanaged Site check box is disabled.

If the NGE Discovery Execution Status is Failed, check the NGE Discovery Execution State for the failure reason.

10

Enable encryption on interfaces of the newly discovered NE by following Step 10 of 82.16 “To configure an NGE domain” (p. 2734).

After encryption has been enabled, the ACL IP exception filters are removed by the NFM-P.

11

Close the forms.

END OF STEPS

82.18 To manually execute a rekeying scheduled task

82.18.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Rekey Schedule tab.

5

Click Execute→Rekey, Execute→Force Rekey or Execute→Progressive Rekey. The NFM-P performs the rekeying operation.

6

If required, view the rekeying results, as described in 82.19 “To view rekeying results and statistics” (p. 2739).

7

Close the Key Group (Edit) form.

8

Close the Manage Network Group Encryption form.

END OF STEPS

82.19 To view rekeying results and statistics

82.19.1 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Rekey Schedule tab.

5

To view results of a completed rekey scheduled task:

1. Click on the Results tab. The scheduled task results for all runs are listed.
2. To view a result, select the result and click Properties. The Rekey Schedule Result form opens.
3. View the information, as required.
4. Close the Rekey Schedule Result form.

6

To view results of an ongoing rekey operation:

1. From the Rekey Schedule tab, select a scheduled rekey task.
2. Click Rekey Progress.

   A Rekey Progress Info dialog box appears, showing the total number of sites to be rekeyed, and remaining sites to be rekeyed.
3. If necessary, you can intervene with the rekey process:
   - Click Force Rekey to switch to a force rekey process. The rekey will proceed as a forced rekey.
   - Click Stop/Fail to manually end the rekey process.
4. Click OK to close the dialog box.

7

To view statistics about the rekeying task runs, perform the following steps:

1. Click on the Rekey Times tab.
2. View the information, which describes the minimum, maximum, and average durations of new key set deployment and previous key set removal.

8

Close the Key Group (Edit) form.

9

Close the Manage Network Group Encryption form.

END OF STEPS

82.20 To disable encryption on an object

82.20.1 Purpose

Perform this procedure to remove the NGE encryption from an encrypted object.

82.20.2 Steps

1

Choose Manage→Network Group Encryption from the NFM-P main menu. The Manage Network Group Encryption form opens.

2

Choose Key Group from the drop-down menu and click Search. The NFM-P NGE key groups are listed.

3

Select a key group and click Properties. The Key Group (Edit) form opens.

4

Click on the Encryption tab.

5

Click on the sub-tab for the type of object on which you need to disable encryption.

6

Select one or more objects and click Disable Encryption. You can disable encryption only on objects with encryption enabled.

7

Click Yes to confirm the action. The NFM-P removes the service association between the key group and the object. The indicators in the Execution Status panel on the General tab display the status of the operation.

If the operation is successful, the indicators appear as follows:

- Execution State—Disable Encryption
- Last Execution Status—Success

8

Close the Key Group (Edit) form.

9

Close the Manage Network Group Encryption form.

END OF STEPS

83 Service PW template policies

83.1 Service PW template policies

83.1.1 Overview

Service PW templates policies enable you to automatically create SDP bindings on NEs that support BGP AD VPLS. PW templates contain the configuration to use when the NFM-P creates SDP bindings automatically through BGP AD. PW templates use the NFM-P policy distribution model to distribute the template to the NEs.

To prevent unwanted service disruptions, NEs do not automatically reconfigure SDP bindings when a change is made to the PW template from which they were generated. You can manually trigger a template re-evaluation to cause the NEs to re-configure SDPs to match the configuration in the template.

PW templates can be re-evaluated at three levels:

- service
- service site
- local definition of PW template policy

When you re-evaluate a PW template from the local definition of the PW template policy, the NFM-P re-evaluates the selected template for all of the VPLSs on the selected site.

When you re-evaluate a PW template from the service site, the NFM-P re-evaluates the selected template on the selected service site.

When you re-evaluate a PW template from the service, the NFM-P re-evaluates the selected template on all sites in the VPLS.

83.2 Workflow to configure and manage PW template policies

83.2.1 Purpose

The following workflow lists the high-level steps required to create and manage a service PW template policy.

83.2.2 Stages

1

Create a PW template policy. See 83.3 “To configure a PW template policy” (p. 2744) .

2

Release and distribute the PW template policy to each NE that is or will be a component of the VPLS or BGP VPLS. See 83.4 “To distribute a PW template policy” (p. 2746) .

3

Reevaluate PW template when a configuration change is made to the PW template. See 83.5 “To reevaluate a PW template policy on a local definition after a configuration change” (p. 2747) .

83.3 To configure a PW template policy

Note: Starting with NE release 15.1 R1, the Name parameter can only be configured during creation of the policy. It cannot be changed after the policy is created. However, the name can be changed if the NE release is 15.0, R4 or later.

83.3.1 Steps

1

Choose Manage→Service→Service PW Template Policies from the NFM-P main menu. The Manage Service PW Template Policies form opens.

2

Click Create. The Service PW Template (Create) form opens.

3

Perform one of the following:

a. To populate the Service PW Template form using an existing pseudowire (SDP Binding) as a base:

   1. Click Populate from Pseudowire and select a pseudowire. The Service PW Template (Create) refreshes with the pseudowire information.
   2. Go to Step 18 .

b. To manually configure the service PW template:

   1. Configure the required general parameters.
   2. Configure the required parameters in the Miscellaneous panel.

      The Block On Peer Fault parameter takes effect only when PW status signaling is enabled.

      The Block On Peer Fault parameter is not configurable on a spoke SDP that is on an MC LAG or an endpoint.

      You can set either the Enable Force-Vlan-VC Forwarding parameter or the Enable Force Q-in-Q-VC Forwarding parameter to true, but not both.
   3. Select an Accounting Policy in the Accounting Policy panel.

4

Click on the MAC tab and configure the required parameters.

5

Click on the Split Horizon Group tab and configure the required parameters.

6

Click on the Ingress Filters tab and configure the Ingress Filter Type parameter.

7

Click on the Egress Filters tab and configure the Egress Filter Type parameter.

8

Click on the IGMP tab and configure the required parameters.

9

Click on the STP tab and configure the required parameters.

10

Click on the QoS tab.

11

Configure the required parameters in the Ingress Forwarding Plane Redirect panel:

1. Select a Network Policy to assign a policy to the PW template.
2. Select an Ingress Queue Group Template Policy to assign a policy to the PW template.
3. Configure the Instance ID parameter.

12

Click on the Tunnel Admin Groups tab.

Note: You can apply the inclusion or exclusion of tunnel administrative groups only if you set the Use Provisioned SDP parameter to true in Step 3 b 1 .

13

Click on the tunnel administrative groups that you need to include in the PW template policy in the Included Admin Groups-Unassigned panel. You can choose multiple groups by holding down the Ctrl key when you click on the group.

14

Click on the right arrow to move the selected administrative groups into the Included Admin Groups-Assigned panel. You can move an assigned administrative group to the unassigned portion of the form if required, by selecting the group and clicking on the left arrow.

15

Click on the tunnel administrative groups that you need to exclude from the PW template policy in the Excluded Admin Groups-Unassigned panel. You can choose multiple groups by holding down the Ctrl key when you click on the group.

16

Click the right arrow to move the selected administrative groups into the Excluded Admin Groups-Assigned panel. You can move an assigned tunnel administrative group to the unassigned portion of the form, if required, by selecting the group and clicking the left arrow

17

Configure the required parameters in the Egress Port Redirect panel:

1. Select a Network Policy Select to assign a policy to the PW template.
2. Select an Egress Queue Group Template Policy to assign a policy to the PW template.
3. Configure the Instance ID parameter.

18

Save your changes and close the forms.

END OF STEPS

83.4 To distribute a PW template policy

83.4.1 Purpose

Use this procedure to manually distribute PW template policies that are used by BGP AD capable NEs and to configure the distribution mode of local policies.

When you distribute a global policy, local policies using the Sync With Global distribution mode allow the NE to receive the policy.

Note: Local policies using the Local Edit Only distribution mode do not allow the NE to receive the distribution of a global policy. You must ensure that the policy distribution mode for the local policy is set to Sync With Global if you need the NE to receive the distribution of a global policy.

83.4.2 Steps

1

Choose Manage→Service→Service PW Template Policies from the NFM-P main menu. The Manage Service PW Template Policies form opens.

2

Select the policy or policies that you need to distribute.

3

Perform 49.6 “To release and distribute a policy” (p. 1476) to release and distribute the Service PW Template Policy.

END OF STEPS

83.5 To reevaluate a PW template policy on a local definition after a configuration change

83.5.1 Purpose

Perform the following procedure to cause the NEs to re-configure SDPs to match the configuration in the template when a PW template policy is modified. This causes a re-evaluation of the selected template for all VPLS services on the selected NE.

You can also reevaluate a PW template policy at the service and service site levels.

83.5.2 Steps

1

Choose Manage→Service→Service PW Template Policies from the NFM-P main menu. The Manage Service PW Template Policies form opens.

2

Choose a PW template policy and click Properties. The Service PW Template Policy (Edit) form opens.

3

Click on the Local Definitions tab.

4

Choose the NE on which you need to re-evaluate the PW template policy. The Service PW Template - Local Policy (Edit) form opens.

5

Click Re-evaluate PW Template to run an evaluation of the modified template. The NFM-P performs a re-evaluation of the selected template for all VPLS services on the selected NE.

6

A dialog box appears indicating if the re-evaluation was successful. If it was not, the reason for the failure is displayed. If you make any subsequent modifications, you can re-evaluate the template again. Click OK.

7

Save the changes and close the forms.

END OF STEPS

84 Service SAP template policies

84.1 Service SAP template policies

84.1.1 Overview

The epipe SAP template policy is used in Wi-Fi offload VLAN to anchor ISA configurations. It allows global ingress and egress qos policy configuration and traffic filtering for access SAPs, allowing users to shape traffic based on the Wi-Fi radio technology of an AP. For more information, see 39.1.7 “VLAN access to anchor ISA” (p. 1321) .

The VPLS SAP template policy is used in L2 wholesale-retail configurations. It provides global ingress/egress policy and other information for the creation of internal VPLS SAPs on a WLAN GW. For more information, see 39.1.8 “L2 wholesale” (p. 1321) .

84.2 To configure an epipe SAP template policy

84.2.1 Steps

1

Choose Manage→Service→SAP Template Policies from the NFM-P main menu. The Manage Service SAP Template Policies form opens.

2

Click Create→Epipe SAP Template or select an existing epipe SAP template policy and click Properties. The Epipe SAP Template (Create|Edit) form opens.

3

On the General tab, configure the Displayed Name parameter.

An epipe SAP template policy cannot share the same object name as a VPLS SAP template policy if both policies are distributed to the same NE. Although both policies could co-exist in the NFM-P database, the NE stores both policy types in the same MIB table, which would cause a conflict.

4

Click on the Egress tab and select the following objects for the epipe SAP template policy:

- IP filter
- IPv6 filter
- MAC filter
- QoS policy

5

Click on the Ingress tab.

1. Select the following objects for the epipe SAP template policy:
   - IP filter
   - IPv6 filter
   - MAC filter
   - QoS policy
2. Configure the Shared Queue parameters, if required.

6

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

84.3 To configure a VPLS SAP template policy

84.3.1 Steps

1

Choose Manage→Service→SAP Template Policies from the NFM-P main menu. The Manage Service SAP Template Policies form opens.

2

Click Create→VPLS SAP Template or select an existing VPLS SAP template policy and click Properties. The VPLS SAP Template (Create|Edit) form opens.

3

On the General tab, configure the required parameters.

An epipe SAP template policy cannot share the same object name as a VPLS SAP template policy if both policies are distributed to the same NE. Although both policies could co-exist in the NFM-P database, the NE stores both policy types in the same MIB table, which would cause a conflict.

4

Select a DoS protection policy, if required.

If you configure a DoS protection policy on the VPLS SAP template policy, you can configure the CPM Monitor MAC parameter or select ETH CFM Monitoring options.

5

Click on the Ingress tab to configure ingress filter, QoS, and scheduler policies for the VPLS SAP template policy.

1. Select the following policy types, as required:
   - IP filter
   - IPv6 filter
   - MAC filter
   - QoS
   - QoS scheduler
   - Policer control
2. Configure the remaining parameters.

6

Click on the Egress tab to configure egress filter, QoS, and scheduler policies for the VPLS SAP template policy.

1. Select the following policy types, as required:
   - IP filter
   - IPv6 filter
   - MAC filter
   - QoS
   - QoS scheduler
   - Policer control
2. Configure the remaining parameters.

7

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

85 Composite service management

85.1 Overview

85.1.1 Purpose

This chapter provides overview and procedural information on composite service management in the NFM-P.

85.1.2 Contents

Composite service management

85.2 Overview

85.2.1 General information

A composite service is a set of linked services. Composite service functionality supports complex applications that require a combination of services, such as VLAN connections to an HVPLS, an IES spoke into a VPLS, or a VPRN-to-VPLS interconnection.

Services that are owned by different customers can be connected to form a composite service. An example is an HVPLS in which the core VPLS belongs to one customer and the satellite VPLS instances belong to other customers. An HVPLS is considered to be a composite service by the NFM-P.

Composite services consist of customer services, called SCs in the context of a composite service, and connectors. A connector is a bidirectional logical link between two SCs, such as a pair of PW spokes that carry traffic in opposite directions between VLL and VPLS instances, a dot1q-encapsulated link between a VLAN and a VPLS, or an internal cross-connect.

The term SCP describes a type of connector endpoint. In the case of the services that are available on the 7210 SAS, 7450 ESS, 7750 SR, or the 7950 XRS, an SCP is a service interface or SAP. For L2 switches, an SCP may be a network interface, such as an uplink port.

Composite services exist only in the context of the NFM-P and are configured through the NFM-P GUI or an OSS application. They are unknown to individual network devices. To simplify composite service configuration and to ensure that non-NFM-P device configuration does not disrupt the management of composite services, the following rules apply to the creation, deletion, modification, and presentation of composite services.

- A composite service can have zero SCs.
- A composite service can have zero connectors.
- Two connected SCs can belong to only one composite service.
- A connector between two SCs belongs to only one composite service.
- An SC cannot be removed from a composite service until its connector to the composite service is removed.
- A group of connected services can be moved from one composite service to another.

The NFM-P supports composite-service configuration using the following methods.

- Tabbed configuration forms with an embedded navigation tree that provides a logical, hierarchical view of the composite service and acts as a configuration interface. When you right-click on an object In the service navigation tree, a menu specific to the object appears. When you choose an object in the menu, the related configuration form opens.
- Connector creation and configuration from within the composite service’s flat topology view

If a service that is specified for inclusion in a composite service does not currently belong to a composite service, it is added to the composite service regardless of its administrative or operational state. If the specified service is part of an existing composite service, it can be moved to a different composite service. However, SCs that are connected to the specified service are also

moved to the new composite service upon confirmation of the action by the NFM-P operator. The NFM-P performs no such action confirmation for OSS applications.

Services within a composite service can have the same service ID, service type, and customer ID. For example, a VPLS with service ID 5 on one NE and an IES service with service ID 5 on another NE can be combined to create a composite service.

85.2.2 Hierarchical organization of composite services

The NFM-P organizes the SCs in a composite service by tiers within a hierarchy during network discovery and for display purposes. When you add an SC to a composite service, the NFM-P assigns a default tier value to the SC according to the service type.

Table 85-1 Default tiers for service types

Service type	Default tier
IES and VPRN	1
VPLS and MVPLS	2
VLL	3
VLAN	4

The default tier values reflect common deployment configurations. The figure below shows a composite service hierarchy. SCs are not restricted to the default tiers, as in the case of the VPLS in Tier 3 that forms an HVPLS with the VPLS in Tier 2.

Figure 85-1 Default SC tier assignments

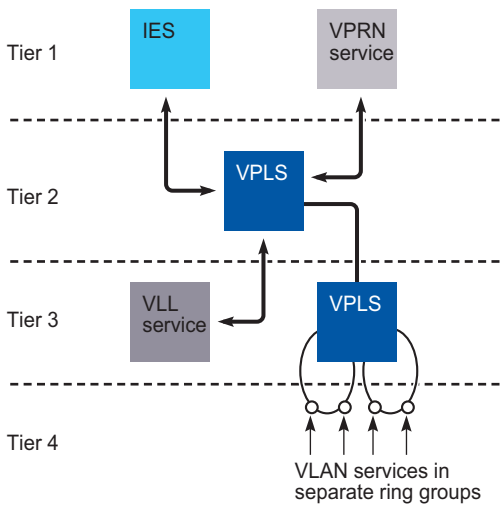

You can change the tier of an SC at any time and can specify a value other than one of the defaults that the NFM-P assigns; composite services can have many tiers.

The tier value of an SC determines the relative position of the SC within a composite service topology map. The NFM-P displays SCs in rows by tier, in numerical order from the top of the panel downward, beginning with tier 1.

You can move SC icons in a composite service topology map from one tier to another for a customized view, then use the Rearrange by Service Tiers button to organize the SC icons in the map panel according to tier. The composite service topology map is redrawn when you click on the Rearrange by Service Tiers button, and you cannot revert to the former layout of the composite service topology map.

See Chapter 4, “Topology map management” for more information about NFM-P map management.

85.2.3 Network discovery of composite services

The NFM-P associates SCs and connectors with composite services during network discovery. The following rules apply to this process.

- When the NFM-P discovers a valid connector between two services that do not belong to a composite service, a composite service that contains the services and connector is automatically created.
- When the NFM-P discovers a valid connector between two services and one of the services belongs to a composite service, the other service is added as an SC of the composite service.
- When the NFM-P discovers a valid connector between two services and the two services are SCs of different composite services, an alarm is raised and the connector is excluded from the NFM-P database.

The NFM-P assigns a default tier value to a service upon creation and to a new SC during network discovery. The tiered hierarchy provides a common framework for service configurations that are provisioned through the NFM-P and CLI. An NFM-P operator can assign a different tier value to an SC after discovery.

A composite service has Aggregated Operational State and Service Component Degraded status indicators. The General tab of the Composite Service management form displays these indicators.

The Aggregated Operational State indicator has four possible values: Up, Down, Partially Down, and Unknown. The value is derived from the aggregated SC operational states as follows.

- Up—All SCs are operationally up
- Partially Down—At least one SC is operationally down
- Down—All SCs are operationally down
- Unknown—The status of at least one SC is undetermined

The Service Component Degraded indicator shows whether there is an operational flag set on any of the service sites under this composite service.

85.3 Connector types

85.3.1 General information

The following types of connectors join SCs in a composite service:

- SCP-to-SCP
- internal cross-connect
- PW spoke
- routed VPLS
- VRF RT connections for VPRN services

Note: The NFM-P does not create connectors between mesh SDP bindings.

85.3.2 SCP-to-SCP connectors

SCP-to-SCP connectors can join any two SC types that have service interfaces on the same device or on different devices. A connector between VPLS and VPRN SAPs is an SCP-to-SCP connector, as is a connector between a dot1q-encapsulated VPLS SAP and L2 switch uplink port in a VLAN ring group. The table below describes the supported encapsulation types.

Table 85-2 Supported encapsulation types

SAP type	Encapsulation type
Ethernet[1]	dot1q
	QinQ
	Null
ATM	VPI/VCI
	VPI
FR	DLCI
SONET/SDH	BCP Null
	BCP dot1q
	IPCP
	PPP Auto
	cHDLC
	WAN Mirror
LAG	Null
	dot1q
	QinQ
	Null

Notes:

1. VLAN connection profiles are supported for Ethernet SAPs with dot1q and QinQ.

The operational status of an SCP-to-SCP connector depends on the operational status of its endpoints. An alarm raised against one of the endpoints causes an alarm to be raised against the connector. Such alarms are aggregated within the composite service.

85.3.3 Internal cross-connect connectors

An internal cross-connect connector can join any SC types. It uses a CCAG to join two SCs that have SAPs or network interfaces on the same device. This functionality is available in the 7450 ESS and 7750 SR. The following rules apply to internal cross-connect connectors.

- A SAP can be connected to another SAP or to a network interface using a CCAG.
- When a SAP or network interface is deleted, the connector associated with it is also deleted.
- The deletion of an internal cross-connect connector causes the associated interfaces and SAPs to be deleted.

The operational state of an internal cross-connect connector depends on the operational state of the CCAG. An alarm raised against the CCAG causes an alarm to be raised against the connector. Such alarms are aggregated within the composite service.

85.3.4 PW spoke connectors

A PW spoke connector generally joins VPLS instances to create an HVPLS. In the 7210 SAS, 7450 ESS, 7750 SR, and 7950 XRS, a PW spoke can, for example, connect IES and VPLS instances to provide distributed Internet access service. The endpoints of a PW spoke connector must be on different devices. PW spoke connectors are subject to restrictions on the SC types that they can join. The table below lists the SC types that can be linked by PW spoke connectors.

Table 85-3 Valid PW spoke interconnections

SC type	Valid PW spoke SC interconnections
VLL	IES, VPLS
VLAN	—
VPLS	IES, VLL, VPLS
MVPLS	MVPLS
IES	VLL, VPLS
VPRN	VLL, VPLS

The operational state of a PW spoke connector depends on the operational state of the underlying SDP bindings. An alarm raised against one of the SDP bindings causes an alarm to be raised against the connector. Such alarms are aggregated within the composite service.

85.3.5 Routed VPLS connectors

A routed VPLS connector joins an L3 access interface within an IES or VPRN service context to a VPLS or I-VPLS on the same site. When an IES or VPRN IP interface is bound to a VPLS site

name, the site name cannot be bound to another IP interface. Although an IES or VPRN IP interface can only be bound to a single VPLS site, the service context that contains the IP interface can have other IP interfaces bound to other VPLS sites. Both the IES or VPRN IP interface and VPLS site must be located on the same NE.

If a VPLS site name does not exist within the system, the binding between the IP interface and the VPLS site remains operationally down until a VPLS site name is assigned to the VPLS site. When an IP interface is bound to a VPLS site, the operational state of the binding depends on the operational state of the VPLS site, or whether the IP interface binding is enabled on the VPLS site.

The operational state of the routed VPLS connector depends on the operational state of the binding and the operational state of the L3 IP interface.

The routed VPLS connector function is supported for IOM3, or later, cards in the following devices:

- 7450 ESS in mixed mode
- 7750 SR
- 7750 SR-c4
- 7705 SAR, all variants
- 7950 XRS

85.3.6 VPRN RT connections

RT connectivity is the matching of import and export VRF targets for VPRN services, and is represented as a dotted line between the services with arrows indicating the direction. You can highlight the VRF targets between VPRN services within a composite service on the topology map and on the flat topology map. The flat topology map shows RT connections between all of the sites, whereas the topology map show RT details between services.

For VPRN RT connections within composite services, you can define:

- the VRF import and export targets on the sites (only one import and one export RT value on each site)
- the VRF target based on import and export policies, where multiple community members can be defined

Consider the following when you highlight the VRF targets between VPRN services within a composite service on the topology map:

- the RT connections are highlighted when the RT defined in the VRF import policy matches the VRF export policy on another site
- if a site is configured with VRF target values and with VRF import and export policies, the NFM-P uses the VRF import and export policy

See 79.27 “To configure a VRF instance on a VPRN site” (p. 2567) for more information about how to configure VRF instances on VPRN sites.

See 79.50 “To configure VRF import and export policies on a VPRN site” (p. 2615) for more information about how to configure VRF import and export policies on a VPRN site.

See Chapter 54, “Routing policies” for more information about routing policies.

85.4 Sample composite service configuration

85.4.1 General information

The figure below shows a sample composite service configuration that involves a variety of customer services and uses the three SC connector types.

Figure 85-2 Sample composite service configuration

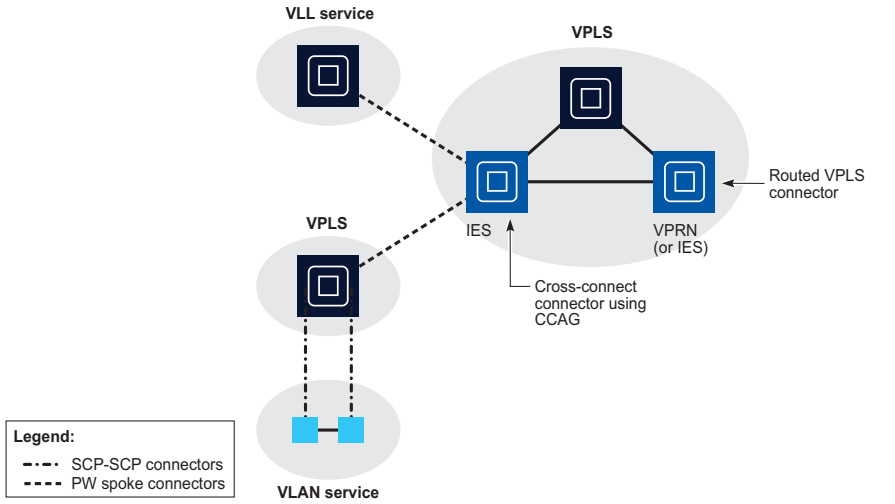

85.4.2 Workflow to create a composite service

The following workflow lists the high-level steps required to create a composite service. As a prerequisite for creating a composite service, this workflow assumes the following:

- a group or customer with the required user access privileges has been configured.
- the IP or IP/MPLS core network exists.
- any required service tunnels are created including the static, dynamic or SR-TE LSP required to create the service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) for more information.
- the access ports for the service are created; see Chapter 16, “Port and channel object configuration” for more information.
- any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite policies are created; see Chapter 49, “Policies overview” for more information. You do not have to create pre-defined policies if policies are created on a per-service basis.
- any required MP-BGP for PE-to-PE routing is configured; see Chapter 28, “Routing protocol configuration” for more information about protocol configuration.
- the other network services that are to be the SCs of the composite service are created.

85.4.3 Stages

1

As required, modify the default system preferences for composite services such as specifying if composite services are auto-discovered or if service alarms are aggregated for composite services.

2

Create composite service.

1. Define the general properties for the composite service. See 85.5 “To create a composite service” (p. 2764) for more information.
2. Specify the services for inclusion in the composite service. You can specify multiple services in one operation. See 85.6 “To add services to a composite service” (p. 2764) for more information.
3. Create connectors to link the SCs of the composite service. See 85.7 “To create a cross connect connector” (p. 2765) , 85.8 “To create an SCP connector” (p. 2766) , 85.10 “To discover VRF route target connections” (p. 2768) , and 85.12 “To create a routed VPLS connector” (p. 2770) for more information.
4. Turn up the composite service.

3

As required, run an OAM validation test for the composite service. See 85.13 “To run an OAM validation test for a composite service” (p. 2771) for more information. Alternatively, you can

also run a One Time Validation on the composite service. See 90.55 “To run a one-time validation test on a service” (p. 3084) for more information.

4

As required, force a manual rediscovery of all composite services in the NFM-P to discover new service connectors that would add SCs to an existing composite service, or for new service connectors that would connect together two SCs to create a new composite service. See 85.14 “To rediscover composite services” (p. 2772) for more information.

5

Manage a composite service.

1. View the service topology map associated with a composite service. See 85.15 “To view the service topology map associated with a composite service” (p. 2773) for more information.
2. As required, modify a composite service:
   - using the navigation tree; see 85.16 “To modify a composite service using the navigation tree” (p. 2773) for more information
   - using the flat topology view; see 85.17 “To modify a composite service using the flat topology view” (p. 2774) for more information
3. As required, delete a composite service. See 85.18 “To delete a composite service” (p. 2778) for more information.

Composite service management procedures

85.5 To create a composite service

85.5.1 Steps

1

Choose Create→Service→Composite Service from the NFM-P main menu. The Composite Service (Create) form opens.

2

Configure the required general parameters.

The Composite ID parameter is configurable only if the Auto-Assign ID parameter is disabled.

3

Click Apply.

4

Save the changes and close the forms.

END OF STEPS

85.6 To add services to a composite service

85.6.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

In the service navigation tree, right-click on the Services icon and choose Add Services. The Add Services form opens.

Note: You can also create a service for inclusion in the composite service by right-clicking on the Services icon and choosing Create *Service Type*.

4

Select one or more services.

5

To add a connector to the composite service:

a. Create a cross connect connector. Perform Step 3 to Step 12 of 85.7 “To create a cross connect connector” (p. 2764) .

b. Create a routed VPLS connector. Perform Step 3 to Step 11 of 85.12 “To create a routed VPLS connector” (p. 2770) .

c. Create an SCP connector. Perform Step 3 to Step 12 of 85.8 “To create an SCP connector” (p. 2766) .

d. Create a Spoke Connector. Perform Step 3 to Step 14 of 85.9 “To create a spoke connector” (p. 2767) .

6

Save the changes and close the forms.

END OF STEPS

85.7 To create a cross connect connector

85.7.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

In the service navigation tree, right-click on the Connectors icon and choose Create CrossConnect. The Cross Connect (Create) form opens.

Note: The service endpoints of a cross-connect connector must be on the same NE.

4

Configure the required general parameters.

5

Select an SC to associate with the connector in the Service A panel.

6

Select an SC to associate with the connector in the Service B panel.

7

Select a site to associate with the connector in the Site A panel.

8

Select a site to associate with the connector in the Site B panel.

9

Click on the Transport tab.

10

Select a CCAG for the connector.

11

Configure the required parameters.

12

Save the changes and close the forms.

END OF STEPS

85.8 To create an SCP connector

85.8.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

In the service navigation tree, right-click on the Connectors icon and choose Create ScpConnector. The ScpConnector (Create) form opens.

4

Configure the required general parameters.

5

Select an SC to associate with the connector in the Service A panel.

6

Select an SC to associate with the connector in the Service B panel.

7

Select a site to associate with the connector in the Site A panel.

8

Select a site to associate with the connector in the Site B panel.

9

Click on the Service Connection Point tab.

10

Select an SCP to associate with the connector in the Service Connector Point A panel.

11

Select an SCP to associate with the connector in the Service Connector Point B panel.

12

Save the changes and close the forms.

END OF STEPS

85.9 To create a spoke connector

85.9.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

In the service navigation tree, right-click on the Connectors icon and choose Create ScpConnector. The SpokeConnector (Create) form opens.

Note: The service endpoints of a spoke connector must be on different NEs.

4

Configure the required general parameters.

5

Select an SC to associate with the connector in the Service A panel.

6

Select an SC to associate with the connector in the Service B panel.

7

Select a site from service A to associate with the connector in the Site A panel.

8

Select a site from service B to associate with the connector in the Site B panel.

9

Click on the Transport tab.

10

Configure the parameters.

The Transport Type parameter is configurable when the Auto Select Tunnels parameter is enabled.

11

Select a service tunnel to associate with the connector in the Tunnel A panel

12

Select a service tunnel to associate with the connector in the Tunnel B panel

13

If one of the SCs in the composite service is an IES, the First L3 Access Interface and Second L3 Access Interface panels are present. Select an access interface to associate with the spoke connector in the First L3 Access Interface and Second L3 Access Interface panels. Otherwise, go to Step 4 .

14

Save the changes and close the forms.

END OF STEPS

85.10 To discover VRF route target connections

85.10.1 Steps

1

Choose Administration→System Preferences and configure the Enable VRF Route Target

Connections parameter in the Composite Services panel on the Service tab. See the *NSP System Administrator Guide* for more information.

2

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

3

Click Rediscover Composite Services→VRF Route Target to update all of the VPRN Route Target connectors.

Any unidirectional RT connectivity between two different VPRN services are automatically moved into one composite service, if any route targets are highlighted.

4

Close the forms.

5

Perform 85.11 “To draw VRF targets between VPRN services” (p. 2768) to highlight VRF targets between VPRN services.

Note: If you disable the Enable VRF Route Target Connections parameter in the system preferences, the NFM-P removes all discovered composite services based on the VRF Route Target. You must repeat 85.10 “To discover VRF route target connections” (p. 2768) to re-discover the VRF route targets.

END OF STEPS

85.11 To draw VRF targets between VPRN services

85.11.1 Purpose

Use this procedure to draw VRF targets between VPRN services on a topology map associated with a composite service.

85.11.2 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and perform one of the following steps:

a. Click on the Topology View button. A Composite Service Topology - *Service Name* map opens.

b. Click on the Topology View Flat button. A Composite Service Flat Topology - *Service Name* map opens. The Flat Topology map is a flattened view of a composite service, meaning that all the service objects are displayed simultaneously, along with the service sites, access interfaces, and the links or groups of links between them.

See Chapter 4, “Topology map management” for more information about composite service topology views.

3

Right-click on the composite services topology map composite services flat topology map and choose Highlight Route Target Topology.

4

Close the forms.

END OF STEPS

85.12 To create a routed VPLS connector

85.12.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

In the service navigation tree, right-click on Connectors and choose Create RoutedVplsConnector. The RoutedVplsConnector (Create) form opens.

4

Configure the required general parameters.

5

Select an SC to associate with the connector in the Service A panel.

6

Select an SC to associate with the connector in the Service B panel.

7

Select a site from service A to associate with the connector in the Site A panel.

8

Select a site from service B to associate with the connector in the Site B panel.

Note: The sites that you select in Step 7 and Step 8 must be located on the same NE.

9

Click on the Routed-VPLS L3 Connection Point tab.

10

Select an L3 access interface.

11

Save the changes and close the forms.

END OF STEPS

85.13 To run an OAM validation test for a composite service

85.13.1 Purpose

A validator test suite must be created for the tested entity. See Chapter 89, “Service Test Manager” for more information about how to create a validator test suite.

Note: As an alternative, you can also run an OAM validation test on the composite service by performing a One Time Validation. This is a mostly automated procedure and is described in 90.55 “To run a one-time validation test on a service” (p. 3084) .

85.13.2 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service (Edit) form opens.

3

Click Action and choose Validate. If a validator test suite is not associated to the composite service, a dialog box appears. Perform the following steps:

1. Click OK to associate the composite service with an existing validator test suite. The Choose Validator Test Suite form opens.
2. Select a validator test suite and click OK The Choose Validator Test Suite form closes.

4

View the Operational Flags indicators. If the validation test fails, a check mark appears beside the OAM Validation Failed indicator.

5

Click on the Tests tab.

6

Click on the Tested Entity Result tab.

7

Select an entry and click Properties. The Tested Entity Result form opens and displays information about the validation test.

8

Close the forms.

END OF STEPS

85.14 To rediscover composite services

85.14.1 Purpose

This procedure forces a manual rediscovery of all composite services in the NFM-P. The Rediscover Composite Services command performs a network discovery that looks for new service connectors that would add SCs to an existing composite service, or for new service connectors that would connect together two SCs to create a new composite service.

Note: If the Auto Discover Composite Services system preference is enabled, composite services are rediscovered as part of routine network discovery in the NFM-P. See 73.5 “To list the services associated with a customer” (p. 1982) for more information.

85.14.2 Steps

1

Choose Manage→Service→Composite from the NFM-P main menu. The Manage Composite Services form opens.

2

Click Rediscover Composite Services.

3

A warning dialog box appears. Click Yes to start the rediscovery.

4

A progress bar appears, indicating the status of the rediscovery process. Upon completion, an information form appears, indicating the number of composite services and connectors discovered. Close the form.

5

Close the Manage Composite Services form.

END OF STEPS

85.15 To view the service topology map associated with a composite service

85.15.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and perform one of the following steps:

a. Click on the Topology View button. A Composite Service Topology - *Service Name* map opens.

b. Click on the Topology View Flat button. A Composite Service Flat Topology - *Service Name* map opens. The Flat Topology map is a flattened view of a composite service, meaning that all the service objects are displayed simultaneously, along with the service sites, access interfaces, and the links or groups of links between them.

See Chapter 4, “Topology map management” for more information about composite service topology views.

3

Close the forms.

END OF STEPS

85.16 To modify a composite service using the navigation tree

Note: Modifying parameters can be service-affecting.

Note: You can also modify composite service components and add service connectors from the flat topology view. See 85.17 “To modify a composite service using the flat topology view” (p. 2774) .

85.16.1 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The Composite Service - *Service Name* (Edit) form opens.

3

In the service navigation tree, right-click on an item and choose Properties.

Using the contextual menu, you can also:

- Add SCs and connectors to a composite service
- Create services for inclusion in the composite service. When you create services by using the contextual menu, the services are automatically added to the composite service.
- Remove SCs and connectors from a composite service
- Move SCs to another composite service
- Delete SCs

  Deleting an SC is not the same as removing an SC from a composite service. Deleting an SC removes the service from the NFM-P database. To avoid a service outage, be certain of the action that you are taking.

4

Save the changes and close the forms.

END OF STEPS

85.17 To modify a composite service using the flat topology view

85.17.1 Purpose

You can modify composite service components and add service connectors from the flat topology view. The main advantage of creating connectors using the flat topology view is that most of the parameters in the associated configuration forms are automatically populated when you select an item in the map.

Note: Modifying parameters can be service-affecting.

Note: You can also modify and add composite service components using the navigation tree. See 85.16 “To modify a composite service using the navigation tree” (p. 2773) .

85.17.2 Steps

1

Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.

2

Select a composite service and click Properties. The CompositeService - *Service Name* (Edit) form opens.

3

Click Flat Topology View. The Composite Service Flat Topology map opens.

4

To configure parameters for any item on the map, right-click on the item and choose Properties. The configuration form for the item is displayed.

5

Edit the parameters as required. See 85.5 “To create a composite service” (p. 2764) for detailed configuration information on creating and configuring components.

6

Add service connectors to the composite service, if required.

Depending on the sites or ports you select, when you right-click on a component, the contextual menu contains one or more of the following choices:

- Create Cross Connect. Go to Step 7 .
- Create Routed Vpls Connector. Go to Step 8 .
- Create Scp Connector. Go to Step 9 .
- Create Spoke Connector. Go to Step 10 .

Click on the required item in the contextual menu and proceed to the associated step indicated in this list.

Note: If you need to modify any of the automatically-populated parameters in the associated configuration forms while creating these connectors, see 85.16 “To modify a composite service using the navigation tree” (p. 2773) for detailed information.

7

To create a Cross Connect Connector:

1. Select the two sites for which you want to create the Cross Connect Connector. These represent the same physical NE that exists in two separate services, since service endpoints of a Cross Connect Connector must be on the same NE.

2. Right-click on either of the NEs and choose Create Cross Connect.

   The CrossConnect (Create) form opens.
3. Configure the parameters.
4. Click on the Transport tab.
5. Select a CCAG for the connector.
6. Save the changes and close the form.
7. Go to Step 14 .

8

To create a Routed VPLS Connector:

1. Select a VPLS site and an L3 Access Interface for which you want to create the Routed Vpls Connector. The site and access interface must be on the same NE.
2. Right-click on either of the icons and choose Create Routed Vpls Connector. The RoutedVplsConnector (Create) form opens.
3. Configure the parameters.
4. Save the changes and close the form.
5. Go to Step 14 .

9

To create an SCP Connector:

1. Select the ports for the two sites for which you want to create the SCP Connector.
2. Right-click on either of the of the ports and choose Create Scp Connector. The ScpConnector (Create) form opens.
3. Configure the parameters.
4. Save the changes and close the form.
5. Go to Step 14 .

10

To create a Spoke Connector:

1. Select the two sites for which you want to create the Spoke Connector. These sites must be two different NEs that exist in two separate services.
2. Right-click on either of the sites and choose Create Spoke Connector. The SpokeConnector (Create) form opens.
3. Configure the parameters.
4. Save the changes and close the form.

11

To create new services within the composite service view:

1. Right-click on the flat topology map background and choose Create. A list of services appears.
2. Choose the service type you want to create. The *Service* (Create) form opens.
3. Configure the required parameters for the service and save the changes and close the form.

   See the appropriate service management chapter for information about creating a specific service type.

12

To add services to the composite service:

1. Right-click on the flat topology map background and choose Add→Service(s). The Add Services - Composite Services form opens.
2. Select a service. The service is displayed on the map. If the service includes service sites, the sites are also displayed on the map.

13

To add service sites to an existing service in the composite service:

1. Right-click on the flat topology map background and choose Add→Service Site(s). The Find Service form opens.
2. Select a service.
3. Select one or more NEs. The *Service* Site (Create) form opens. If you select more than one NE, the *Service* Site, Multiple Instances (Create) form opens.
4. Configure the service site and save and close the form. The new service site is displayed on the map.

   See the appropriate service management chapter for information about configuring a specific service site type.

14

Save the changes and close the forms.

END OF STEPS

85.18 To delete a composite service

CAUTION

Service disruption

Deleting a service may result in a service disruption for customers.

Consider the implication of deleting the service before proceeding.

85.18.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

As required, configure the filter criteria to locate the service or range of services to be deleted. A list of services appears at the bottom of the Manage Services form.

3

Choose a service or a range of services from the list.

4

Click Delete. A warning form appears. This form is dynamic based on the priority of the service. Perform one of the following:

a. For services with a low priority, go to Step 5 .

b. For services with a medium priority, configure the “Enter the highest priority of the service being deleted” text field by typing: Medium. Go to Step 5 .

c. For services with a high priority, configure the “Enter the highest priority of the service being deleted” text field by typing: High. Go to Step 5 .

5

For all services regardless of how their priority is configured, acknowledge the check box that prompts you confirm that you understand the implications of deleting the service.

Note: If you select multiple services with different priorities, you must enter the highest priority level of selected services before you can delete the services.

6

Click Yes to confirm the action. The service is deleted and removed from the list.

7

Close the Manage Services form.

END OF STEPS

86 Dynamic service management

86.1 Overview

86.1.1 Purpose

This chapter provides overview and procedural information on dynamic service management in the NFM-P.

86.1.2 Contents

Dynamic service management

86.2 Overview

86.2.1 General information

A dynamic service is a service in which sites are created using information contained in RADIUS Change of Authorization messages instead of manual operator input.

You can configure dynamic services for the following service types:

- Epipe
- IES
- VPLS
- VPRN

86.2.2 Dynamic service structure

The creation of a dynamic service is triggered by a CPE device establishing a control channel (PPPoE, IPoE, or DHCP) with an L3 NE. The control channel is authenticated using RADIUS, and the RADIUS server sends Access-Accept or Change of Authorization messages that contain the parameters for the dynamic service required by the triggering device.

The L3 NE passes the information from the RADIUS messages to a Python script, which returns CLI commands that create the required dynamic sites on the NE.

Data-triggered RADIUS authentication is possible through association of the dynamic service policy with a VPLS capture SAP, allowing dynamic service policy binding to an MSAP. The dynamic service policy is associated with a RADIUS server policy for authentication purposes. In deployments where RADIUS is used for accounting and dynamic changes (CoA) but cannot provide actual service provisioning parameters, you can configure a local authentication database object and associate it with the dynamic services policy, in place of a RADIUS server policy. In this case, the local authentication database is also associated with the dynamic service-enabled NE.

Figure 86-1 Basic dynamic service

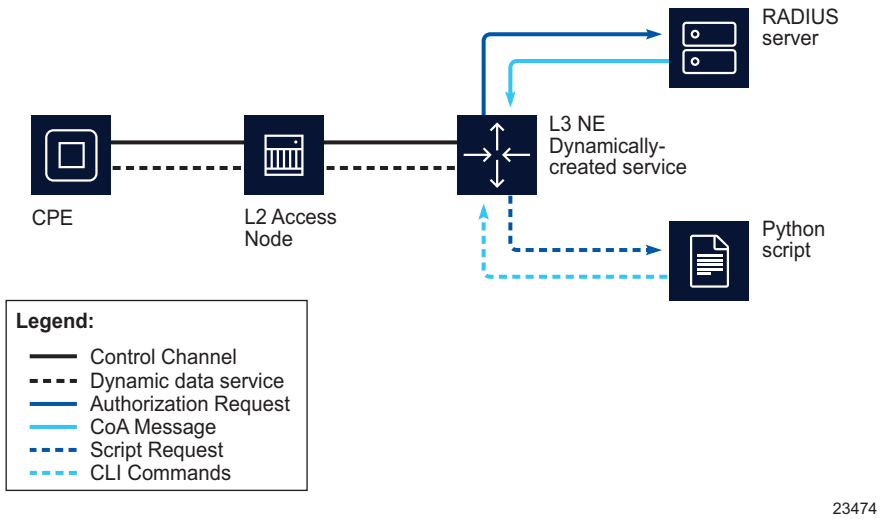

When the control channel fails - for example, when the CPE device is removed - the NE deletes any dynamic services that were created for the control channel.

86.3 Dynamically-created objects

86.3.1 General information

Network objects in a dynamic service are read-only on the NE and in the NFM-P, and cannot be modified except through RADIUS CoA messages. As a result, some operations are not supported on dynamically-created sites. When a dynamic object is created or deleted, the action is recorded in the dynamic service activity log. This section describes the unsupported operations, and the dynamic service activity log.

Dynamically created service sites are discovered and synchronized with the NFM-P in the same manner as manually created service sites. A service that contains a dynamically created object has the Contains Dynamically Created Sites parameter enabled, and cannot be deleted until the dynamically created object is removed from the service.

86.3.2 Unsupported operations

Table 86-1 Unsupported operations on dynamically created objects

Operation	Description
Object configuration operations	
Delete	You cannot delete dynamically created objects. A RADIUS CoA message must request their deletion, or the control channel for the dynamic service that created the object must be disconnected.

Table 86-1 Unsupported operations on dynamically created objects (continued)

Operation	Description
Create child objects	You cannot create objects within dynamically created sites; for example, subscriber interfaces or spoke SDP bindings.
Topology map actions	You cannot interact with dynamically-created objects that appear on the service topology map, except to display their parameters.
SAP copy/move	Dynamically created SAPs are not included in SAP copy or move operations. The operations proceed normally, but skip any dynamically created SAPs.
MDA or card removal	You cannot delete MDAs or cards with dynamic sites.
Automated operations	
Automatic MEP/MIP creation	MEGs skip MEP or MIP creation on SAPs and SDP bindings that are on a dynamically-created site.
SDP binding auto-completion	Dynamically-created sites are skipped during SDP binding auto-completion.
Service Throughput site list	Dynamically-created sites do not appear on the Service Throughput form

86.3.3 Dynamic services activity log

The NFM-P maintains a log of dynamic service actions for each NE, adding an entry whenever an object is created or deleted on the NE. Log entries are deleted after a configurable amount of time, as specified by the ageout constraint policy defined in the NFM-P. For information about configuring the ageout constraint policy, see the *NSP System Administrator Guide*; for information about viewing the dynamic services activity log, see 86.9 “To view the dynamic services activity log” (p. 2789) .

86.3.4 ID range behavior

Each NE that is configured for dynamic service must specify an ID range for dynamically-created sites. Dynamically-created sites can only use IDs from the specified range, and non-dynamic sites cannot use IDs from that range. When a dynamic service is created on an NE, the next available ID from the specified range is used.

86.4 Workflow to create dynamic services

86.4.1 Purpose

The following workflow lists the high-level steps required to configure an NE to process dynamic service requests. As a prerequisite for creating a dynamic service, this workflow assumes that the following appropriate preconfigurations have been performed:

- pre-discovery CLI modifications
- discovery including mediation configuration with CLI user name and password
- creation of python scripts to be called by the NE; for more information about creating dynamic service python scripts, and CLI operations supported for python scripts, see the *7750 SR OS Triple Play Guide*.
- creation of RADIUS policies

86.4.2 Stages

1

Configure a RADIUS script policy; see 64.23 “To configure a RADIUS script policy” (p. 1865).

2

Configure a dynamic service policy; see 86.5 “To configure a dynamic service policy” (p. 2786).

3

Where required, configure a local authentication database, with binding to the dynamic service policy; see 86.6 “To configure a local authentication database” (p. 2787).

4

Configure the NE for dynamic service; see 86.7 “To configure an NE for dynamic service” (p. 2788).

5

Where data-triggered authentication is required, configure the dynamic service policy on a VPLS capture SAP; see 74.26 “To configure a capture SAP” (p. 2049).

6

Configure a RADIUS-authenticated control channel. See the *7750 SR OS Triple Play Guide* for more information about dynamic service control channels.

Dynamic service management procedures

86.5 To configure a dynamic service policy

86.5.1 Steps

1

Choose Policies→Dynamic Services→Dynamic Service Policy from the NFM-P main menu. The Dynamic Service Policy form opens.

2

Click Create or choose a dynamic service policy and click Properties. The Dynamic Service Policy (Create|Edit) form opens.

3

Configure the required general parameters.

4

Select a RADIUS script policy for the dynamic service policy.

5

Select a CLI user for the dynamic service policy.

6

To configure RADIUS or local authentication for the dynamic services policy:

RADIUS authentication and local authentication are mutually exclusive.

1. Click on the Authentication tab and do one of the following:
   - Select a RADIUS server policy and configure the RADIUS Password parameter.
   - Select a local authentication database.
2. Save your changes and close the form.

7

To configure the accounting instances of the dynamic services policy:

1. Click on the Accounting tab. Each dynamic services policy has two accounting instances.
2. Choose an accounting instance and click Properties. The Accounting Instance form opens.
3. Configure the parameters.
4. Select a RADIUS server policy for the accounting instance.
5. Configure the parameters in the Update Interval panel.
6. Save your changes and close the form.

8

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs

END OF STEPS

86.6 To configure a local authentication database

86.6.1 Steps

1

Choose Policies→Dynamic Services→Local Authentication Database from the NFM-P main menu. The Local Authentication Database form opens.

2

Click Create or choose a local authentication database and click Properties. The Local Authentication Database (Create|Edit) form opens.

3

Configure the required general parameters.

4

To configure users in the local authentication database:

1. Click on the Users tab.
2. Click Create or select an existing user and click Properties. The Dyn Svc Local Auth Db User form opens.
3. Configure the required general parameters.
4. Click on the SAP User tab to configure per-SAP users for the database user.
5. Click Create or select an existing SAP user and click Properties. The Dyn Svc Local Auth Db Sap form opens
6. Configure the required general parameters.
7. Select a dynamic service policy.
8. Click on the Accounting tab to configure accounting instances for the SAP user.
9. Click Create or select an existing accounting instance and click Properties. The Dyn Svc Local Auth Db Acct form opens.
10. Configure the required parameters.
11. Save your changes and close the forms.

5

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs

END OF STEPS

86.7 To configure an NE for dynamic service

86.7.1 Steps

1

Choose Manage→Service→Dynamic Service NE Configuration from the NFM-P main menu. The Dynamic Service NE Configuration form appears.

2

Choose an NE and click Properties. The Dynamic Service NE Configuration (Edit) form opens.

3

Disable the Disable Creation parameter to allow the dynamic creation of network objects on the NE.

4

Define an ID range using the Start and End parameters. Dynamically-created objects are created using IDs from within the specified range.

5

Configure the Access-Accept parameter in the Timers panel.

6

Save the changes and close the form.

END OF STEPS

86.8 To list dynamically created objects on an NE

86.8.1 Steps

1

Choose Manage→Service→Dynamic Service NE Configuration from the NFM-P main menu. The Dynamic Service NE Configuration form appears.

2

Configure the filters and click on the Search button. A list of NEs that support dynamic services appears.

3

Choose an NE and click Properties. The Dynamic Service NE Configuration (Edit) form opens.

4

Click on one of the following tabs to search for and list the related dynamically created objects on the NE:

Click on one of the following tabs to search for and list the related dynamically created objects on the NE:

- Dynamic SAP Info
- Root Objects
- Script Snippets
- Service Sites

Note: The list of objects in the Dynamic SAP Info tab is automatically updated.

5

Close the forms.

END OF STEPS

86.9 To view the dynamic services activity log

86.9.1 Steps

1

Choose Manage→Service→Dynamic Service NE Configuration from the NFM-P main menu. The Dynamic Service NE Configuration form appears.

2

Choose an NE and click Properties. The Dynamic Service NE Configuration (Edit) form opens.

3

Click on the Dynamic Service Activity tab.

4

Configure the filter and click Search. A list of dynamic service actions appears.

5

Close the forms.

END OF STEPS

87 Application assurance

87.1 Overview

87.1.1 Purpose

This chapter provides overview and procedural information about the NFM-P application assurance functions.

87.1.2 Contents

Application assurance

87.2 Overview

87.2.1 General information

Application Assurance (AA) is a service-enabling technology that enhances QoS functions by providing additional traffic control, traffic diversion, application identification, statistics collection, and data reporting.

AA and dynamic subscriber policy control allow a broadband network to provide application-based subscriber management for Internet access.

The fundamental elements of AA processing are:

- identification of the traffic on a per-flow or per-session basis
- policy-based treatment of the identified traffic

Shallow packet inspection (SPI), inspects Layer 1 to Layer 3 traffic information and does not provide application-based QoS information. AA enables deep packet inspection (DPI) of subscriber traffic using policies that define the action to perform for specific traffic types on a per-subscriber basis.

Subscriber traffic is selected for processing by AA, then inspected by an ISA-AA MDA or ESA VM in a logical group. See 87.3 “ISA-AA groups and partitions” (p. 2794) for more information about ISA-AA groups and partitions.

Note: AA policies can be configured on an NE only when an ISA-AA group exists on the NE.

The benefits of AA for residential and business service providers are:

- enhanced application-level QoS using policy-based traffic management
- traffic flow thresholds and threshold-crossing notifications
- application-aware reporting and performance analysis
- greater network security
- high scalability and flexibility to control network costs

The NFM-P supports AA functions on the 7450 ESS and 7750 SR.

Note: AA is not supported on a single-slot chassis.

87.3 ISA-AA groups and partitions

87.3.1 General information

The NFM-P supports the configuration of ISA-AA groups and partitions. An ISA-AA partition is a unique child object of an ISA-AA group. The partition can be assigned an AA policy. There is no relationship between partitions of different ISA-AA groups. You can configure up to 128 partitions per group.

You can divide an ISA-AA group into partitions that are dedicated to VPN-specific AA services. A partition can have a set of VPN-specific custom protocols, applications, application group definitions, policy definitions, and reporting. Each partition policy can be divided into multiple application QoS policies using ASOs.

To apply one policy to multiple groups or partitions, use a policy sync group; see 87.16 “Policy sync groups” (p. 2808).

87.3.2 ISA-AA groups

You can perform the following operations on an ISA-AA group:

- Assign ISA-AA MDAs or ESA VMs and create AA partitions.
- Specify FCs to be diverted for inspection, and choose the AA policy to apply to the group.
- Configure redundancy and a bypass mode to protect against equipment failure.
- Configure QoS on IOMs that host ISA-AA traffic.
- Configure ISA capacity planning using low and high thresholds.

Residential services are an example where all AA services can be configured as part of a single group that includes all ISA-AAs. The configuration provides the management of common applications and reporting for all subscribers and services, with common or per-customer AQP using ASO characteristics to divide the AQP of the ISA-AA into per-application profile QoS policies.

Multiple ISA-AA groups can also be used to create separate services based on different sets of common applications, traffic diversion needs, or different redundancy models.

Multiple ISA-AA groups can be used for:

- a mix of residential and business customers
- different business VPN verticals
- business services with a common template base but different levels of redundancy, FC diversion, or scaling per group

See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) for information about configuring ISA-AA groups.

87.3.3 ISA-AA partitions

ISA-AA groups and partitions improve the scaling of policies. When partitions are configured, application identification and policy configuration apply only to the specified partitions in the ISA-AA group, and not to any other partition in the AA group. When partitions are not configured, the ISA-AA group acts as a single partition.

Although the definition of application profiles and related ASO characteristics is in the context of a partition, each application profile name in an NE configuration must be unique. Application, application group, and AQP definitions are specific to a partition.

ISA-AA partitions support accounting and customized reporting for every AA subscriber associated with a partition, and allow you to do the following:

- Define different types of reporting and accounting policies for different partitions in a single AA group.

- Display AA group protocol statistics with partition visibility; for example, you can view protocol counts for each partition in a group.

When you create or delete an ISA-AA partition, a default AA accounting policy is automatically created in or deleted from the ISA-AA partition.

See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) for information about configuring ISA-AA partitions.

87.4 AA components

87.4.1 General information

The NFM-P supports AA component creation and configuration using configuration forms and scripts. The following figure shows the AA components.

Figure 87-1 AA components

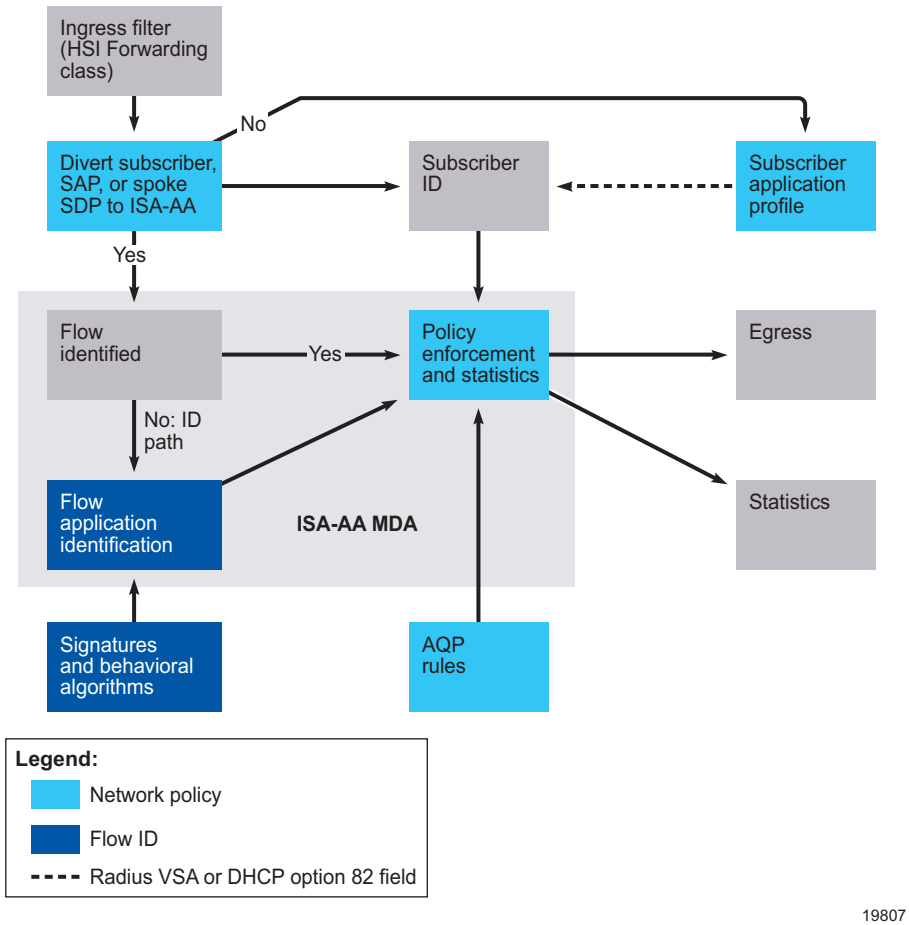

The following figure shows the relationship between AA components.

Figure 87-2 AA policy structure

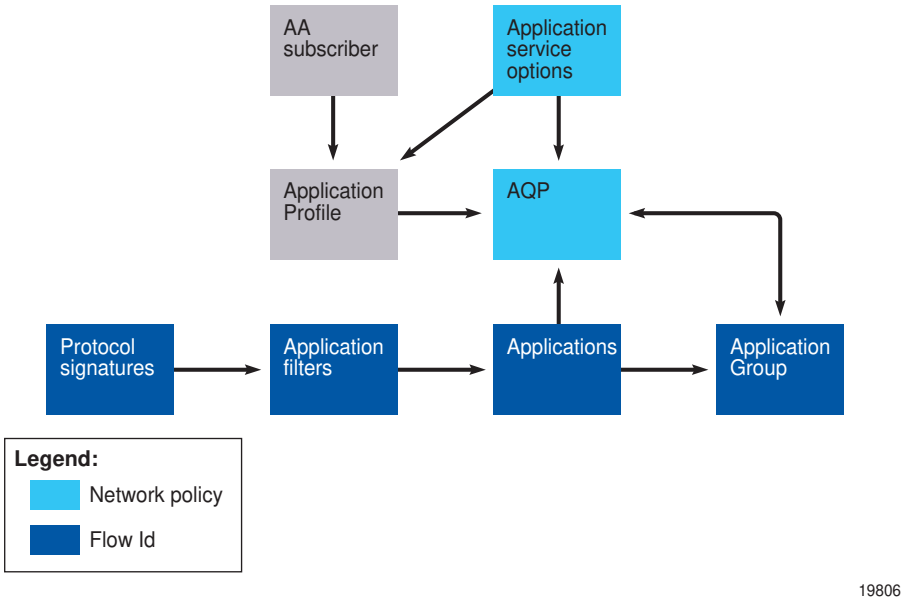

87.5 AA group policies

87.5.1 Components

An AA group policy includes the following components, as required:

- applications
- application groups
- charging groups
- application profiles
- application filters
- application service options (ASOs)
- application QoS policies (AQPs)
- custom protocols

The following table describes the components.

Component	Description
Application	Identifies the type of IP payload in the subscriber flow An application provides a definition and description to the application names supported by an application filter.
	Defaults One predefined application (“unknown”) is provided. The unknown application cannot be modified. All other applications must be configured.
	Examples/use cases Webscale companies, CORBA, DNS Server
	Notes The application name is a key match criterion within the AQP rules that are applied to the IP traffic of a subscriber. The application name is also the unique identifier of the application object for reporting systems such as NSP Analytics Reports. Network operators can change the application group that an application is associated with, or view the application filters that define the application.
Application group	A container for multiple applications
	Defaults A set of default application groups is provided. At least one default application group (“unknown”) must be associated with each NE. The default application groups cannot be modified.
	Examples/use cases File hosting, gaming, mail, peer to peer
	Notes Multiple applications can be assigned to an application group. An application can only belong to one group. Applications that are not assigned to an application group are automatically placed in the “unknown” application group.

Component	Description
Charging group	A set of applications that you need to bill for in a specific way, such as usage-based billing
	Defaults If an application is not associated with a charging group, a default charging group is assigned. If tethering detection or flow attribute matching is in use with 7450 ESS, a default tethered charging group can be configured.
	Examples/use cases Applications that are free to the end user can all be associated with a zero-rate charging group. A use case for charging groups is bill-shock avoidance.

Component	Description
Application profile	Enables AA service for traffic to and from an ESM subscriber, ESM subscriber host, SAP, or spoke SDP. Each application profile is unique and defines the AA service that the AA subscriber receives. The type of traffic is configured in the system-wide configuration of QoS FCs to be diverted to the ISA-AA MDA for subscribers with AA enabled. The FC is used for any subscriber traffic that a service provider needs to inspect using AA. A subscriber can be assigned an application profile that affects every host of the subscriber. For SAP or spoke SDP AA subscribers, an application profile can be assigned that affects the traffic originating from or destined for the SAP or spoke SDP. For subscribers with application profiles that enable AA, traffic is diverted to the active ISA-AA MDA using ingress QoS policy filters. The filters identify classes that can be diverted for AA. The system identifies and diverts traffic for any subscriber to the ISA-AA MDA, according to the application profile. Diversion to the ISA-AA MDA depends on the ISA-AA MDA status. If a subscriber is not configured to divert traffic to the ISA-AA MDA, normal ingress forwarding occurs.
	Defaults One predefined application profile (“default”) is provided. By default, subscribers are not assigned to an application profile and traffic is not diverted for AA analysis.
	Notes • Application profiles are customer-defined and created using the configured ASO characteristics. • Application profiles allow ASO characteristics to be associated with AA subscribers. • Local instances of the application profile are configured to bind to specific objects of an NE (for example, L2 access interface, SAP, local subscriber explicit map entry, and so on). • Application profiles can be assigned only when ISA-AA cards are assigned to an ISA-AA group. • Global application profiles must be manually distributed to the NEs. • Application profiles can be assigned a capacity cost for subscriber load balancing among ISAs within the AA group.

Component	Description
Application filter	Application filters are provided as an indirect action between protocols and applications to allow the addition of variable parameters, for example, port numbers and IP addresses, to an application definition. An application filter contains numbered entries that define the use of protocol signatures and other application criteria. Multiple filter entries can be used to define an application, but each application filter entry maps to one application.
	Notes A traffic flow may have multiple filter entries. The entries are applied to the flow sequentially, and the matching entry with the lowest ID value is applied first. A traffic flow can be assigned to only one application. You can accommodate the insertion of new entries in a list of entries by renumbering one or more of the existing entries using a GUI or OSS client. See87.31 “To renumber application filter entries” (p. 2842) for information about renumbering application filter entries using the GUI. An application must be configured before the associated application filter is defined.

Component	Description
Application service option (ASO)	A series of operator-defined application characteristics that define the service provider and customer network functions that are common to sets of subscribers ASOs prevent subscribers from requiring each subscriber-specific entry in the application QoS policies for standard network services. ASOs are assigned one or more characteristic values to define service offerings to customers. ASO characteristics are used: • as input to application profiles • by application QoS policies to influence how specific traffic is checked and how policies are applied
	Examples/use cases • Entry for each managed application group; for example, VoIP, P2P, and HTTP • Multiple entries where specific applications in an application group can be individually managed as service parameters; for example, HTTP content from a specific content provider, or streaming video from network television or games • HSI tiers, for example, Gold, Silver, and Bronze, that specify bandwidth levels • Bandwidth parameters for each service option
	Notes You can configure ASO characteristics for a subscriber using an AA subscriber policy override. An AA subscriber policy override can be configured for a SAP or spoke SDP binding. The AA subscriber must have an application profile assigned, or the subscriber policy override is rejected. You can retrieve on demand the ASO values that are assigned to a SAP or spoke SDP binding subscriber from the properties form of the subscriber. In a typical application, ASOs define the HSI service parameters. • ASOs are optional; AA can check the subscriber IP traffic without the use of application service options. • ASOs can be configured for each AA group policy. • ASO characteristics are used to define application profiles and operator-defined AQP rules. • The set of ASOs represent network-wide menus of service capabilities that are available to subscribers.

Component	Description
Application QoS policy (AQP)	List of rules that define the match criteria and action to be performed on all traffic flows. The AQP rules use the application groups, applications, and so on, as the match criteria. The output of the AQP rules defines the policy actions to perform.
	Notes AQP rules consist of match and action criteria: • Match—Refers to the criteria used to identify a flow in order to apply actions such as dropping, forwarding, mirroring, and policing of bandwidth and flow. Matches cannot be made against protocols. Match criteria can be a combination of the following: - applications or application groups - ASO characteristics - flow direction - AA subscribers - flow source or destination IP address and/or port - flow attributes • Action—Defines AA actions to be applied to traffic. For example, you can apply a set of actions such as bandwidth policing, packet discards, QoS remarking, and rate limiting to a flow.
Custom protocols	Configurable strings of up to 16 hexadecimal octets for pattern-matched application identification in the payload of TCP- and UDP-based applications The match applies to the client-to-server, server-to-client, or any direction for TCP-based applications, and to any direction for UDP-based applications.
	Notes You can configure a custom protocol description, custom protocol ID, and shutdown. When a custom protocol is administratively disabled, traffic is identified as though the protocol is not configured. Custom protocols and Nokia protocols are distinct; the span of a custom protocol is limited to the group or partition associated with an AA policy. Nokia protocols span all groups and partitions. All application filter entries, except strings, are supported. Custom protocol statistics collection on an ISA-AA partition group or special study subscriber is supported.

87.6 AA Cflowd

87.6.1 Data collection

ISA-AA groups and partitions support Cflowd sampling and TCP performance data collection for AA applications and application groups. AA Cflowd statistics are collected by the NSP, if flow collection

is enabled in the NSP deployment. The NSP can forward the data to an OSS or to NSP Analytics Reports for reporting and analysis.

Cflowd analysis is performed on a per-flow basis. Statistics are created at the end of a flow session. Standard Cflowd sampling collects the source and destination addresses and the protocol used, AA Cflowd also collects the start time, duration, and context IP address. This additional information can be used to customize network usage reporting.

If an NFM-P deployment includes an auxiliary database, you can store the AA Cflowd statistics for use as the source data for reports in NSP Analytics Reports.

An ISA-AA group supports one Cflowd instance. See Chapter 13, "Logical group object configuration" for information about configuring AA Cflowd on an ISA-AA group, and the *NSP Statistics Management Guide* for AA flow configuration information.

AA Cflowd policies

You can use the NFM-P to create AA Cflowd group policies for AA groups and partitions. The parameters in an AA Cflowd group policy are the same as the Cflowd parameters that you configure on an ISA-AA group.

When you use an AA Cflowd group policy to specify the AA Cflowd collection criteria, you can apply the policy to multiple AA sites. You can also use the policy as the master policy in a policy sync group to ensure that the distributed AA Cflowd group policies remain synchronized with the master, and to facilitate the deployment of a policy update to multiple local policy instances at one time.

AA Cflowd templates

You can apply a cflowd template to an AA Cflowd policy on the 7750 SR NE. The policy must either be for an ISA-AA group with partitions disabled, or be a group level policy for an ISA-AA group with partitions enabled. Templates are not available for a partition level AA Cflowd policy. Cflowd template allows for groups of parameters and fields to be collected together under a combined heading such as Volume.

See 87.33 "To configure an AA Cflowd group policy" (p. 2844) for AA Cflowd group policy configuration information. See 87.16 "Policy sync groups" (p. 2808) for information about policy sync groups.

87.7 AA protocol signatures

87.7.1 Protocol signature set

The NFM-P generates a set of signatures that identify AA protocols.

The signature set includes:

- protocol support summary—list of protocols that can be identified with the load using a combination of pattern and behavioral techniques. The protocols are used to generate statistics by protocol and as input in combination with other information to identify applications.
- pattern signatures—set of pattern-match signatures used in analysis
- behavior signatures—set of diagnostic techniques used in analysis

Because protocol signatures are intended to be the most basic block of application identification, other AA components, such as application filters, are provided to further customize protocol signatures. Customization reduces the need for a new protocol signature load when a new application may need to be identified.

Each protocol can be referenced in the definition of one or more applications by the application filter definition. The assignment of each supported protocol to an application filter or application is optional, and allows the addition of new signature protocols without the need to update the application filter and applications.

Protocol signature upgrades

The NFM-P supports protocol signature upgrades without affecting policy behavior. You can obtain new protocol signatures by dynamically upgrading only the ISA-AA MDA software on an NE, if the new software is from the same major release. All new signatures in R2 or later of a major release are disabled during an upgrade to ensure that policies and services are not affected. See 26.16 “To upgrade the ISA-AA MDA software” (p. 795) for information about how to upgrade the ISA-AA MDA software.

If the software is from a different major release, you must upgrade the entire device to obtain the new signatures. See “NE software upgrade overview” (p. 771) for information about performing device software upgrades.

The protocols in R1 of a release are designated as parent signatures and cannot be disabled. The protocols must be enabled on a per-protocol basis to take effect. See 87.70 “To enable an AA protocol signature” (p. 2884) for information about enabling AA protocols.

87.8 Dynamic experience management

87.8.1 ANL capacity monitoring

Dynamic Experience Management (DEM) employs DPI and congestion control to maintain QoE. For each ANL, or bandwidth bottleneck point, the system continuously monitors the capacity at the ANL. If congestion is predicted or detected at an ANL, the DEM gateway uses bandwidth policers to give priority to delay-sensitive applications.

Multiple ANLs can be configured when multiple ISAs are configured within an ISA-AA group; see 87.36 “To configure an AA Access Network Location” (p. 2849).

87.9 AA policers

87.9.1 Bandwidth or flow limitation

Policers allow IP traffic on an interface to be limited.

AA policers can be bandwidth or flow limiting, and can have one of the following scopes:

- system scope—limits all traffic entering an ISA-AA MDA
- subscriber scope—limits apply only to the traffic of a subscriber

After a policer is assigned by an AQP for one traffic direction, the same policer cannot be assigned in the other direction. AQP rules with policer actions must specify a traffic direction other than “both”.

Congestion override

Congestion override can be configured on an AA policer with a subscriber scope. The policer is triggered when the subscriber ANL is in a congestion state. You can configure the congestion thresholds in the policer configuration form. Congestion override configuration can assist with DEM.

Note: If both congestion override and ToD override are configured and conflict, the congestion override will take precedence.

See 87.26 “To configure an AA policer” (p. 2830) for configuration information.

87.10 AA GTP firewalls

87.10.1 Stateless GTP firewall

You can use the NFM-P to create a stateless GTP firewall using GTP and SCTP traffic filters, and enforce a GTP flow count limit in an AA flow-count policer. See 87.27 “To configure an AA GTP firewall” (p. 2831) for information about AA GTP firewall configuration.

You can view the local GTP and SCTP filter statistics from the Statistics tab of a local filter properties form. The GTP firewall statistics for an ISA-AA group or partition are available on the Statistics tab of the group or partition properties form.

87.10.2 GTP-c firewall for S8 and Gn

You can use the NFM-P to create a firewall for AA on Gn or S8 interfaces. This type of firewall ensures that total traffic to or from the SGSN is enforced at a configured rate.

See 87.28 “To configure an AA GTP-c firewall for S8 or Gn” (p. 2835) for information about AA GTP-c firewall configuration.

You can view the local GTP filter statistics from the Statistics tab of a local filter properties form. An ISA card has to be assigned as a group member for the statistics to be collected. The GTP firewall statistics for an ISA-AA group or partition are available on the Statistics tab of the group or partition properties form.

87.11 AA statistics TCAs

87.11.1 Threshold-crossing alarms

Threshold-crossing alarms (TCAs) can be raised for events related to excessive flow counts or flow rates on an ISA-AA group or partition. The TCA mechanism enhances firewall security using AA policers and the watermarks specified in GTP, SCTP, or session filters. See 87.41 “To configure an AA statistics TCA” (p. 2854) for configuration information.

87.12 AA flow watermarks

87.12.1 Flow-table consumption monitoring

AA supports the configuration of high and low thresholds (known as “watermarks”) to monitor flow-table consumption by traps and logs, and also to monitor data path CPU usage. The thresholds are established to alert operators when the flow table approaches maximum capacity, or when AA CPU usage is excessive. When a high threshold is reached, the NFM-P raises an alarm that clears when the alarm condition is no longer present. The alarm is viewable from the Faults tab of the associated ISA-AA group properties form.

AA flow watermark statistics are available from the Statistics tab of an ISA-AA group properties form.

See 87.39 “To configure an AA flow watermark” (p. 2852) for configuration information.

87.13 AA transit IP and transit prefix policies

87.13.1 Transit subscriber discovery

AA transit IP and transit prefix policies define how AA transit subscribers are created by an NE. Transit IP policies can be configured to discover dynamic transit subscribers via DHCP or RADIUS, and can also be configured with a list of static transit subscribers. Transit IP policies can also be configured to discover dynamic transit subscribers when traffic is detected on a parent SAP or spoke SDP binding.

Transit prefix policies allow you to specify transit subscribers by using network or subscriber IP ranges.

AA transit subscribers can persist in the NFM-P database. Transit subscriber aggregators allow you to aggregate multiple database persisted transit subscriber instances for VPN sites.

Because of the potentially high rate at which AA transit subscribers are created and deleted, a throttling mechanism is available on the NE for the associated creation and deletion traps. See 87.50 “To configure trap throttling for AA transit subscriber creation and deletion” (p. 2864) .

87.14 AA HTTP redirect policies

87.14.1 HTTP/HTTPS redirection

AA HTTP redirect policies allow you to configure HTTP/HTTPS redirection, specify templates for local NE termination of HTTP errors, and configure web browser redirection to customized error pages according to HTTP error codes. See “Application assurance procedures” (p. 2819) for configuration information.

87.15 AA URL Filter policies

87.15.1 URL filtering

AA URL filter policies allow you to configure filtering of URL access by category. This will allow for parental controls for household internet access, and control of access to the Internet for smart devices.

An AA URL filter can filter based on an AA URL list, or use an ICAP server or a web service to communicate with an external server or database to determine filtering decisions. See 87.68 “To configure an AA URL filter” (p. 2880) for configuration information.

87.16 Policy sync groups

87.16.1 Centrally defined policy components

Policy sync groups allow AA group policy or AA Cflowd policy components to be centrally defined and applied to multiple ISA-AA groups and partitions. A policy sync group includes a master policy that defines a set of common policy components, and a list of member AA group or AA Cflowd policies that are synchronized with the master.

A synchronization action that you perform using a policy sync group affects only the classes named in the group, and applies to each AA group and partition that is a member of the policy sync group.

The following synchronization options are available:

- The master policy classes override the member policy classes.
- The master policy classes are added to each member policy.

After you create a policy sync group, you can use the policy sync group to do the following:

- Distribute AA group policy or AA Cflowd group policy updates to all members using the Synchronize function.
- Audit the member policies to identify variances from the master policy settings.
- Compare two member policies.

If a policy audit identifies a difference between the master policy and a member policy, the NFM-P raises an alarm. The differences are listed as the affected objects of the alarm.

See 87.34 “To configure a policy sync group” (p. 2846) for policy sync group configuration information. See 87.35 “To audit, compare, or synchronize policies using a policy sync group” (p. 2848) for information about how to use a policy sync group to compare, audit, or synchronize member policies.

AA reporting

87.17 Overview

87.17.1 General information

AA reporting attributes are metadata parameters that specify how application performance data is reported to systems such as NSP Analytics Reports; the parameters are not deployed to NEs.

The parameters are configurable on the following object properties forms:

- AA application
- AA application group
- subscriber objects:
  - subscriber instance
  - service
  - SAP
  - customer
  - spoke SDP binding
  - transit subscriber

See 87.54 “To enable application performance reporting on a service” (p. 2867) , 87.55 “To configure application performance reporting on a SAP or SDP binding” (p. 2868) , 87.56 “To configure application performance reporting for a transit subscriber” (p. 2869) , and 87.57 “To disable application performance reporting on a service” (p. 2870) for configuration information.

87.17.2 Enabling reporting

Application performance reporting enables AA performance and flow-based volume reporting to functions such as NSP Analytics Reports.

For increased reporting granularity, and to reduce the amount of reported information, you can use default DCP groups to segregate the Internet or intranet traffic of an entire network, an IP address range, or a single IP address.

You can create custom DCP groups for service objects to report only the required business application data for a service that has an application profile. See “Default and custom DCP groups” (p. 2811).

A reporting system can use the collected data to perform Apdex analysis. Apdex measures the quality of end-user experience with respect to Internet applications.

See “AA reporting” (p. 2809) for more information about application performance reporting.

87.17.3 Apdex reporting

Application performance index, or Apdex reporting, is a standard method of estimating the end-user satisfaction associated with online application delivery; for example, you can configure Apdex to monitor a VoIP or video stream, which requires very low delay.

Apdex measures quality indicators that include the following:

- round-trip time
- mean delay
- standard deviation delay
- packet loss

Visit the Apdex Alliance website for detailed Apdex information.

The NFM-P supports Apdex threshold configuration based on the following:

- application type
- application delay sensitivity
- application tier level

The NFM-P uses tiers to classify the application types. An Apdex tier signifies the traffic importance and specifies which set of application performance thresholds are in effect.

Each tier has default end-user frustration and tolerance thresholds for quantifying the application performance. You can modify the threshold values for each tier to attain the desired quality of end-user experience.

For residential Apdex, the thresholds are configurable for an AA group or partition, application group, and application. A child object, for example, an application, can inherit the threshold values from the parent application group or AA partition. A child object can also override the parent thresholds.

You can also configure tier-based Apdex reporting for the following service objects:

- spoke SDP bindings
- access interfaces
- business transit subscribers

Note: You must associate an application profile with an object in order to configure AA reporting on the object.

87.17.4 IP detail reporting

You can configure service-level IP detail reporting, which collects information about the most active end users, by volume, of applications and applications groups.

The following service types support IP detail configuration:

- IES
- VLL Epipe
- VLL Ipipe

- VPLS
- VPRN

87.17.5 Usage-based billing

The NFM-P supports application-specific reporting to systems such as NSP Analytics Reports based on configurable charging groups. You create charging groups as child objects of an AA group policy and can assign them to the following:

- applications
- application groups
- charging group thresholds
- AA subscriber statistics objects

You can assign one charging group to multiple applications or application groups, and can associate an application or application group with only one charging group.

You can configure charging groups with specific usage quotas and notification thresholds, and assign the thresholds to application profiles.

87.17.6 DCP groups

An NFM-P operator can configure business or residential DCP groups that specify how to summarize the collected reporting data for upstream reporting systems, which use the group specifications as an aggregation framework for analysis.

A DCP group organizes subscriber traffic based on the IP address range in a subnet rule that represents a subscriber address range or the server address range of a content provider. A DCP group can contain up to 100 subnet rules.

A residential DCP group is configured in an AA group policy. A business DCP group is configured as part of a service, for example, on the Application Assurance tab of an Epipe. Either can contain up to 65 535 DCP groups.

You can use DCP groups to help identify traffic patterns such as the following:

- the volume of application traffic delivered to a specific regional market, for product management or policy development
- the application performance in a specific regional market; for example, how streamed video that is external in origin affects a local video-on-demand service
- the source of continuous high-volume traffic that is external in origin

You can use the NFM-P policy sync group function to synchronize the DCP groups in one AA group policy with other AA group policies. See 87.16 “Policy sync groups” (p. 2808) in 87.4 “AA components” (p. 2796) for more information.

Default and custom DCP groups

When you enable application performance reporting in an AA group policy or service, the NFM-P creates default IPv4 and IPv6 Intranet and Internet DCP groups that separately aggregate all

subscriber and content-server traffic. You can add DCP subnet rules to a default DCP group, or create multiple custom DCP groups per service.

You can associate one of the following subscriber objects with a custom DCP group in a service:

- service
- SAP
- customer
- spoke SDP binding
- business transit subscriber

A SAP or spoke SDP binding requires an application profile before you can associate it with a custom DCP group; a business transit subscriber is associated with an AA group that is configured for use in the VPN context.

If application performance reporting is disabled in a global AA group policy, the NFM-P retains the associated default and custom DCP groups for future use.

87.18 Flow attributes

87.18.1 Additional AA classification

Flow attributes allow for further classification of flows, other than according to application or application group. Flow attributes can be marked against packets of a flow.

Attributes are assigned to a flow based on characteristics of the flow. A flow may have zero or more attributes assigned to it including competing or conflicting attribute values such as Video and Video ABR.

Flow attributes can be based on the following:

- application filter provisioned: the flow attribute is based on the application filter definition
- flow state derived: the flow attribute is derived by an algorithm based on the protocol in use

Including flow attributes in an AQP allows match criteria and actions based on flow attribute, for example, application of a bandwidth policer or usage-based billing. To include flow attributes in an AQP, flow attribute matching must be enabled in the ISA-AA Group.

AA accounting statistics collection

87.19 Overview

87.19.1 General information

AA accounting statistics provide information about application use in a network. You can configure AA accounting to collect and report statistics when at least one ISA-AA MDA is active. The AA accounting statistics provide information about application use on a SAP or spoke SDP, or by a subscriber.

Note: You cannot use a local AA accounting policy on an NE of a different release to update a global policy using the policy synchronization function; you must use a local policy from the same release of an NE.

AA accounting collects statistics on traffic flows. The NFM-P can collect the following AA statistics types:

- AA application
- AA application group
- AA protocol
- AA subscriber protocol (special study)
- AA subscriber application (special study)
- AA subscriber custom record

The statistics data can be viewed in graphical or tabular form using the NFM-P GUI, or forwarded for reporting and analysis to systems such as a third-party server, and used to create reports in NSP Analytics Reports. You can also configure the NFM-P to export the data to XDR-encoded IPDR files for OSS client retrieval. See 87.19.2 “IPDR data export” (p. 2813) in this section for more information.

AA uses the NFM-P and NE accounting statistics and logging capabilities to collect AA accounting statistics. See “Statistics collection in the NFM-P” in the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting accounting statistics. See 87.37 “To configure an AA accounting policy” (p. 2849) for information about configuring an AA accounting policy.

The NFM-P can be configured to collect statistics for each protocol and application of a specific subscriber. For more comprehensive traffic monitoring, you can enable AA statistics collection on a specific NE for a subset of subscribers. An NE can have only one policy for each AA statistics type that is enabled.

87.19.2 IPDR data export

To support data analytics and reporting functions, the NFM-P can export AA accounting statistics data to files on a main or auxiliary server station in one or both of the following formats:

- IPDR format, using XDR encoding
- format for NSP Analytics Reports

The exported files are stored in the aaAccountingStats directory under the OSS XML output directory on a main or auxiliary server station. Each filename has an AA_*type* prefix, where *type* represents the statistics type. Each filename also includes the associated NE IP address.

Note: The tethering statistics type is modeled and represented differently in each format:

- IPDR:
  TETHER—filename prefix is AA_TETHER
- NSP Analytics Reports::
  SUMMARY—filename prefix is AA_SUMMARY

You can enable and configure the export function during an NFM-P system installation or upgrade, or afterward, using 87.38 “To configure AA accounting file export” (p. 2850).

87.19.3 Special study statistics

The AA subscriber protocol and AA subscriber application statistics are special study statistics for detailed accounting statistics collection on a limited number of subscribers, SAPs, or spoke SDPs on an NE. The number of subscribers, SAPs, or spoke SDPs is limited to constrain the volume of generated statistics.

Special study statistics collection is required for certain reports in NSP Analytics Reports.

Special study statistics are enabled by adding subscribers, SAPs, or spoke SDPs to a list in an ISA-AA group or partition on an NE for detailed traffic monitoring. When a subscriber is on a special study list, the ISA-AA creates one statistics record for each application and application group that is associated with the subscriber. When a SAP or spoke SDP is on a special study list, the ISA-AA creates one statistics record for each application, application group, and protocol flow on the SAP or spoke SDP. See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) for information about configuring special study objects in an ISA-AA group.

87.19.4 Subscriber custom record statistics

An AA subscriber custom record accounting policy applies to all subscribers, SAPs, and spoke SDPs in an ISA-AA group or partition for a specified set of AA protocols, applications, and application groups. The policy enables statistics collection on only the specified objects, which limits the volume of collected data and the statistics collection processing load. For example, an NFM-P operator may require statistics for only three application groups and 10 applications. You can view and configure the AA statistics objects of an AA group or partition from the AA Subscriber Stats Objects tab of an application, application group, or charging group properties form.

DSM AA statistics export

When the Subscriber Scale of an ISA-AA Group is set to Lightweight Internet, you can enable the export of AA statistics according to the settings in a RADIUS accounting policy. You enable the function in the application, application group, and charging group configured as subscriber statistics objects in an AA group policy, and in the AA Subscriber Custom Record accounting policy configuration of an AA accounting policy.

87.19.5 Monitoring usage by charging group

You can enable usage monitoring in an AA Subscriber Custom Record accounting policy, and optionally export the usage statistics to a RADIUS accounting policy. See 87.73 “To configure subscriber usage monitoring” (p. 2887) for configuration information.

Workflows to configure AA

87.20 Workflow to perform hardware procedures for AA configuration

87.20.1 Stages

1

Configure the required ISA-AA groups and partitions on NEs. See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) .

2

Configure AARP instances on an NE. See 87.25 “To configure an AARP instance on an NE” (p. 2829) .

87.21 Workflow to manage AA policies

87.21.1 Stages

1

Configure AA group policies using 87.23 “To configure an AA group policy” (p. 2819) .

As required, configure the following components:

- ASOs and ASO overrides
- application profiles
- application groups
- charging groups
- applications
- application filters
- custom protocols

2

Configure AQPs for AA group policies. See 87.24 “To configure an AQP” (p. 2825) .

3

Configure AA policers. See 87.26 “To configure an AA policer” (p. 2830) .

4

Create a policy sync group. See 87.34 “To configure a policy sync group” (p. 2846) .

5

Configure an AA flow watermark policy. See 87.39 “To configure an AA flow watermark” (p. 2852) .

6

Configure AA transit IP policies and transit prefix policies as required.

1. Configure AA transit IP policies. See 87.46 “To configure an AA transit IP policy” (p. 2860) .
2. Configure AA transit prefix policies. See 87.47 “To configure an AA transit prefix policy” (p. 2861) .
3. Configure database persisted transit subscriber aggregators. See 87.48 “To configure a database persisted transit subscriber aggregator” (p. 2862) .
4. Associate database persisted transit subscribers with aggregators. See 87.49 “To associate a database persisted transit subscriber with an aggregator” (p. 2863) .
5. As required, view the database persisted transit subscriber information. See 87.51 “To view database persisted transit subscriber information” (p. 2864) .

7

Configure usage-based billing as required:

1. Enable usage-based billing at the application profile level. See 87.52 “To configure usage-based billing for an application profile” (p. 2865) .
2. Associate applications with charging groups. See 87.53 “To associate an application with a charging group” (p. 2866) .

8

Configure HTTP policies as required:

1. Configure AA HTTP redirect policies. See 87.59 “To configure an AA HTTP redirect policy” (p. 2872) .
2. Configure AA HTTP error redirect policies. See 87.58 “To configure an AA HTTP error redirect policy” (p. 2871) .

9

Configure AA protocol signatures and enable inactive protocols, as required. See 87.70 “To enable an AA protocol signature” (p. 2884) .

10

As required, delete AA objects or records.

a. Delete an AA application, application group, or custom protocol; see 87.77 “To delete an AA application, application group, or custom protocol” (p. 2892) .

b. Remove the record of an inactive AA transit subscriber from the NFM-P database; see 87.78 “To delete an inactive AA transit subscriber instance” (p. 2893) .

87.22 Workflow to manage AA reporting

87.22.1 Purpose

This workflow provides general information about configuring AA reporting. See the *NSP Statistics Management Guide* for additional information about statistics collection, and the *NSP Analytics Report Catalog* for information about NSP Analytics Reports.

87.22.2 Stages

1

Create an AA accounting policy to specify when statistics are collected from the ISA-AA MDA. See 87.37 “To configure an AA accounting policy” (p. 2849) .

2

Configure application performance reporting on a service, and configure DCP groups for application performance reporting, if required. See 87.54 “To enable application performance reporting on a service” (p. 2867) .

3

Configure application performance reporting on SAPs and SDP bindings, as required. See 87.55 “To configure application performance reporting on a SAP or SDP binding” (p. 2868) .

4

Configure reporting attributes for AA reporting on a database persisted transit subscriber. See 87.56 “To configure application performance reporting for a transit subscriber” (p. 2869) .

5

As required, view AA results.

a. View AA summary information for subscribers, transit subscribers, SAPs, and spoke SDPs on ISA-AA MDAs; see 87.72 “To view AA summary information for an ISA-AA group or partition” (p. 2886) .

b. View AA special study statistics data on an ISA-AA MDA; see 87.75 “To view AA special study statistics data” (p. 2889) .

c. View application filter hit counts on local definitions of AA group policies; see 87.74 “To view AA statistics data for an ISA-AA group or partition” (p. 2888) .

d. View application filter hit counts on local definitions of AA group policies; see 87.76 “To view AA statistics data for application filters” (p. 2891) .

Application assurance procedures

87.23 To configure an AA group policy

Note: AQP creation for an AA group policy is described separately in 87.24 “To configure an AQP” (p. 2825) .

87.23.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Group Policy or select an AA group policy and click Properties. The AA Group Policy (Create | Edit) form opens.

3

Configure the parameters.

Note: You can set the Minimum ISA Generation parameter to 2 only if the ISA-AA group contains exclusively revision 2 ISA-AA MDAs.

4

Select a default charging group and default tethered charging group, as required.

Configure default Apdex thresholds

5

An Apdex threshold in an application or application group overrides a default threshold.

To modify a default Apdex threshold:

1. Click on the NSP Analytics Parameters tab, then on the Default Apdex Configuration tab.
2. Select an Apdex entry and click Create. The Apdex Configuration (Edit) form opens.
3. Click on the Thresholds tab.
4. Select a tier entry and click Properties. The Apdex Threshold (Edit) form opens.
5. Configure the parameters.

   A threshold value of -1 means that no value is specified.

   A frustrated threshold must be greater than the associated tolerated threshold, unless one or both values are inherited.
6. Save your changes and close the forms.

6

Repeat Step 5 to modify an additional threshold set, if required.

Create application groups

7

You can also manage application groups using the right-click contextual menu options for the Application Groups object on the AA Identification Components tab.

To create one or more application groups:

1. Click on the Application Groups tab.
2. Click Create. The Application Group (Create) form opens.
3. Configure the parameters.

   You must select the check box beside the Export ID parameter in order to configure that parameter.
4. Click on the Charging Group tab.
5. Select a charging group.
6. To configure Apdex thresholds for the application group, click on the NSP Analytics Parameters tab. Otherwise, go to 12 .

   An Apdex threshold in an application group overrides the associated default threshold configured in Step 5 .
7. Click Create. The Apdex Configuration (Create) form opens.
8. Configure the parameters.
9. Click on the Thresholds tab.
10. Select a tier entry and click Properties. The Apdex Threshold (Create) form opens.
11. Configure the parameters.

    A threshold value of -2 means that the value is inherited.

    A threshold value of -1 means that no value is specified.

    A frustrated threshold must be greater than the associated tolerated threshold, unless one or both values are inherited.
12. Save your changes and close the forms.

Create charging groups

8

To create one or more charging groups:

1. Click on the Charging Groups tab.
2. Click Create. The Charging Group (Create) form opens.
3. Configure the parameters.

You must select the check box beside the Export ID parameter before you can configure that parameter.

4. Save your changes and close the forms.

Create applications

9

You can also manage applications using the right-click contextual menu options for the Applications object on the AA Identification Components tab.

To create one or more applications:

1. Click on the Applications tab.
2. Click Create. The Application (Create) form opens.
3. Configure the parameters.

   You must select the check box beside the Export ID parameter before you can configure that parameter.
4. Click on the Application Group tab.
5. Select an application group.
6. Click on the Charging Group tab.
7. Select a charging group.
8. To configure Apdex thresholds for the application group, click on the NSP Analytics Parameters tab. Otherwise, go to 16 .

   An Apdex threshold in an application overrides the associated application group threshold configured in Step 7 .
9. Configure the Threshold Administrative State parameter on the General tab.
10. Click on the Apdex Configuration tab.
11. Click Create. The Apdex Configuration (Create) form opens.
12. Configure the parameters.
13. Click on the Thresholds tab.
14. Select a tier entry and click Properties. The Apdex Threshold (Create) form opens.
15. Configure the parameters.

    A threshold value of -2 means that the value is inherited.

    A threshold value of -1 means that no value is specified.

    A frustrated threshold must be greater than the associated tolerated threshold, unless one or both values are inherited.
16. Save your changes and close the forms.

Create application filters

10

You can also manage application filters using the right-click contextual menu options for the Application Filters object on the AA Identification Components tab.

To create one or more application filters:

1. Click on the Application Filters tab.
2. Click Create. The Application Filter (Create) form opens.
3. Configure the parameters.
4. Click on the General Properties tab.
5. Configure the Flow Set-up Direction and HTTP Match All Requests parameters.
6. Configure the parameters in the IP Protocol panel.
7. To specify an application protocol, configure the parameters in the Protocol panel and select a protocol.
8. You can specify server address properties or network address properties in an application filter, but not both.

   To specify server address properties, configure the Server Address Operator parameter. You can specify one server address range by configuring the Server Address and Server Address Mask parameters, or choose an IP prefix list to specify multiple server address ranges.

   Specify server address with netmask by configuring the Masked IP Address and Masked IP Address Mask parameters.You can specify server address properties or masked IP address properties in an application filter, but not both.

   Choose a DNS IP cache, if required, and go to Step 10 10 .
9. To specify network address properties, configure the Network Address Operator parameter. You can specify one network address range by configuring the Network Address and Network Address Mask parameters, or choose an IP prefix list to specify multiple network address ranges.
10. Configure the parameters in the Server Port panel.
11. Configure the parameters in the HTTP Port panel.
12. Click on the Application tab and choose an application.
13. To create one or more application filter expressions, click on the Application Filter Expressions tab and click Create. The Application Filter Expression (Create) form opens.
14. Configure the parameters.
15. Save your changes and close the forms.

Create ASOs

11

To create one or more ASOs:

1. Click on the Application Service Options tab.
2. Click Create. The Application Service Option (Create) form opens.
3. Configure the parameters.

   The Displayed Name of an ASO cannot contain a forward slash (/) character.
4. To create one or more ASO value entries, click on the Application Service Option Value Entries tab.
5. Click Create. The Application Service Option Value Entry (Create) form opens.
6. Configure the ASO Characteristic Value parameter.
7. Save your changes and close the form.

Create application profiles

12

To create an application profile:

1. Click on the Application Profiles tab.
2. Click Create. The Application Profile (Create) form opens.
3. Configure the parameters.
4. Click on the NSP Analytics Parameters tab.
5. Configure the Billing Reset Date parameter.
6. Configure the parameters.
7. Click on the Charging Group Thresholds tab.
8. Click Create. The Charging Group Thresholds (Create) form opens.
9. Configure the parameters.
10. Select a charging group.
11. Save your changes and close the form.
12. Click on the Characteristics tab to create one or more application profile characteristics.
13. Click Create. The Application Profile Characteristic (Create) form opens.
14. Choose an ASO characteristic.
15. Choose an ASO characteristic value.
16. Save your changes and close the forms.

Create AQPs

13

To configure one or more AQPs, perform 87.24 “To configure an AQP” (p. 2825) .

Create custom protocols

14

To create one or more custom protocols:

1. Click on the Custom Protocols tab.
2. Click Create. The Custom Protocol (Create) form opens.
3. Configure the parameters.
4. Click on the Custom Protocol Expressions tab.
5. Click Create. The Custom Protocol Expressions (Create) form opens.
6. Configure the parameters.
7. Save your changes and close the forms.

Configure default DCP group address rules

15

To configure one or more address rules for a default DCP group:

1. Click on the NSP Analytics Parameters tab, then on the Default DCP Groups tab.
2. Select a group and click Properties. The Default DCP Group (Edit) form opens.
3. Click on the Address Rules tab.
4. Click Create. The DCP Address Rule (Create) form opens.
5. Configure the parameters.
6. Save your changes and close the form.

Configure custom DCP groups

16

To configure one or more custom DCP groups:

1. Click on the NSP Analytics Parameters tab, then on the Custom Residential DCP Groups tab.
2. Configure the parameters.
3. Click on the Address Rules tab.
4. Click Create. The DCP Address Rule (Create) form opens.
5. Configure the parameters.

6. Save your changes and close the form.

Configure IP address-based application identification assistance

17

To enable AA to assign an application to a traffic session based on the IP addresses in the session:

1. Click on the IP Identification Assist tab.
2. Configure the parameters on the General tab.
3. Click on the DNS Trust Servers tab.
4. Click Create. The IP Identity Assist DNS Server (Create) form opens.
5. Configure the parameters.

18

Save your changes and close the forms.

END OF STEPS

87.24 To configure an AQP

87.24.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Select an AA group policy and click Properties. The AA Group Policy (Edit) form opens.

3

Click on the Application QoS Entries tab.

4

Click Create. The Application QoS Policy (Create) form opens.

5

Configure the parameters.

6

Click on the Match Criteria tab.

7

To configure General Attributes:

1. Configure the Traffic Direction parameter.
2. To specify an application as a criterion:
   a. Configure the Application Operator parameter.
   b. Choose an application.
3. To specify an application group as a criterion:
   a. Configure the Application Group Operator parameter.
   b. Choose an application group.
4. To specify DSCP criteria, configure the parameters in the DSCP panel.
5. To specify IP criteria, configure the parameters in the IP Protocol panel.
6. To specify a charging group as a criterion:
   a. Configure the Charging Group Operator parameter.
   b. Choose a charging group.

8

To configure subscriber criteria:

1. Click on the AA Subscriber tab.
2. Configure the parameters.

   The SAP and Spoke SDP Binding options for the Subscriber Type parameter are selectable only during local AQP entry creation.

   If the Subscriber Type parameter is set to None, the other parameters on the AA Subscriber tab are reset to the default values.
3. Perform one of the following.
   - If the Subscriber Type parameter is set to ESM, configure the ESM Subscriber parameter.
   - If the Subscriber Type parameter is set to SAP, configure the SAP Subscriber parameter.
   - If the Subscriber Type parameter is set to Spoke SDP Binding, configure the Spoke Subscriber parameter.
   - If the Subscriber Type parameter is set to Transit, configure the Transit Subscriber parameter.
   - If the Subscriber Type parameter is set to ESM-MAC, configure the ESM Subscriber Host parameter.

9

To configure source or destination criteria:

1. Click on the Source and Destination tab.
2. To configure source criteria, configure the Address Operator parameter in the Source Address panel. You can specify a single IP prefix or address by configuring the Address and Address Length parameters, or select an IP prefix list.

3. Configure the parameters in the Source Port panel.
4. To configure destination criteria, configure the Address Operator parameter in the Destination Address panel. You can specify a single IP prefix or address by configuring the Address and Address Length parameters, or select an IP prefix list.
5. Configure the parameters in the Destination Port panel.

10

To configure ASO characteristic criteria:

1. Click on the Characteristics tab.
2. Click Create. The Application QoS Policy Characteristic (Create) form opens.
3. Configure the parameters.
4. Save your changes and close the form.

11

Click on the Action tab and configure the Drop parameter. If you select the Drop parameter, go to Step 28 .

12

Configure flow attribute criteria if required.

To include flow attribute criteria in an AQP, the Flow Attributes parameter must be enabled in the ISA-AA group associated with the AA Group policy.

To configure flow attribute criteria:

1. Click on the Flow Attributes tab.
2. Click Create. The Application QoS Flow Attribute (Create) form opens.
3. Click Select to choose a flow attribute name.
4. Configure the parameters.
5. Save your changes and close the form.

13

Choose one or more agents in any or all of the following panels.

- Bandwidth Limit Policer
- Flow Rate Limit Policer
- Flow Count Limit Policer
- HTTP Error Redirect

14

To configure HTTP redirect:

1. Choose an HTTP redirect policy.

2. Configure the Flow Type parameter.

15

Choose an HTTP enrichment policy, if required.

16

Choose an HTTP notification policy, if required.

17

Choose a session filter, if required.

18

If required, configure the Subscriber Cut Through and Mirror Source All Inclusive parameters.

Note: The Mirror Source All Inclusive parameter is configurable only when you configure an AQP entry in a local definition of the AA group policy.

19

Choose a mirror service, if required.

Note: You can specify a mirror service only when you configure an AQP entry in a local definition of the AA group policy.

20

Choose a URL filter, if required.

21

Choose an ASO characteristic name, if required. The name is correlated to the ICAP Custom X-Header value in an AA URL filter.

22

Choose a DNS IP cache, if required.

23

Choose a GTP filter, if required.

24

Choose an SCTP filter, if required.

25

Configure the remarking parameters in the Remark panel.

26

Choose a Multi-Path TCP Proxy policy, if required.

27

If required, enable and configure the Value parameter in the TCP Maximum Segment Size panel.

28

Save your changes and close the forms.

END OF STEPS

87.25 To configure an AARP instance on an NE

87.25.1 Purpose

Application Assurance Redundancy Protocol (AARP) provides data-plane connectivity to dynamically maintain dual-homed AA subscriber traffic on the same ISA-AA for processing. An AARP instance is configured between dual-homed routers to establish connectivity with the same AARP instance number on each NE. See the *7750 SR Advanced Configuration Guide* and *7750 SR Multiservice Integrated Service Adapter Guide* for more information about AARP instances. See 78.18 “To add an AARP interface to an IES or a VPRN site” (p. 2448) for information about how to configure an AARP interface in an IES or VPRN service.

87.25.2 Steps

1

Choose Manage→ISA Functions→ISA-AA from the NFM-P main menu. The Manage ISA-AA form opens.

2

Click Create→AARP. The Select Network Elements form opens.

3

Select a Network Element and click OK. The AARP (Create) form opens.

4

Configure the required parameters.

a. If the AARP instance endpoint is a SAP:

1. Set the Peer Endpoint Type parameter to SAP.
2. If the SAP has a name, click on the L3 Access Interface tab and choose the SAP. Otherwise, configure the SAP Terminating Port, SAP Encapsulation Type, SAP Outer Encapsulation Value, and SAP Inner Encapsulation Value parameters.

The SAP Outer Encapsulation Value parameter is configurable when the SAP Encapsulation Type parameter is set to Dot1 Q:

The SAP Inner Encapsulation Value parameter is configurable when the SAP Encapsulation Type parameter is set to Q in Q:

b. If the AARP instance endpoint is a spoke SDP binding:
   1. Set the Peer Endpoint Type parameter to Spoke SDP Binding.
   2. If the spoke SDP binding has a name, click on Spoke SDP Bindings tab and choose the spoke SDP binding. Otherwise, configure the Spoke SDP Bind ID parameter.

5

Click Apply.

6

View the Operational Flag and Peer Operational Flag indicators to ensure that the configuration is correct.

7

Modify the configuration, if required.

8

Save your changes and close the forms.

END OF STEPS

87.26 To configure an AA policer

87.26.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Policer or select an AA policer and click Properties. The AA Policer (Create | Edit) form opens.

3

Configure the parameters.

4

Save your changes and close the forms.

END OF STEPS

87.27 To configure an AA GTP firewall

87.27.1 Purpose

Perform this procedure to configure a firewall for AA GTP traffic associated with an AQP that is bound to an ISA-AA group or partition.

87.27.2 Steps

Configure AA GTP filter

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA GTP Filter or select an AA GTP filter and click Properties. The AA GTP Filter (Create | Edit) form opens.

3

Configure the parameters.

4

Click on the Message Type tab.

5

Click Create, or select an entry and click Properties. The AA GTP Message Type (Create | Edit) form opens.

6

Configure the parameters.

7

Click OK to save your changes and close the form.

8

Click on the General tab.

9

Click Apply.

10

Distribute the policy to NEs, as required.

11

Close the AA GTP Filter (Create | Edit) form.

Configure AA SCTP filter

12

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

13

Click Create→AA SCTP Filter or select an AA SCTP filter and click Properties. The AA SCTP Filter (Create | Edit) form opens.

14

Configure the parameters.

15

Click on the PPID tab.

16

Click Create, or select an entry and click Properties. The AA SCTP Message Type (Create | Edit) form opens.

17

Configure the parameters; click Select or type a value to specify a PPID.

18

Click OK to save your changes and close the form.

19

Click on the General tab.

20

Click Apply.

21

Distribute the policy to NEs, as required.

22

Close the AA SCTP Filter (Create | Edit) form.

Bind GTP and SCTP filters to AQP

23

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

24

Choose AA Group Policy (Application Assurance) from the drop-down menu and click Select. A list of AA group policies is displayed.

25

Select a policy and click Properties. The AA Group Policy (Edit) form opens.

26

Click on the Application QoS Entries tab.

27

Select an entry and click Properties. The Application QoS Policy (Edit) form opens.

28

Click on the Action tab.

29

Use the Select buttons to specify the GTP and SCTP filters.

30

Click OK to save your changes and close the form.

31

Click Apply.

32

Distribute the policy to NEs, as required.

33

Close the AA Group Policy (Edit) form.

Configure GTP flow count monitoring in AA policer

34

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

35

Choose AA Policer (Application Assurance) from the drop-down menu and click Select. A list of AA policers is displayed.

36

Select a policer and click Properties. The AA Policer (Edit) form opens.

37

Configure the GTP Flow Count Limit parameter.

38

Click Apply.

39

Distribute the policer to NEs, as required.

40

Close the AA Policer (Edit) form.

41

Close the Application Assurance Policies form.

Configure firewall on AA group or partition

42

Right-click on an ISA-AA group in the Equipment view of the network navigation tree and choose Properties. The ISA-AA Group (Edit) form opens.

43

To configure the GTP firewall for the group, configure the parameters in the GTP panel.

You can configure a GTP firewall for an ISA-AA group only if the group is not partitioned.

44

To configure the GTP firewall for a partition:

1. Click on the ISA-AA partitions tab.
2. Select a partition and click Properties. The ISA-AA Group Partition (Edit) form opens.
3. Configure the parameters in the GTP panel.
4. Click OK to save your changes and close the form.

45

Click OK to save your changes and close the ISA-AA Group (Edit) form.

END OF STEPS

87.28 To configure an AA GTP-c firewall for S8 or Gn

87.28.1 Purpose

Perform this procedure to configure a firewall for AA GTP-c traffic associated with an S8 or Gn interface that is bound to to an ISA-AA group or partition.

87.28.2 Steps

Configure GTP parameters on the ISA-AA group and partition

1

Configure the GTPC Database parameter:

1. On the Equipment tree, expand *[NE]*→Logical Groups→ISA-AA Groups→ISA-AA Group *n*.
2. Right-click on the ISA-AA Group *n* icon and choose Properties. The ISA-AA Group (Edit) form opens.
3. On the General tab, in the Group panel, change the GTP Tunnel Database parameter value from 0 to 100 or vice versa.

2

If you are using 7750 SR NEs, configure the Minimum ISA Generation parameter value to 2.

This configuration is not required for VSR NEs.

3

Click on the ISA-AA Partitions tab.

4

Click Create or select an existing partition entry and click Properties. The ISA-AA Group Partition (Create|Edit) form opens.

5

On the General tab, in the GTP panel, configure the GTPC Inspection parameter value to Enabled.

6

Close the ISA-AA Group and ISA-AA Group Partition forms.

Configure AA GTP filter

7

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

8

Click Create→AA GTP Filter or select an AA GTP filter and click Properties. The AA GTP Filter (Create | Edit) form opens.

9

Configure the following parameters to Enabled:

- Validate GTP Tunnels
- Validate Sequence Number
- Validate Source IP Address

10

Click on the GTP in GTP tab and configure the default action for GTP in GTP packets.

11

Click on the GTP V2 Message Type tab and configure the default action for packets that do not match any GTP V2 message type entries.

12

Configure GTP V2 message type entries:

1. Click on the Entries subtab.
2. Click Create and configure the parameters.
3. Create additional GTP V2 message type entries as needed.

13

Click on the IMSI APN Filter tab and configure the default action for packets that do not match any IMSI APN filter entries.

14

Configure IMSI APN Filter entries:

1. Click on the Entries subtab.
2. Click Create and configure the parameters.
3. Create additional IMSI APN Filter entries as needed.

15

Distribute the policy to NEs, as required.

16

Close the AA GTP Filter (Create | Edit) form.

END OF STEPS

87.29 To configure AA TCP validation

87.29.1 Purpose

Perform this procedure to configure AA to monitor TCP packet discards on an ISA-AA group or partition for reasons such as packet corruption, malformation, or incorrect sequencing. The NFM-P can use the collected statistics to raise AA statistics TCAs.

87.29.2 Steps

Configure AA TCP validation policy

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA TCP Validation, or select an AA TCP validation policy and click Properties. The AA TCP Validation (Create | Edit) form opens.

3

Configure the parameters.

4

Click Apply.

5

Distribute the policy to NEs, as required.

6

Close the AA TCP Validation (Create | Edit) form.

Configure AA TCP validation TCA

7

Click Create→AA Statistics TCA→TCP Validation, or select a TCP validation AAsStatistics TCA

and click Properties. The AA Statistics TCA TCP Validation (Create | Edit) form opens.

8

Use the Select button to choose the new AA TCP validation policy.

9

Configure the parameters.

10

Click Apply.

11

Distribute the policy to NEs, as required.

Note: You can view the TCAs for an NE from the Faults tab of the local policy definition, which is listed on the Local Definitions tab of the AA TCP Validation form.

12

Close the AA Statistics TCA TCP Validation (Create | Edit) form.

Bind TCA to AQP entry

13

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

14

Choose AA Group Policy (Application Assurance) from the drop-down menu and click Select. A list of AA group policies is displayed.

15

Select a policy and click Properties. The AA Group Policy (Edit) form opens.

16

Click on the Application QoS Entries tab.

17

Select an entry and click Properties. The Application QoS Policy (Edit) form opens.

18

Click on the Action tab.

19

Use the Select button in the TCP Validation panel to specify the new TCP validation policy.

20

Click OK. The Application QoS Policy (Edit) form closes, and the AA Group Policy (Edit) form is again displayed.

21

Click Apply.

22

Click on the General tab..

23

Distribute the policy to NEs, as required.

24

Close the AA Group Policy (Edit) form.

Enable TCP validation statistics collection

25

Click Create→AA Accounting Policy, or select an AA accounting policy and click Properties. The AA Accounting Policy (Create | Edit) form opens.

26

If you are creating a policy, configure the parameters and click Apply.

27

Click on the Accounting Policies tab.

28

Select the AA Admit Deny entry and click Properties. The AA Accounting Configuration - AA Admit Deny (Create | Edit) form opens.

29

Select the Include TCP Validation Statistics parameter.

30

Click OK. The AA Accounting Configuration — AA Admit Deny (Create | Edit) form closes, and the AA Accounting Policy (Edit) form is again displayed.

31

Click on the General tab.

32

Distribute the accounting policy to NEs, as required.

Note: You can view the collected statistics for an NE from the Statistics tab of the local policy definition, which is listed on the Local Definitions tab of the AA TCP Validation form.

33

Close the AA Accounting Policy (Edit) form.

34

Close the Application Assurance Policies form.

END OF STEPS

87.30 To configure AA TCP optimization

87.30.1 Steps

Configure TCP optimization policy

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA TCP Optimization, or select an AA TCP optimization policy and click Properties. The AA TCP Optimization (Create | Edit) form opens.

3

Configure the parameters.

Note: To configure a dynamic initial slow start threshold, enter 1000001 for the initial slow start threshold parameter. This will set the initial slow start threshold to Auto on the NE.

4

Click Apply.

5

Distribute the policy to NEs, as required.

6

Click OK. The AA TCP Optimization (Create | Edit) form closes.

Bind policy to AA session filter

7

Click Create→AA Session Filter, or select an AA session filter and click Properties. The AA Session Filter (Create | Edit) form opens.

8

If you are creating an AA session filter, configure the parameters and click Apply.

9

Click on the Entries tab.

10

Click Create, or select a session filter and click Properties. The AA Session Filter Params (Create | Edit) form opens.

11

Set the Action parameter to TCP Optimizer.

12

Use the Select button in the TCP Optimization panel to specify the new TCP optimization policy.

13

Click OK. The AA Session Filter Params (Edit) form closes, and the AA Session Filter (Create | Edit) form is again displayed.

14

Click on the General tab.

15

Click Apply.

16

Distribute the policy to NEs, as required.

17

Close the AA Session Filter (Create | Edit) form.

18

Close the Application Assurance Policies form.

END OF STEPS

87.31 To renumber application filter entries

87.31.1 Purpose

The application filter entries in an AA group policy are applied to a traffic flow sequentially; the matching entry with the lowest ID value is applied first. Perform this procedure to assign new ID values to one or more application filter entries in an AA group policy.

Note: The NFM-P renumbers a filter entry ID only when the following are true.

- The new value is unused.
- The new value is less than or equal to the maximum allowed value.
- The renumbering does not change the filter order.

87.31.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Select an AA group policy and click Properties. The AA Group Policy (Edit) form opens.

3

Click on the Application Filters tab. The application filter entries are listed.

4

To renumber one filter entry:

1. Select the entry and click Renumber ID. The Renumber Entry ID form opens.
2. Configure the New Entry ID parameter.
3. Save your changes and close the form. The NFM-P renumbers the entry.

5

To renumber multiple filter entries:

1. Select the lowest-numbered entry that you want to renumber and click Renumber IDs. The Renumber Entry ID form opens.
2. Configure the Start Entry ID and End Entry ID parameters.

An End Entry ID value of 0 specifies that there is no ending entry ID; the range of entry IDs to renumber extends to the highest-numbered entry.

3. Perform one of the following:
   - Select Add Factor to add an offset to each entry ID in the specified range.
   - Select Multiply Factor to multiply the ID of the selected entry and each higher-numbered entry in the specified range by an offset.
4. Configure the Factor Value parameter by specifying the required addition or multiplication offset.
5. Save your changes and close the form. The NFM-P renumbers the entries.

6

Close the open forms.

END OF STEPS

87.32 To bind an application filter to an AA Port List Policy

87.32.1 Purpose

An AA Port List Policy can be bound to an application filter configured for the same ISA-AA Group and Partition.

87.32.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose AA Group Policy (Application Assurance) from the object drop-down menu.

3

Enter the ISA-AA Group ID and ISA-AA Partition ID in the fields and click Search. The AA group policies for the ISA-AA Group and Partition are listed.

4

Select an AA group policy and click Properties. The AA Group Policy (Edit) form opens.

5

Click on the Application Filters tab. The application filter entries are listed.

6

Click Create or select an application filter and click Properties. The Application Filter (Create | Edit) form opens.

7

Configure the parameters.

8

Click on the General Properties tab.

9

Configure the Server Port parameters:

1. Set the Server Port Value Type parameter to List.
2. Set the Server Port Operator parameter to Equal or Not Equal
3. Click Select for the Server Port List parameter and choose a policy from the list of released AA Port List (Application Assurance) Policies.

10

Configure the HTTP Port parameters:

1. Set the HTTP Port Value Type parameter to List.
2. Set the Server Port Operator parameter to Equal or Not Equal
3. Click Select for the Server Port List parameter and choose a policy from the list of released AA Port List (Application Assurance) Policies.

11

Close the open forms.

END OF STEPS

87.33 To configure an AA Cflowd group policy

87.33.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Cflowd Group Policy, or select an AA Cflowd group policy and click Properties. The AA Cflowd Policy (Create | Edit) form opens.

3

Configure the parameters.

When you configure the ISA-AA Partition ID parameter, the Partitions parameter is automatically set to Enabled.

4

Click Apply.

5

If the Partitions parameter is set to Enabled and you specified an ISA-AA Partition ID value, go to Step 14 .

6

Click on the Cflowd tab and configure the parameters.

7

Specify the Cflowd collectors.

1. Click on the Collector tab.
2. Click Create or select an existing collector entry and click Properties. The Cflowd Collector (Create | Edit) form opens.
3. Configure the parameters.
4. Save your changes and close the form.

   An ISA-AA group can contain a maximum of sixteen Cflowd collectors.

8

If the Partitions parameter on the General tab is set to Disabled, go to Step 9 .

9

To add an AA application or application group for Cflowd performance monitoring, click on the Performance tab. Otherwise, go to Step 12 .

10

To add an AA application for Cflowd performance monitoring, perform the following steps.

1. Click Add Applications Comprehensive, or Add Applications TCP. The Select Applications form opens.
2. Select one or more applications and click OK.

11

To add an AA application group for Cflowd performance monitoring, perform the following steps.

1. Click Add Application Groups Comprehensive, or Add Application Groups TCP. The Select Application Groups form opens.
2. Select one or more application groups and click OK.

12

To set the Cflowd sampling administrative states, click on the State tab and configure the required parameters.

13

To configure a template:

1. Click on the Template tab.
2. Select a template and click Properties.
3. In the form that opens, configure the General parameters.
4. Click on the Dynamic Fields tab and click Create.
5. In the Create form, click Select to choose a dynamic field and click OK to add it to the template.
6. Add additional fields as needed and click OK.

14

Save your changes and close the forms.

END OF STEPS

87.34 To configure a policy sync group

87.34.1 Steps

1

Choose Policies→Policy Sync Group from the NFM-P main menu. The Policy Sync Groups form opens.

2

Click Create→Policy Sync Group or select a policy sync group entry and click Properties. The Policy Sync Group (Create | Edit) form opens.

3

Configure the Policy Type parameter.

4

Use the Select button to choose a master policy.

5

Configure the remaining parameters.

6

Specify the AA group policy or AA Cflowd group policy classes to include in the policy sync group.

1. Click on the Included Classes tab. The tab lists the classes that are included in the policy sync group.

   The Sync Members, Add to Members, and Update Master actions performed on an AA group policy affect only the classes listed on the Included Classes tab.
2. To add one or more classes, click Add and specify the classes.
3. To delete one or more classes, select the classes and click Delete.

7

Add one or more member AA group or AA Cflowd group policies.

1. Click on the Members tab.
2. Click Add and select one or more policies. The Confirm form opens.

   The form displays the following options: “Listed master policy classes will overwrite member classes” and “Listed master policy classes will be added to member classes.”

   The Sync Members button performs the same function as the “Listed master policy classes will overwrite member classes” option.

   The Add to Members button performs the same function as the “Listed master policy classes will be added to member classes” option.
3. Choose an option and click OK.

8

To remove a member AA group policy:

1. Click on the Members tab.
2. Select a member policy.
3. Click Delete. The member policy is removed from the list.

9

Save your changes and close the forms.

END OF STEPS

87.35 To audit, compare, or synchronize policies using a policy sync group

87.35.1 Steps

1

Choose Policies→Policy Sync Group from the NFM-P main menu. The Policy Sync Groups form opens and lists the policy sync groups.

2

Select a policy sync group and click Properties. The Policy Sync Group (Edit) form opens.

3

To display the differences between the master policy and the policy members, click Audit.

Alarms are raised against member policies if mismatches in the included classes exist.

4

To compare two member policies, or compare the master policy to a member policy:

1. From the Members tab, select two policies and click Compare. The Compare - AA Group Policy or Compare - AA Cflowd Policy form opens.

   You can use the Swap button to switch Policy A and Policy B.
2. Click Compare to view the differences between Policy A and Policy B.
3. Choose a result and click Properties. The Difference - AA Group Policy (*class type*) or Difference - AA Cflowd Policy (*class type*) form opens.
4. View the differences between the two policies.

5

To add the classes defined in the master policy to each member policy, click Add To Members. The master policy classes are added to each member policy.

Note: After you click Add to Members, the member policy mode changes to Draft.

6

To add the contents of a selected AA group policy to the included classes of the master policy, click Update Master. The member policy classes are added to the master policy.

7

To synchronize each member policy with the master policy, click Sync Members. The master policy class list overwrites the class list in each member policy.

Note: After you click Sync Members, the member policy mode changes to Draft.

END OF STEPS

87.36 To configure an AA Access Network Location

87.36.1 Before you begin

An ISA-AA group must be configured; see 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415).

87.36.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create or select an AA Access Network Location entry and click Properties. The AA Access Network Location (Create | Edit) form opens.

3

Configure the required parameters.

4

If the Source Level parameter is set to Cell or mac-vlan, configure RTT Threshold parameters.

Note: The Source Level parameter cannot be changed after the policy is created.

5

Save your changes and close the forms.

END OF STEPS

87.37 To configure an AA accounting policy

Note: See “Statistics collection in the NFM-P” in the *NSP NFM-P Statistics Management Guide* for general information about configuring and collecting accounting statistics.

87.37.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create or select an AA accounting policy entry and click Properties. The AA Accounting Policy (Create | Edit) form opens.

3

Configure the parameters.

4

Click on the Accounting Policies tab. The accounting policies are listed.

5

To configure an accounting policy:

1. Select the accounting policy and click Properties. The AA Accounting Configuration form opens.
2. Configure the parameters.

   Enabling the Export Using parameter enables aggregate counter export using the method specified in the parameter options; for example, Accounting Policy.
3. Choose an accounting policy.
4. Configure the required policy-specific parameters.

   The Collect Usage Monitoring Statistics parameter is configurable only for a 7450 ESS or 7750 SR, and only when the accounting policy type is AA Subscriber Custom Record.
5. Click OK.

6

Save your changes and close the forms.

END OF STEPS

87.38 To configure AA accounting file export

87.38.1 Purpose

Perform this procedure to configure the export of AA accounting statistics to files on NFM-P main or auxiliary server stations.

CAUTION

Service Disruption

If you are enabling or disabling the export function, performing this procedure requires a server restart, which is service-affecting.

Perform this procedure only during a scheduled maintenance period.

Note: If the statistics are to be exported to a main server, and the NFM-P system is redundant, you must perform the procedure on each main server in the system, and must perform the procedure on the standby main server station first.

87.38.2 Steps

1

Log in to the main or auxiliary server station as the nsp user.

2

Open a console window.

3

Enter one of the following:

- On a main server station:
  `bash$ cd /opt/nsp/nfmp/server/nms/config ↵`
- On an auxiliary server station:
  `bash$ cd /opt/nsp/nfmp/auxserver/nms/config ↵`

4

Enter the following to make a backup copy of the server configuration file:

`bash$ cp nms-server.xml nms-server.xml.backup ↵`

5

Open the nms-server.xml file with a plain-text editor such as vi.

6

Locate the <aaStatsProcessing section.

7

To enable or disable the export function, set the following parameter to true or false, as required; for example, to enable the function:

`enabled="true"`

8

Specify the file format by editing the following line, as required:

`fileFormat="`*`format`*`"`

where *format* is one of the following:

- ipdr—IPDR format
- ram—format for NSP Analytics Reports
- ipdr, ram—IPDR and ram formats

9

Save and close the nms-server.xml file.

10

If you are configuring the standby main server in a redundant system, go to Step 12 .

11

If you changed only the fileFormat value, enter one of the following to put the configuration change into effect:

- On a main server station:

  bash$ **/opt/nsp/nfmp/server/nms/bin/nmsserver.bash read_config** ↵
- On an auxiliary server station:

  bash$ **/opt/nsp/nfmp/auxserver/nms/bin/auxnmsserver.bash auxread_config** ↵

The server reads the updated configuration, and saves subsequently collected AA accounting statistics to files on the server station.

12

If you are enabling or disabling the function, or are configuring the standby main server in a redundant system, enter one of the following to restart the server:

- On a main server station:

  bash$ **/opt/nsp/nfmp/server/nms/bin/nmsserver.bash force_restart** ↵
- On an auxiliary server station:

  bash$ **/opt/nsp/nfmp/auxserver/nms/bin/auxnmsserver.bash auxforce_restart** ↵

The server restarts, reads the updated configuration, and saves subsequently collected AA accounting statistics to files on the server station.

13

Close the console window.

END OF STEPS

87.39 To configure an AA flow watermark

87.39.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Select AA Flow Watermark (Application Assurance) from the object drop-down menu and click Search.

3

Select an AA flow watermark entry and click Properties. The AA Flow Watermark (Edit) form opens.

4

Configure the parameters.

5

Click Apply. The Configuration Mode changes to Draft.

6

Distribute the flow watermark to NEs, as required.

7

Close the forms.

END OF STEPS

87.40 To configure an AA RADIUS accounting policy

87.40.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create or select an AA RADIUS accounting policy entry and click Properties. The AA RADIUS accounting policy (Create | Edit) form opens.

3

Configure the parameters.

You must select the Enable check box beside the IP and Value (minutes) parameters before you can configure the parameters.

4

If you set the Router Instance parameter to VPRN, choose a VPRN service to associate the policy with.

5

Click on the RADIUS Accounting Servers tab and click Create. The AA RADIUS Accounting Server (Create) form opens.

6

Configure the parameters.

7

To add another RADIUS accounting server, click Apply and repeat Step 6 .

8

Save your changes and close the forms.

END OF STEPS

87.41 To configure an AA statistics TCA

87.41.1 Steps

1

Create an accounting policy and distribute the policy to the NEs that require the TCA configuration. See the *NSP NFM-P Statistics Management Guide* for information about creating accounting policies.

Note: You must set the accounting policy Type to AA Admit Deny.

2

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

3

Click Create→AA Accounting Policy, or select an AA accounting policy and click Properties. The AA Accounting Policy (Create | Edit) form opens.

4

Configure the parameters.

5

Click Apply.

6

Click on the Accounting Policies tab.

7

Select the AA Admit Deny entry and click Properties. The AA Accounting Configuration — AA Admit Deny (Edit) form opens.

Note: If the AA Admit Deny entry is not listed, you must synchronize the global policy with a local policy definition that has the AA Admit Deny policy type configured using a CLI.

8

Click Select and choose the previously configured AA Admit Deny accounting policy.

9

Click OK. The AA Accounting Configuration — AA Admit Deny (Edit) form closes.

10

Click OK. The AA Accounting Policy (Create | Edit) form closes.

11

Click Create→AA Statistics TCA Config, or select an AA statistics TCA and click Properties. The AA Statistics TCA Config (Create | Edit) form opens.

12

Configure the parameters.

13

Click OK to save your changes and close the AA Statistics TCA Config (Create | Edit) form.

14

Perform 87.42 “To configure an AA statistics TCA policer” (p. 2856).

15

Perform 87.43 “To configure an AA statistics TCA filter” (p. 2856).

16

Perform 87.44 “To configure an AA statistics TCA filter entry” (p. 2858).

17

Close the Application Assurance Policies form.

END OF STEPS

87.42 To configure an AA statistics TCA policer

87.42.1 Steps

1

Configure an AA policer; see 87.26 “To configure an AA policer” (p. 2830) for information.

Note: You must set the policer Type to Flow Rate Limit or Flow Count Limit.

2

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

3

Click Create→AA Statistics TCA Policer Config, or select an AA TCA policer and click Properties. The AA Statistics TCA Policer Config (Create | Edit) form opens.

4

Click Select and choose the previously created policer.

5

Configure the parameters.

6

Click OK to save your changes and close the AA Statistics TCA Policer Config (Create | Edit) form.

7

Close the Application Assurance Policies form, if required.

END OF STEPS

87.43 To configure an AA statistics TCA filter

87.43.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Statistics TCA Filter Config, or select an AA TCA filter and click Properties. The AA Statistics TCA Filter Config (Create | Edit) form opens.

3

Configure the TCA Config Filter Type parameter.

4

Click Select and choose a filter.

Note: Only filters of the type specified as the TCA Config Filter Type are listed.

5

Configure the Watermark Type parameter.

If the TCA Config Filter Type is set to GTP Filter, the valid options are the following:

- Payload Range
- Filter Default Action
- Header Sanity

Note: If you select Payload Range, the Max Payload Length parameter must be set on the GTP filter associated with the AQP.

If the TCA Config Filter Type is set to SCTP Filter, the valid options are the following:

- PPID Range
- Filter Default Action
- Packet Sanity

Note: If you select PPID Range, the PPID Range parameter must be set on the SCTP filter associated with the AQP.

If the TCA Config Filter Type is set to Session Filter, the valid option is Filter Default Action.

6

Configure the remaining parameters.

7

Click OK to save your changes and close the AA Statistics TCA Filter Config (Create | Edit) form.

8

Close the Application Assurance Policies form, if required.

END OF STEPS

87.44 To configure an AA statistics TCA filter entry

87.44.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Statistics TCA Filter Entry Config, or select an AA TCA filter entry and click Properties. The AA Statistics TCA Filter Entry Config (Create | Edit) form opens.

3

Configure the following parameters:

- TCA Config Filter Type
- ISA-AA Group ID
- ISA-AA Partition ID

4

Click Select and choose a filter entry.

Note: Only filter entries for the type specified as the TCA Config Filter Type are listed, and only if the filters contain at least one:

- message type, for GTP filters
- PPID, for SCTP filters
- protocol entry, for session filters

5

Configure the remaining parameters.

6

Click OK to save your changes and close the AA Statistics TCA Filter Config (Create | Edit) form.

7

Close the Application Assurance Policies form, if required.

END OF STEPS

87.45 To configure an AA Tether Detection policy

87.45.1 Purpose

An AA Tether Detection policy is created automatically when an ISA-AA group and its associated AA Group policy is created on a 7450 ESS NE. The Subcriber Scale parameter of the ISA-AA group must be configured to Mobile Gateway or another Gateway setting, depending on the NE release.

See 13.5 “To configure an ISA-AA group and ISA-AA partitions” (p. 415) to create an ISA-AA group. Use this procedure to complete the configuration.

87.45.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Select an AA Tether Detection policy and click Properties.

3

Configure the required parameters on the General tab.

4

Click on the TTL Monitoring tab.

5

Configure the required parameters.

6

Click on the Single Device Expected TTLs tab.

7

Click Create. The Expected TTL, Global Policy (Create) form opens.

8

Configure up to 16 TTLs.

9

Click Apply.

10

Distribute the policy to NEs, as required.

11

Close the forms.

END OF STEPS

87.46 To configure an AA transit IP policy

87.46.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Transit IP Policy or select an AA transit IP policy and click Properties. The AA Transit IP Policy (Create | Edit) form opens.

3

Configure the parameters.

4

Choose a default application profile.

Only application profiles with the same Group ID and Partition ID values as the AA transit IP policy are listed.

5

Choose a subscriber identification policy.

6

Configure the following required parameters.

The following parameters cannot be enabled at the same time:

- DHCP, and RADIUS or Seen IP
- Auto Create, and DHCP or RADIUS

7

If the RADIUS parameter is set to Enabled, select a subscriber authentication policy and an AA RADIUS accounting policy.

8

If the Diameter parameter is set to Enabled, select a diameter application policy.

9

To add one or more static subscribers to the AA transit IP policy:

1. Click on the Subscribers tab.
2. Click Create. The Transit Subscriber (Create) form opens.
3. Configure the Displayed Name parameter.
4. Choose an application profile.
5. Click on the IP Addresses tab.
6. Click Create. The Transit IP Address (Create) form opens.
7. Configure the Address parameter.
8. Save your changes and close the form.

10

Save your changes and close the forms.

END OF STEPS

87.47 To configure an AA transit prefix policy

87.47.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Transit Prefix Policy or select an AA Transit Prefix Policy entry and click Properties. The AA Transit Prefix Policy (Create | Edit) form opens.

3

Configure the parameters.

4

To configure one or more transit prefix subscribers:

1. Click on the Subscribers tab and click Create. The Transit Prefix Subscriber (Create) form opens.
2. Configure the parameters.
3. Choose an application profile.
4. Save your changes and close the form.

5

Note: Each entry ID must be assigned according to the following hierarchy, or the configuration fails:
local IPv4 < remote IPv4 < local IPv6 < remote IPv6
For example:

- Each local IPv6 entry ID must be higher than all IPv4 entry IDs.
- Local IPv4 entry IDs are lowest; remote IPv6 entry IDs are highest.

To create one or more transit prefix entries:

1. Click on the Entries tab and click Create. The Transit Prefix Entry (Create) form opens.
2. Choose a subscriber.
3. Configure the parameters.

   The Subscriber Address and Network Address values must both be IPv4 or IPv6; you cannot configure one address as IPv4 and the other as IPv6.
4. Save your changes and close the form.

6

Save your changes and close the forms.

END OF STEPS

87.48 To configure a database persisted transit subscriber aggregator

87.48.1 Steps

1

Choose Manage→AA Transit Subscribers from the NFM-P main menu. The Manage AA Transit Subscribers form opens.

2

Click Create→Transit Subscriber Aggregator or select a transit subscriber aggregator entry and click Properties. The Database Persisted Transit Subscriber Aggregator (Create | Edit) form opens.

3

Configure the parameters.

4

Save your changes and close the forms.

END OF STEPS

87.49 To associate a database persisted transit subscriber with an aggregator

Note: You can associate an aggregator only with database persisted business transit subscribers. A business transit subscriber is associated with an AA group that is configured for use in the VPN context.

87.49.1 Steps

1

Choose Manage→AA Transit Subscribers from the NFM-P main menu. The Manage AA Transit Subscribers form opens.

2

Choose a database persisted transit subscriber entry and click Properties. The Database Persisted Transit Subscriber (Edit) form opens.

3

Click on the NSP Analytics Parameters tab.

4

Click Select. The Select Database Persisted Transit Subscriber Aggregator form opens.

5

Configure the filter for the Keyword List column and click Search.

6

Select an aggregator and click OK.

7

Repeat Step 2 to Step 8 to associate more database persisted transit subscribers with aggregators, as required.

8

Save your changes and close the forms.

END OF STEPS

87.50 To configure trap throttling for AA transit subscriber creation and deletion

87.50.1 Purpose

Perform this procedure to specify the trap throttling rate on one or more NEs for AA transit subscriber creation and deletion events. When the trap rate exceeds the specified rate, the NFM-P resynchronizes the AA transit subscriber table on the NE to obtain the updated list of AA transit subscribers.

87.50.2 Steps

1

To enable the default rate limit of ten traps per second, create a CLI script that contains the following lines:

```
configure log event-control "application_
assurance" tmnxBsxTransIpPolAaSubCreated generate disable-specific-throttle
configure log event-control "application_
assurance" tmnxBsxTransIpPolAaSubDeleted generate disable-specific-throttle
```

2

To specify a different rate limit, create a CLI script that contains the following lines:

```
configure log event-control "application_
assurance" tmnxBsxTransIpPolAaSubCreated generate throttle specific-throttle-rate traps interval int

configure log event-control "application_
assurance" tmnxBsxTransIpPolAaSubDeleted generate throttle specific-throttle-rate traps interval int
```

where

traps is the number of traps

interval is the number of seconds over which the traps are generated

3

Use Workflows to execute workflows on one or more NEs, as required.

END OF STEPS

87.51 To view database persisted transit subscriber information

87.51.1 Steps

1

Choose Manage→AA Transit Subscribers from the NFM-P main menu. The Manage AA Transit Subscribers form opens.

2

Choose a database persisted transit subscriber entry and click Properties. The Database Persisted Transit Subscriber (Edit) form opens.

3

View the subscriber information, as required.

4

Close the Database Persisted Transit Subscriber (Edit) form.

5

Close the Manage AA Transit Subscribers form.

END OF STEPS

87.52 To configure usage-based billing for an application profile

87.52.1 Purpose

Perform this procedure to associate a charging group with an application profile. You can perform 87.53 “To associate an application with a charging group” (p. 2866) to assign an application to a charging group, if required.

87.52.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose an AA group policy entry and click Properties. The AA Group Policy (Edit) form opens.

3

Click on the Application Profiles tab.

4

Select an application profile and click Properties. The Application Profile (Edit) form opens.

5

Click on the NSP Analytics Parameters tab.

6

Configure the Billing Reset Date parameter.

7

Configure the parameters in the Apdex Thresholds and MOS Thresholds panels.

8

Click on the Charging Group Thresholds tab.

9

Click Create or select a charging group threshold and click Properties. The Charging Group Thresholds (Create | Edit) form opens.

10

Configure the parameters.

11

Select a charging group.

12

Save your changes and close the forms.

END OF STEPS

87.53 To associate an application with a charging group

87.53.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose an AA group policy entry and click Properties. The AA Group Policy (Edit) form opens.

3

Click on the Applications tab.

4

Select an application and click Properties. The Application (Edit) form opens.

5

Click on the Charging Group tab.

6

Choose a charging group.

7

Save your changes and close the forms.

END OF STEPS

87.54 To enable application performance reporting on a service

87.54.1 Purpose

Perform this procedure to enable application performance reporting on a service, and to configure default and custom DCP groups.

87.54.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The *service_type* Service (Edit) form opens.

3

Click on the Application Assurance tab.

4

Click on the NSP Analytics Parameters tab.

5

Select the Enable Application Performance Reporting parameter.

6

Click Apply. The NFM-P creates default IPv4 and IPv6 Internet and Intranet DCP subnet groups, which are listed on the Default DCP Groups tab.

7

To configure one or more address rules for a default DCP group:

1. Click on the Default DCP Groups tab.
2. Select a group and click Properties. The Default DCP Group (Edit) form opens.
3. Click on the Address Rules tab.

4. Click Create. The DCP Address Rule (Create) form opens.
5. Configure the parameters.
6. Save your changes and close the form.

8

To configure one or more custom DCP groups:

1. Click on the Custom DCP Groups tab.
2. Configure the parameters.
3. Click on the Address Rules tab.
4. Click Create. The DCP Address Rule (Create) form opens.
5. Configure the parameters.
6. Save your changes and close the form.

9

Save your changes and close the forms.

END OF STEPS

87.55 To configure application performance reporting on a SAP or SDP binding

87.55.1 Purpose

Perform this procedure to configure an Apdex tier value and the associated thresholds on a service object.

87.55.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The *service_type* Service (Edit) form opens.

3

Click on the Application Assurance tab.

4

Click on the NSP Analytics Parameters tab.

5

Select the Enable Application Performance Reporting parameter.

6

Click Apply.

7

Open the SAP or spoke SDP binding properties form.

8

Click on the Application Assurance tab.

9

Choose an application profile.

10

Click on the NSP Analytics Parameters tab.

11

Click Create. The AA Reporting (Create) form opens.

12

Configure the parameters.

13

Save your changes and close the forms.

END OF STEPS

87.56 To configure application performance reporting for a transit subscriber

87.56.1 Purpose

Perform this procedure to configure an Apdex tier value and the associated thresholds for a transit subscriber.

Note: You can perform this procedure only for a business transit subscriber. A business transit subscriber has an associated AA group policy with a Subscriber Scale setting of VPN.

87.56.2 Steps

1

Choose Manage→AA Transit Subscribers from the NFM-P main menu. The Manage AA Transit Subscribers form opens.

2

Choose a database persisted transit subscriber and click Properties. The Database Persisted Transit Subscriber (Edit) form opens.

3

Click on the NSP Analytics Parameters tab.

4

Click on the Reporting tab.

5

Click Create. The AA Reporting (Create) form opens.

6

Configure the parameters.

7

Save your changes and close the forms.

END OF STEPS

87.57 To disable application performance reporting on a service

87.57.1 Purpose

Perform this procedure to disable application performance reporting on a service and to delete the default DCP subnet groups.

87.57.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service and click Properties. The *service_type* Service (Edit) form opens.

3

Click on the Application Assurance tab.

4

Deselect the Enable Application Performance Reporting parameter.

5

Click Apply.

6

Click on the Default DCP Groups tab.

7

Click Delete All Result. The default DCP subnet groups are deleted.

8

Save your changes and close the forms.

END OF STEPS

87.58 To configure an AA HTTP error redirect policy

87.58.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA HTTP Error Redirect or select an AA HTTP error redirect policy entry and click Properties. The AA HTTP Error Redirect (Create | Edit) form opens.

3

Configure the parameters.

4

Choose a template.

5

To configure an HTTP error code:

1. Click on the HTTP Error Codes tab.
2. Click Create. The HTTP Redirect Error Code (Create) form opens.

3. Choose an error code.
4. Configure the Message Size (octets) parameter.
5. Save your changes and close the form.

6

Save your changes and close the forms.

END OF STEPS

87.59 To configure an AA HTTP redirect policy

87.59.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA HTTP Redirect or select an AA HTTP redirect policy entry and click Properties. The AA HTTP Redirect (Create | Edit) form opens.

3

Configure the parameters.

Note: The TCP Client Reset parameter is configurable only for a 7750 SR.

4

Configure the Redirect URL parameter to specify the location to which the subscriber hosts are redirected. If required, click Insert Macros to add keywords that specify an associated identifier such as the subscriber ID.

5

Choose a template.

6

Save your changes and close the forms.

END OF STEPS

87.60 To configure an AA HTTP Enrichment (Application Assurance) policy

87.60.1 Purpose

Perform this procedure to create an AA HTTP Enrichment (Application Assurance) policy for an AQP.

HTTPS enrichment is supported for connection to auto-authentication services with web portals that use HTTPS. If an HTTP TLS extension is configured, a customized extension is inserted in the message sent by the client to the server during the TLS handshake. The web portal should include an add-on that will remove the extension and perform integrity checks on the packet without the extension.

Note: HTTP enrichment policy restrictions may differ considerably between NE types and releases. To avoid deployment failures, Nokia recommends that you create a global HTTP enrichment policy for each combination of NE type and release in your network.

For example, create one global policy for each of the following:

Different NE types:	Different NE revisions:
• 7750 SR	• 7750 SR, Release *x* R*y* • 7750 SR, Release *x* R*z*

87.60.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA HTTP Enrichment or select an AA HTTP Enrichment (Application Assurance) policy entry and click Properties. The AA HTTP Enrichment (Create | Edit) form opens.

3

Configure the parameters.

4

Click on the HTTP Enrichment Fields tab.

5

Click Create. The HTTP Enrichment Field (Create) form opens.

6

Perform one of the following.

a. Choose a field.

b. Configure the Name parameter.

7

Configure the remaining parameters and select a certificate profile.

The Static String parameter is configurable when the Name parameter is set to static-string.

8

To create an HTTP Enrichment field with a certificate profile, select Certificate for the Encode Type parameter and choose a certificate profile.

Note: If the certificate profile is associated with an HTTP Enrichment field, the association must be removed before the certificate profile can be deleted.

9

Repeat Step 5 to Step 7 to add another field, if required.

10

To configure HTTPS enrichment, click on the HTTP Enrichment TLS Extension tab. The HTTP Enrichment TLS Extension (Create) form opens.

11

Click Select.

12

Choose one or more subtypes and click OK.

13

To configure HTTP RAT enrichment, click on the HTTP RAT Enrichment tab and perform the following to create RAT-type entries:

1. Click Create. The HTTP RAT Enrichment form opens.
2. Configure the RAT Type and RAT String, then click OK.

14

Save your changes and close the forms.

END OF STEPS

87.61 To configure an AA Certificate Profile

87.61.1 Purpose

Use this procedure to create a certificate for use in certificate based encryption. When certificate based encryption is configured for HTTP header encryption, the NE extracts the key from the certificate and uses it to perform encryption.

87.61.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Certificate Profile or select an AA Certificate Profile (Application Assurance) policy entry and click Properties. The AA Certificate Profile (Create | Edit) form opens.

3

Configure the parameters.

The certificate profile is added to the lists in the Application Assurance Policies form and the ISA AA Group (Edit) form.

4

Save your changes and close the forms.

END OF STEPS

87.62 To configure an HTTP notification policy

87.62.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA HTTP Notification or select an AA HTTP notification policy entry and click Properties. The AA HTTP Notification (Create | Edit) form opens.

3

Configure the parameters.

The Interval (minutes) parameter is configurable when the Interval Type parameter is set to Minimum Interval.

4

Choose a template.

5

Save your changes and close the forms.

END OF STEPS

87.63 To configure an AA Port List Policy

87.63.1 Purpose

Use this procedure to create an AA Port List Policy. A port list can contain a maximum of 32 port members. Members can be either an individual port or a range of ports.

87.63.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Port List Policy, or select an AA Port List policy and click Properties. The AA Port List Policy (Create | Edit) form opens.

3

Configure the parameters in the General tab.

4

Specify the port list members.

1. Click on the Port List Members tab.
2. Click Create or select an existing collector entry and click Properties. The AA Port List Member (Create | Edit) form opens.
3. Configure the parameters.
4. Save your changes and close the form.

5

Save your changes and close the forms.

END OF STEPS

87.64 To configure an AA IP prefix list policy

87.64.1 Purpose

Perform this procedure to create a policy that contains a list of IP prefixes. You can associate a prefix list with an IP session filter, application filter, or AQP.

87.64.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA IP Prefix List Policy or select an AA IP prefix list policy entry and click Properties. The AA IP Prefix List Policy (Create | Edit) form opens.

3

Configure the parameters.

4

Click on the IP Prefix tab.

5

Click Create. The AA Prefix Entry (Create) form opens.

6

Configure the parameters.

7

Save your changes and close the form.

8

Repeat Step 5 to Step 7 to create another prefix entry, if required.

END OF STEPS

87.65 To configure an AA Multi-path TCP policy

87.65.1 Purpose

Perform this procedure to create a policy for multi-path TCP scheduling. You can associate a prefix list with an AQP.

87.65.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Multi-Path TCP Policy or select an AA multi-path TCP policy entry and click Properties. The AA Multi-Path TCP Policy (Create | Edit) form opens.

3

Configure the parameters.

4

Save your changes and close the form.

END OF STEPS

87.66 To configure an AA session filter

87.66.1 Purpose

Perform this procedure to create an AA session filter for an AQP. See 87.24 “To configure an AQP” (p. 2825) for information about configuring an AQP.

87.66.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA Session Filter or select an AA session filter entry and click Properties. The AA Session Filter (Create | Edit) form opens.

3

Configure the parameters.

4

Click on the Entries tab and click Create. The AA Session Filter Params (Create) form opens.

5

Configure the parameters.

6

If you set the Action parameter to HTTP Redirect, choose an HTTP redirect policy.

7

Click on the Match Criteria tab and configure the Protocol Number parameter.

8

Perform one of the following in either or both of the Source Address and Destination Address panels.

a. Specify a single IP address or prefix by configuring the Address and Length parameters.

b. Choose an IP prefix list.

9

If required, choose a DNS IP cache in the Destination Address panel.

10

If the Protocol Number parameter is set to SCTP, TCP, or UDP, configure the parameters in the Source Port and Destination Port panels.

11

Save your changes and close the form.

12

Repeat Step 5 to Step 11 to add another filter entry, if required.

13

Save your changes and close the forms.

END OF STEPS

87.67 To configure and manage an AA URL list policy

87.67.1 Purpose

Perform this procedure to configure a URL list policy for an AA URL filter. An AA URL list is a list of URLs that is stored locally in a file on an NE.

87.67.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA URL List Policy, or select an AA URL list policy entry and click Properties. The AA URL List Policy (Create | Edit) form opens.

3

Configure the parameters.

Note: You can set the File Size parameter to Extended only when the Minimum ISA Generation value on the associated ISA-AA group is greater than 1.

4

Click Apply.

5

Distribute the policy to NEs, as required.

6

To check the status of a local policy instance, or remove a local decryption key:

1. Click on the Local Definitions tab.
2. Select a local instance and click Properties. The AA URL List Policy form (Edit) opens.
3. View the general status shown by the Operational Flags indicators on the General tab.
4. To remove the local decryption key, click Remove Decrypt Key.
5. Click on the Status tab.
6. Click Search, select an entry, and click Properties. The AA URL List Status Entry form opens. The form displays detailed, NE-specific status information.

7

Close the forms.

END OF STEPS

87.68 To configure an AA URL filter

87.68.1 Purpose

Perform this procedure to configure a URL filter and optional ICAP server or web service for an AQP.

Web services are supported on the following NEs: 7750 SR and 7450 ESS in mixed mode.

Note: To use a web service for AA URL filtering on an ISA-AA Group, the following parameters must be configured on the ISA-AA group:

- For 7450 ESS or 7750 SR NEs, the Minimum ISA Generation parameter should be set to 2.
  For VSR or VMG NEs, the parameter is set by default.
- The Web Service Cache Size parameter should be set to 100.

87.68.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA URL Filter or select a URL filter entry and click Properties. The AA URL Filter (Create | Edit) form opens.

3

Configure the parameters.

4

If the Default Action parameter is set to Block Http Redirect, choose an HTTP Redirect policy.

5

Choose an ICAP HTTP Redirect policy, if required.

6

If a web service will be used, configure a Web Server Classify ID and Category Set ID and click Apply.

7

To configure a local filtering allow list:

1. Click Select in the Local Filtering panel.
2. Click Search to populate the list of available URL List policies.
3. Select a list, or click Create and configure the parameters to create a URL list.
4. Click OK.

8

To configure a local filtering deny list:

1. Click on the Local Filtering Deny List tab.
2. Click Create. The AA URL Filter Deny List (Create) form opens.
3. Click Select to select a URL List policy, an HTTP Redirect policy, or both.

Notes:

The deny list cannot be the same as the allow list.

To block HTTP Redirect, the Default Action parameter must be set to Block Http Redirect.

4. Click OK.

9

To create an ICAP server entry:

1. Click on the ICAP Server tab.
2. Click Create. The AA ICAP Server Entries form opens.
3. Configure the parameters.

10

Repeat Step 9 to create an additional ICAP server entry.

11

To create a web service profile:

1. Click on the Web Service Profile tab and click Create.
2. Configure the General parameters.
3. Click on the Web Service Category sub-tab and click Create.
4. Click Select and select the URL categories to block.

12

Repeat Step 11 to create an additional web service profile if needed.

13

To configure a web service classification override:

1. Configure the Classifier ID parameter in the Web Service panel on the General tab.
2. Click on the Web Service Classification Overrides tab.
3. Click Create. the AA URL Filter Web Service Classification Override (Create) form opens.
4. Configure the parameters and click OK.

14

On the General tab, select a Web Service profile as the default profile.

15

Click Apply to save your changes.

16

Distribute the policy to NEs, as required.

17

Close the forms.

END OF STEPS

87.69 To configure an AA DNS IP cache

Note: You can specify an AA DNS IP cache in an application filter or AQP of an AA group policy.

87.69.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Click Create→AA DNS IP Cache, or select an AA DNS IP cache and click Properties. The DNS IP Cache (Create | Edit) form opens.

3

Configure the parameters.

4

Specify one or more cache servers.

1. Click on the AA DNS Cache Server tab and click Create. The AA DNS IP Cache Server (Create) form opens.
2. Configure the parameters.
3. Save your changes and close the form.

5

Specify one or more DNS domains and domain substrings.

1. Click on the AA DNS IP Cache Domain tab and click Create. The AA DNS IP Cache Domain (Create) form opens.
2. Configure the parameters.
3. Save your changes.

6

Distribute the policy to NEs, as required.

7

Close the forms.

END OF STEPS

87.70 To enable an AA protocol signature

87.70.1 Purpose

Perform this procedure to administratively enable an AA protocol signature for use by AA functions.

87.70.2 Steps

1

Choose Manage→ISA Functions→ISA-AA from the NFM-P main menu. The Manage ISA-AA form opens.

2

Select an AA protocol entry and click Properties. The AA Protocol (Edit) form opens.

3

Configure the Protocol Administrative State parameter.

You can configure the Protocol Administrative State parameter only for a local protocol.

4

Save your changes and close the form.

END OF STEPS

87.71 To update the AA application database on multiple NEs

87.71.1 Purpose

Perform this procedure to distribute an updated AA application database to multiple NEs using the NFM-P policy synchronization function.

87.71.2 Steps

1

Contact Nokia technical support to obtain the required application database delta file.

2

Using a CLI, execute the delta file on an NE.

Note: If the “Switch Distribution Mode to Local Edit Only on CLI Change” parameter on the System Preferences form is enabled, the distribution mode of each local policy changes to Local Edit Only. Otherwise, the distribution mode remains Sync With Global.

Each local policy on the NE is updated with the delta content, and differs from the global policy.

3

Perform an NE resync audit to ensure that the application database content on the NE and in the local NE policies is synchronized.

4

Perform the following steps to update the global policy from the local policy.

1. Navigate to the global policy by choosing Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.
2. Select AA Group Policy (Application Assurance) from the object drop-down menu and click Search.
3. Select the required global AA group policy and click Synchronize. The Synchronize form opens and the global policy is listed in the Destination Policy Instance(s) panel.
4. In the Available Local Policies list, select the NE that has the updated local policy with the delta content and click the right-pointing arrow. The NE moves to the Selected Source Local Policy panel.
5. Click Synchronize. The NFM-P synchronizes the global policy with the local policy, and the global policy distribution mode changes to Draft.
6. Close the Synchronize form.

5

Perform the following steps to distribute the global policy to one or more NEs.

1. Select the global policy in the Application Assurance Policies form and click Properties. The AA Group Policy (Edit) form opens.
2. Click Switch Mode. The Release - AA Group Policy form opens with the associated NEs listed in the Select Objects panel.
3. Click Distribute. The NFM-P distributes the policy to the NEs.
4. Close the forms.

Note: If a local policy is in Local Edit Only distribution mode, the release has no effect on the local policy; if the distribution mode of a local policy is Sync With Global, the global policy overwrites the local definition.

6

Perform the following steps to switch any local policies in Local Edit Only mode to Sync With Global mode.

1. Choose Local from the Policy scope drop-down menu on the Application Assurance Policies form. The local policies are listed.
2. Select the required policies and click Switch Distribution Mode. The Distribute form opens.
3. Choose Local Edit Only from the drop-down menu. The NEs that have policies in Local Edit Only mode are listed in the Available Nodes with Local Policies panel.
4. Select the NEs in the Available Nodes with Local Policies panel and click the right-pointing arrow. The NEs move to the Selected Nodes with Local Policies panel.
5. Click Synch With Global. The distribution mode of each policy changes to Synch With Global.

7

Close the forms.

END OF STEPS

87.72 To view AA summary information for an ISA-AA group or partition

87.72.1 Purpose

CAUTION

Service Disruption

The AA summary information on the NE may be slightly out of synchronization with the NFM-P database. Before you attempt to resynchronize an ISA-AA group, Nokia recommends that you consider the effect of resynchronizing an ISA-AA group; such an operation may require the retrieval of a large volume of information, and may affect NE performance.

Perform this procedure to view the AA summary information for the following objects associated with an ISA-AA group or partition:

- subscribers
- transit subscribers
- SAPs
- spoke SDP bindings

AA summary information is typically used for AA debugging purposes. You can view debug statistics information on the Statistics tab of an AA object.

87.72.2 Steps

1

In the equipment view navigation tree, expand Network→NE→Logical Groups→ISA-AA Group→ISA-AA Group.

2

Right-click on the ISA-AA Group icon and choose Properties. The ISA-AA Group (Edit) form opens.

3

If partitions are enabled in the ISA-AA group:

1. Click on the Partitions tab. The AA group partitions are listed.
2. Select a partition and click Properties. The ISA-AA Group Partition (Edit) form opens.

4

To view the summary information for an object:

1. Click on the AA Summary tab. The following tabs are displayed; each tab lists the corresponding objects in the ISA-AA group: Subscribers, SAPs, Spoke SDP Bindings, and Transit Subscribers.
2. Click on the appropriate tab.
3. Select an object and click Properties. The AA *object_type* (Edit) form opens.
4. View the information.

5

Close the forms.

END OF STEPS

87.73 To configure subscriber usage monitoring

87.73.1 Purpose

Perform this procedure to configure usage monitoring for a subscriber instance associated with an AA Subscriber Custom Record accounting policy.

87.73.2 Steps

1

Configure the Collect Usage Monitoring Statistics parameter in the AA Subscriber Custom Record accounting policy, as described in Step 5 of 87.37 “To configure an AA accounting policy” (p. 2849) .

2

Distribute the accounting policy to NEs, as required.

3

To configure the export method for a local policy instance:

1. Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.
2. Select AA Group Policy from the object drop-down menu and click Search.
3. Select a policy and click Properties. The AA Group Policy (Edit) form opens.
4. Click on the Local Definitions tab.
5. Select an instance and click Properties. The properties form of the local policy instance opens.
6. Click on one of the following tabs, as required: Application, Application Group, Charging Group.
7. Select an entry and click Properties. The properties form of the entry opens.
8. Click on the AA Subscriber Stats Objects tab.
9. Select a statistics object and click Properties. The AA Subscriber Stats Object Config form opens.
10. Deselect the Enable Selection parameter.
11. Save your changes and close the form.

END OF STEPS

87.74 To view AA statistics data for an ISA-AA group or partition

87.74.1 Steps

1

In the equipment view, expand Network→NE→Logical Groups→ISA-AA Groups-ISA-AA Group.

2

Right-click on the ISA-AA Group icon and choose Properties. The ISA-AA Group (Edit) form opens.

3

Perform one of the following.

a. View AA statistics data for an ISA-AA group.
   1. Click on the Statistics tab.
   2. Choose an AA statistics class from the object drop-down menu.

b. View AA statistics data for an ISA-AA partition.
   1. Click on the ISA-AA Partitions tab. The ISA-AA partitions are listed.
   2. Select an entry and click Properties. The ISA-AA Group Partition (Edit) form opens.
   3. Click on the Statistics tab.
   4. Choose an AA statistics class from the object drop-down menu.

4

Perform one of the following.

a. Click Collect to perform an on-demand collection of the current performance statistics data. The collected statistics entries are listed on the form.

b. Click Collect All to collect one on-demand statistics record for each statistic type that the object supports. The collected statistics entries are listed on the form.

5

Select an entry and click Properties. The Statistics Record form opens.

6

View the statistics data.

7

Close the forms.

END OF STEPS

87.75 To view AA special study statistics data

87.75.1 Purpose

Perform this procedure to view the AA special study statistics data on an ISA-AA MDA. If the ISA-AA group is partitioned, the AA special study tabs are on the partition properties form.

87.75.2 Steps

1

In the equipment view, expand Network→NE→Logical Groups→ISA-AA Groups-ISA-AA Group.

2

Right-click on the ISA-AA Group icon and choose Properties. The ISA-AA Group (Edit) form opens.

3

Click on the AA Special Study tab. The following tabs are displayed; each tab lists the corresponding objects in the ISA-AA group:

- Subscribers
- SAPs
- Spoke SDP Bindings
- Transit Subscribers
- ESM Subscriber Host

4

Click on the appropriate tab. A list of objects is displayed.

5

Select an object and click Properties. The AA *object_type* Config (Edit) form opens.

6

Click on the Statistics tab.

7

Choose an AA statistics class.

8

Perform one of the following.

a. Click Collect to perform an on-demand collection of the current performance statistics data. The collected statistics entries are listed on the form.

b. Click Collect All to collect one on-demand statistics record for each statistic type that the object supports. The collected statistics entries are listed on the form.

9

Select an entry and click Properties. The Statistics Record form opens.

10

View the statistics data.

11

Close the forms.

END OF STEPS

87.76 To view AA statistics data for application filters

87.76.1 Purpose

Perform this procedure to view the application filter hit counts for a local AA group policy.

87.76.2 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose AA Group Policy (Application Assurance) from the object drop-down menu.

3

Click Search. The AA group policies are listed.

4

Select an AA group policy and click Properties. The AA Group Policy - Global Policy (Edit) form opens.

5

Click on the Local Definitions tab.

6

Click Search. The local policy instances are listed.

7

Select a policy and click Properties. The AA Group Policy - Local Policy (Edit) form opens.

8

Click on the Application Filters tab.

9

Click Search. The application filters are listed.

10

Select an application filter and click Properties. The Application Filter - Local Policy (Edit) form opens.

11

Click on the Statistics tab and click Collect All.

12

Click Search. The collected statistics records are listed.

13

Select a statistics object and click Properties. The Statistics Record - AA Application Filter Stats form opens.

14

View the statistics data.

15

Close the forms.

END OF STEPS

87.77 To delete an AA application, application group, or custom protocol

87.77.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose AA Group Policy (Application Assurance) from the object drop-down menu and click Search. The AA group policies are listed.

Note: The NFM-P does not create or delete a local AA group policy on each applicable managed NE. A local AA group policy is created or deleted when the ISA-AA group is created or deleted.

3

Select a policy and click Properties. The AA Group Policy (Edit) form opens.

4

Click on the Applications, Application Groups, or Custom Protocols tab, as required.

5

Select an entry and click Delete.

6

Click View Dependencies to review the dependencies, if required.

When an entry is associated with a global AA group policy, the dependency count includes all dependencies on the global and local policy instances. When an entry is in a local AA group policy, the counts include the dependencies on the local policy instance only.

7

Save your changes and close the form.

8

Distribute the updated policy to each NE that has a local definition of the policy.

END OF STEPS

87.78 To delete an inactive AA transit subscriber instance

87.78.1 Purpose

Perform this procedure to remove the record of an inactive AA transit subscriber from the NFM-P database. An AA transit subscriber instance become inactive when the subscriber is deleted from an AA transit IP policy.

Note: You can use an ageout constraint policy to ensure that the number of inactive AA transit subscriber instances in the NFM-P database remains at a manageable level. See the *NSP System Administrator Guide* for information about ageout constraint policies.

AA statistics data can be lost when you delete an AA transit subscriber instance before the NFM-P finishes collecting and processing the AA statistics. Nokia recommends that you wait at least two statistics collection intervals before you delete an inactive AA transit subscriber instance.

87.78.2 Steps

1

Choose Manage→AA Transit Subscribers from the NFM-P main menu. The Manage AA Transit Subscribers form opens.

2

Click Search. The AA transit subscriber instances are listed.

3

Click on the Active column header to sort the list. An inactive AA transit subscriber instance has no check mark in the Active column.

4

Select the AA transit subscriber instance to delete and click Delete. The AA transit subscriber instance is removed from the list.

5

Close the Manage AA Transit Subscribers form.

END OF STEPS

88 Tunnel administrative groups

88.1 Overview

88.1.1 Purpose

The NFM-P supports SDP admin groups—or steering parameters—on all NEs and on all service tunnel types. However, steering parameters are NFM-P objects that are not deployed to NEs. The 7750 SR supports the use of SDP admin groups, which enable services to use PW templates for automatic inclusion or exclusion of specific SDPs.

See the 7750 SR documentation for information about SDP admin groups

In the context of NFM-P management, steering parameters are called tunnel administrative groups. A tunnel administrative group uses steering parameter and SDP admin group concepts.

Tunnel administrative groups use the NFM-P policy distribution model. You can create a global tunnel administrative group and distribute it locally to NEs. Consider the following when you create a tunnel administrative group:

- you can create a maximum of 32 tunnel administrative groups
- the group value is uniquely associated with the group name
- an SDP can be added to more than one tunnel administrative group
- tunnel administrative groups are supported on:
  - GRE service tunnels
  - MPLS:BGP service tunnels
  - MPLS:LDP service tunnels
  - RSVP-LSP service tunnels
  - mixed LSP mode service tunnels
- you cannot assign tunnel administrative groups to an SDP using an LDP IPv6 FEC

SDPs that share a specific characteristic or attribute can be made members of the same tunnel administrative group. When you create a service PW template, you can include and exclude one or more tunnel administrative groups. When a service is bound to the template, the SDP selection rules enforce the specified tunnel administrative group inclusion and exclusion constraints. Tunnel administrative groups allow you to control the SDP, or set of SDPs, that are selected while spokes are established among service sites using PW templates in BGP-AD, BGP-VPLS, or MS-PW applications. Tunnel administrative groups can be assigned to services that use the PW template policy, such as BGP-AD VPLS services, BGP-VPLS services, and Epipe spoke SDP FEC services. See Chapter 83, “Service PW template policies” for more information about PW template policies.

In addition, you can apply tunnel administrative groups to tunnel selection profiles. Tunnel selection profiles assign transport tunnels for a service when the service has been configured for automatic SDP binding creation. When a tunnel selection profile is used within a service to create SDP bindings, any tunnels that include the tunnel administrative group in that profile become eligible for consideration in the tunnel selection process. See Chapter 33, “Service tunnels” for more information about tunnel selection profiles.

88.1.2 Contents

Tunnel administrative group procedures

88.2 Workflow to configure tunnel administrative groups

88.2.1 Overview

The following workflow describes the high-level tasks required to create and configure tunnel administrative groups.

88.2.2 Stages

1

Create a tunnel administrative group. See 88.3 “To create a tunnel administrative group” (p. 2897) .

2

Release and distribute the tunnel administrative group to the NEs. See 88.3 “To create a tunnel administrative group” (p. 2897) .

3

Assign the tunnel administrative group to one or more of the following:

- A service tunnel. See 33.9 “To create an IP/MPLS service tunnel” (p. 1190) in Chapter 33, “Service tunnels” .
- A tunnel selection profile. See 33.13 “To create a tunnel selection profile” (p. 1201) in Chapter 33, “Service tunnels” .
- A service PW template. Ensure that the Use Provisioned SDP parameter on the service PW template is enabled. See 83.3 “To configure a PW template policy” (p. 2744) in Chapter 83, “Service PW template policies” .
- A BGP or BGP AD VPLS. See 77.113 “To assign tunnel administrative groups to a BGP or BGP AD VPLS” (p. 2418) in Chapter 77, “VPLS management ” .

88.3 To create a tunnel administrative group

88.3.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Click Create→Tunnel Admin Group. The Tunnel Admin Group, Global Policy (Create) form opens.

3

Configure the required general parameters and click Apply.

4

Click Switch Mode to release the policy for distribution. The Release - Tunnel Admin Group form opens.

5

Choose the NEs in the Available Objects to which you want to distribute the tunnel admin group and click on the arrow button to move the NEs to the Selected Objects panel.

6

Click Distribute and close the form.

7

Save the changes and close the forms.

END OF STEPS

88.4 To list and view tunnel administrative groups

88.4.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Choose Tunnel Admin Group (Service Tunnel Management) from the object drop-down menu and click Search.

3

Choose a tunnel administrative group and click Properties. The Tunnel Admin Group - Global Policy (Edit) form opens.

4

Click on the following tabs for additional information:

- Local Definitions
- Service Tunnels—lists the service tunnels to which this tunnel administrative group is assigned, based on ID only
- Tunnel Selection Profiles— lists the tunnel selection profiles that include or exclude this tunnel administrative group, based on ID only

- Service PW Templates— lists the service PW templates that include or exclude this tunnel administrative group, based on ID only
- Faults—lists the alarms raised for the policy

5

Close the forms.

END OF STEPS

Part VII: Service assurance

Overview

Purpose

This part provides information about service assurance using the NFM-P.

Contents

89 Service Test Manager

89.1 Overview

89.1.1 Purpose

This chapter describes the NFM-P Service Test Manager, and provides information about configuring and using the STM.

89.1.2 Contents

Service Test Manager description

89.2 STM concepts and components

89.2.1 OAM diagnostic tests

The NFM-P Service Test Manager (STM) provides access to a set of configurable in-band or out-of-band, packet-based OAM diagnostic tests that allow you to perform on-demand or scheduled verifications of SLA compliance and for network troubleshooting.

You can perform the following tasks with the STM:

- Create on-demand OAM diagnostic tests to provide proactive detection of service degradation and SLA verification and to verify end-to-end service performance. See Chapter 90, “OAM diagnostic tests”for information about each OAM diagnostic test that can be configured with the NFM-P STM.
- Run discrete OAM diagnostic tests on concurrent groups of objects.
- Group OAM tests in an integrated STM test suite for concurrent execution to provide continual performance feedback such as latency, delay, packet loss, and threshold-crossing alerts. The test results are logged for monitoring and trend analysis.
- STM test suites can be immediately processed, scheduled for the automatic execution of tasks using NFM-P-based schedules at designated times, or retained for future use. An NFM-P schedule is configurable for one-time or ongoing task execution.
- Configure threshold-crossing parameters to generate alarms when rising or falling threshold values are reached due to the reach, latency, or jitter issues discovered by the OAM diagnostic tests.
- Automatically generate OAM tests based on object or topology changes.

Note: You can customize some STM-specific user and system preferences to change the default STM settings to meet your operational requirements. See the workflow in 89.9 “STM workflow” (p. 2945) for additional information.

Note: The NFM-P STM allows the deployment of the maximum number of tests to a 7450 ESS, 7750 SR, or 7950 XRS. The NFM-P raises an alarm when the number of tests on an NE is 60% of the configured maximum. Attempts to create or execute a test using the NFM-P fail when the number of deployed tests on an NE is too high.

Note: OAM tests on 7705 SAR-Hm NEs will not execute if the packet size is greater than 1600 octets.

89.2.2 STM test policies

To enable the automatic generation of tests within an STM test suite, the NFM-P requires that the STM test policy contains a set of test definitions that defines the pre- and post-processing rules. An STM test policy also specifies the order of execution for the generated tests. An STM test policy is

applied to an STM test suite during test suite creation. See 89.10 “To configure an STM test policy” (p. 2947) for information about how to create an STM test policy.

You can re-apply an STM test policy under the following conditions:

- by clicking Generate on the Test Policy (Create|Edit) form when any test entities have changed
- by clicking Update Test Suites on the Test Policy (Create|Edit) form

An STM test policy can be used by multiple STM test suites but an STM test suite can have only one associated STM test policy.

Note: You can apply changes made to an STM test policy to all referencing STM test suites by clicking Update Test Suites on the Test Policy (Create|Edit) form.

STM test policy parameters can be configured to only display test results if a test fails or generates a threshold-crossing alarm. In large networks, this can substantially reduce the amount of test data that the NFM-P needs to collect.

An STM test policy is specific to one type of entity; for example, a VLL service or service tunnel. The test definitions in the policy are restricted to the tests that apply to the entity type specified in the policy.

89.2.3 STM test suites

The grouping of OAM diagnostic tests into an STM test suite allows an NFM-P operator to use one schedule for the periodic execution of multiple tests against multiple network objects; for example, services, NEs, or transport components.

An operator can choose to include existing tests, use the NFM-P to generate the OAM diagnostic tests that comprise an STM test suite, or both. Groups of tests in a suite can be configured to execute sequentially or concurrently. In addition, you can configure an STM test suite as an OAM validator to verify the operational status of a service. This can also be done on a one-time basis using the one-time service validation test; see “OAM diagnostic tests” (p. 2980) for more information.

A test suite contains three test groups:

- **First-run tests**

  First-run tests are the tests in a suite that the NFM-P executes before the tests in the other groups. First-run tests are chosen from a list of existing tests and might typically include high-level diagnostics; for example, a service site ping or VPRN ping. No restrictions apply to the types of tests that are selectable as first-run tests.
- **Generated tests**

  Generated tests are created by the NFM-P for use against a specific network entity, based on the entity type specified in the suite and the specific tested entities that are named in the associated test policy. For example, a service site ping test policy associated with a three-site VPRN test suite causes the NFM-P to generate six tests: one site ping test from each site in the VPRN to the other two sites. When you change the configuration of a network entity, such as a service, you must regenerate the generated tests that apply to the entity. Test regeneration removes previously generated tests from a test suite.
- **Last-run tests**

Last-run tests are the tests in a suite that the NFM-P executes after the tests in the other groups. Last-run tests are chosen from a list of existing tests and might typically include transport-layer diagnostics; for example, an LSP trace or a tunnel ping. No restrictions apply to the types of tests that are selectable as last-run tests.

Note: First-run and last-run tests do not apply to STM test suites where the MEF35 Mode parameter is enabled. See 89.12 “To create an STM test suite” (p. 2951) for more information.

To create an STM test suite that contains tests for different entity types, you can specify that the test suite applies to no specific entity type. In this way, you can create a group of disparate tests to which no test policy restrictions apply. Specifying None as the entity type in a test suite has the following effects:

- It allows you to choose any predefined test as a first-run or last-run test.
- It disables test generation in the test suite, because test generation requires a test policy that is based on a specific entity type.

Note: The NFM-P does not attempt to discover tests or test suites that are configured locally on an NE, for example, using a CLI.

Note: By default, the NFM-P suppresses alarms for suspended NEs. See 9.3.3 “Suspending device management” (p. 280) for more information. If you attempt to modify an OAM test or test suite that is deployed to a suspended NE, the NFM-P raises an alarm.

To manage the system resources that test execution consumes, the NFM-P assigns a weight value to a test. When the NFM-P executes a test, it attempts to reserve the test weight from a resource pool, performs the test, then returns the test weight to the pool. The weight of a test suite is the sum of the weights of the individual tests in the suite. The NFM-P attempts to reserve the weight of the whole suite for the duration of suite execution. If the required weight for a test or test suite is unavailable, execution is halted and the Status value contained in the test result is set to Not Enough Resources.

You can create an OAM service validation test to verify the operational status of a service. The operational status of a service depends on the operational states of its service sites or instances. It is possible for a service to be operationally up when communication between sites is not operational. For example, a VRF can be operationally up but the routes to its peers might not be populated because of the routing policy, route target, or ACL configuration. The State Cause of the service indicates the success or failure of the OAM validation test, and therefore, service connectivity.

You can configure an OAM validator when you create a test suite and run the OAM validation test from the service configuration form.

Note: OAM validation tests are not supported for HVPLS.

89.2.4 STM test suite design considerations

Consider the following when you create, schedule, or run STM test suites.

- **General design considerations:**
  - A test, whether pre-existing or generated, is associated with only one STM test suite.

- You can execute a generated test on demand, not just in the context of a test suite.
- An NFM-P user who is assigned the admin or QoS/ACL management scope of command role can create and modify all tests, test policies, and test suites. A user who is assigned the service management scope of command role can create and modify only STM components that are related to services. A user who is assigned the topology management role can create and modify only STM components that are related to network transport elements.

- **STM Test Suite (Create) form parameter considerations**
  - OAM test suites in which the Validation Test Suite parameter is enabled are used to test the operational status of a service or service-related entity such as a service tunnel. The result of this validator test is indicated by the OAM Validation Failed state cause indicator on the General tab of the management form for the object.
  - The Entity Type parameter specifies that only test policies created for the same entity type are available for the test suite. The parameter also restricts the predefined tests that are available as first-run or last-run tests to those that apply to the entity type.
  - To create a test suite that includes tests for different entity types, specify None as the Entity Type parameter value. This allows you to choose any predefined test as a first-run or last-run test.
  - Specifying None as the Entity Type parameter value disables the generation of tests in a test suite; the policy the STM uses to generate tests must be associated with a specific entity type.
  - The Validation Test Suite parameter specifies that the test suite is used to validate the connectivity of the tested service entity to which it is applied. This is also referred to as a validator test suite. OAM validation tests are not supported for HVPLS. See 90.55 “To run a one-time validation test on a service” (p. 3084) for information about how to run a one-time validation test on a service. See 33.26 “To run an OAM validation test on a service tunnel” (p. 1225) for information about how to run a one-time validation test on a service tunnel.
  - If you are creating a test suite that contains an Ethernet CFM test definition, you must enable the NE Schedulable parameter.
- **Scheduling and execution considerations:**
  - The execution of STM test suites can be configured as scheduled tasks. See Chapter 5, “NFM-P-based schedules” for additional design consideration information about using the NFM-P to schedule STM test suites.
  - Using an NE schedule does not ensure that the target NEs perform the actions in an associated STM scheduled task in the order specified; the NEs that execute an NE scheduled task operate independently and are not directed by the NFM-P. As soon as an NE completes an action in an NE scheduled task, it performs the next action.
  - First-run and last-run tests that have the Continuously Executed and Accounting Files options enabled can be added to test suites.
  - When issuing an Execute or Stop Execution command for a test suite that has the Continuously Executed and Accounting Files options enabled, the execution or stopping of the first-run and last-run tests will be triggered.
  - You must click Validate on the associated service or service tunnel configuration form to run the OAM validator test suite. The result of this OAM validation is indicated by the OAM Validation Failed operational flag. See 90.55 “To run a one-time validation test on a service” (p. 3084) for information about how to run a one-time validation test on a service. See

33.26 “To run an OAM validation test on a service tunnel” (p. 1225) for information about how to run a one-time validation test on a service tunnel.

- **Test result considerations:**
  - You can view test suite results for an object from the Tests tab of the configuration form for the object.
  - The NFM-P database logs test results only for scheduled test suites that are associated with NFM-P-defined schedules. Logging for test suites that are associated with NE schedules is performed locally on each target NE.
  - Size constraint policies control the number of historical records that the NFM-P retains in the database. You can configure size constraint policies to limit the number of database objects that are associated with specific test types. See the *NSP System Administrator Guide* for information about size constraint policies.

89.3 Sample STM implementation

89.3.1 STM test suites

The following figure shows a sample STM implementation that illustrates how you can use STM suites to verify service operation and identify network problems.

Figure 89-1 Sample STM implementation

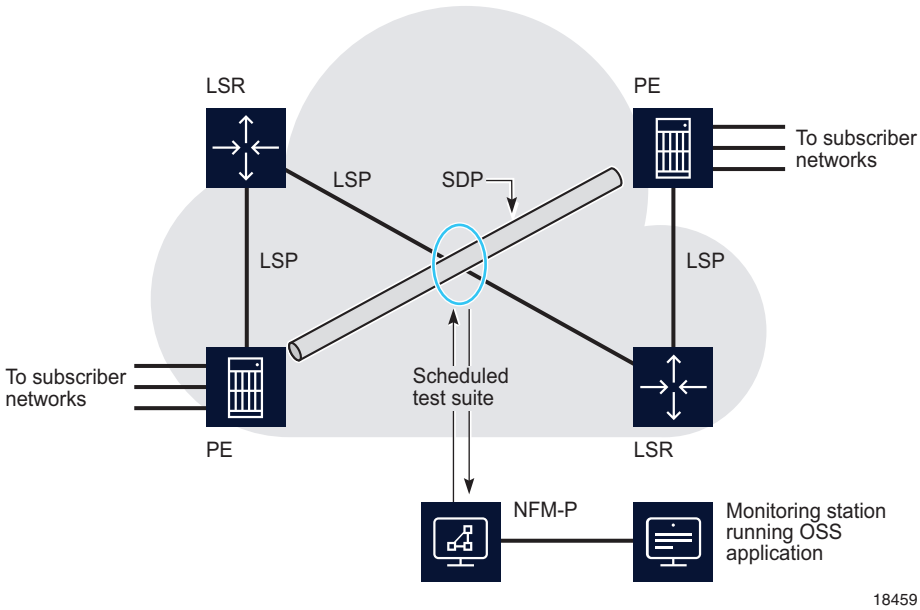

In the sample network, two scheduled STM test suites run constantly to ensure that the transport elements are functioning properly. One test suite monitors the SDPs, and the other test suite monitors the LSPs. An NFM-P administrator can quickly update the test suites, as required, to

reflect network topology changes by revising the list of tested entities in the suite and regenerating the tests. An OSS application collects the test result information which is then available for monitoring by NOC staff.

Note: This sample configuration uses an NFM-P schedule to create a scheduled task. See Chapter 5, “NFM-P-based schedules” for information about NFM-P schedules.

The SDP test suite contains the following generated tests for each SDP:

- tunnel ping
- MTU ping

The LSP test suite contains the following generated test for each LSP path:

- LSP ping (tested entity type is LSP path)

The NOC monitoring staff become aware that tunnel ping operations fail occasionally on one SDP. Packet loss is not yet significant enough to affect SLAs, but threatens to become so, based on the observed trend. NOC staff run a test suite against the LSPs in the affected LSP path.

The test suite contains the following generated test for each LSP in the LSP path:

- LSP ping (tested entity type is LSP)

Test results show that the packet loss is related to a specific LSP in the path. Investigation of the LSP traffic pattern indicates that a recently provisioned service is causing the LSP to be oversubscribed. The problem is addressed by a network designer and is corrected through configuration changes.

The following table lists the high-level tasks necessary to configure and use the STM elements in this sample.

Table 89-1 Sample STM implementation configuration sequence

Task	Description
1. Scheduled STM test suite creation	Create STM test policies for transport-layer elements. • Create a test policy for the SDPs. Specify Tunnel (SDP) as the Entity Type for the policy, and choose Tunnel Ping and MTU Ping as the test definitions. Create separate tunnel ping definitions for different forwarding classes, as required. • Create a test policy for the LSP paths. Specify LSP as the Entity Type for the policy and choose LSP Ping as the test definition. In the test definition, specify LSP Path as the Target Type. Create separate LSP ping definitions for different forwarding classes, as required. Create STM test suites for transport-layer elements. • Create a test suite for the SDPs. Specify Tunnel (SDP) as the Entity Type for the suite, choose the SDP test policy, choose the SDPs against which the suite is to run, and use the NFM-P to generate the tests in the suite. Create a schedule for the test suite according to network monitoring requirements, and apply the schedule to the test suite to create a scheduled task. • Create a test suite for the LSP paths. Specify LSP as the Entity Type for the suite, choose the LSP test policy, choose the LSP paths against which the suite is to run, and use the NFM-P to generate the tests in the suite. Create a schedule for the test suite according to network monitoring requirements, and apply the schedule to the test suite to create a scheduled task.

Table 89-1 Sample STM implementation configuration sequence (continued)

Task	Description
2. Data presentation	Customize the tabular display of test results in the NFM-P, or create an OSS application that retrieves test-result data from the NFM-P and presents it as information in graphical format for NOC staff.
3. Creation of non-scheduled STM test suite or individual OAM tests for network troubleshooting	Create an STM test policy for the LSPs for use during network troubleshooting. • Create a test policy for the LSPs. Specify LSP as the Entity Type for the policy and choose LSP Ping as the test definition. In the test definition, specify LSP as the Target Type. Create separate LSP ping definitions for different forwarding classes, as required. Create a non-scheduled STM test suite (or individual OAM tests, depending on the number of LSPs involved) for the LSPs in the LSP path. • Specify LSP as the Entity Type for the suite, choose the created LSP test policy, choose the LSPs in the LSP path against which the suite is to run, and use the NFM-P to generate the tests for the suite.
4. Network monitoring	Use the NFM-P to display the scheduled test results in tabular format, or use an OSS application to monitor the LSP ping, tunnel ping, and MTU ping diagnostic results in real time. Record inconsistencies and trends, and troubleshoot the network as required. Use the threshold-crossing alarm capabilities to raise alarms when traffic characteristics rise above or fall below specific values.
5. Network troubleshooting	When potential trouble arises, use a non-scheduled STM test suite or individual STM tests to help identify the cause. • Create individual tests or edit the existing non-scheduled test suite. Include as tested entities the LSPs that comprise the LSP path and regenerate the tests for the suite, as required. • Execute the non-scheduled test suite and monitor the test results. • Edit the test definition in the test policy. Change test parameter values as required, regenerate the tests in the suite, and re-execute the suite. Continue to refine the diagnostic test results until the problem manifests itself in the test results.

89.4 Sample STM network SLA monitoring configuration

89.4.1 Continual network topology monitoring

The following figure shows the high-level flow of tests and test results in a simple network topology that the network provider wants to continually monitor using the STM. See the sample network monitoring configuration steps that follow for the steps required to create the STM configuration.

Figure 89-2 Continual network topology monitoring

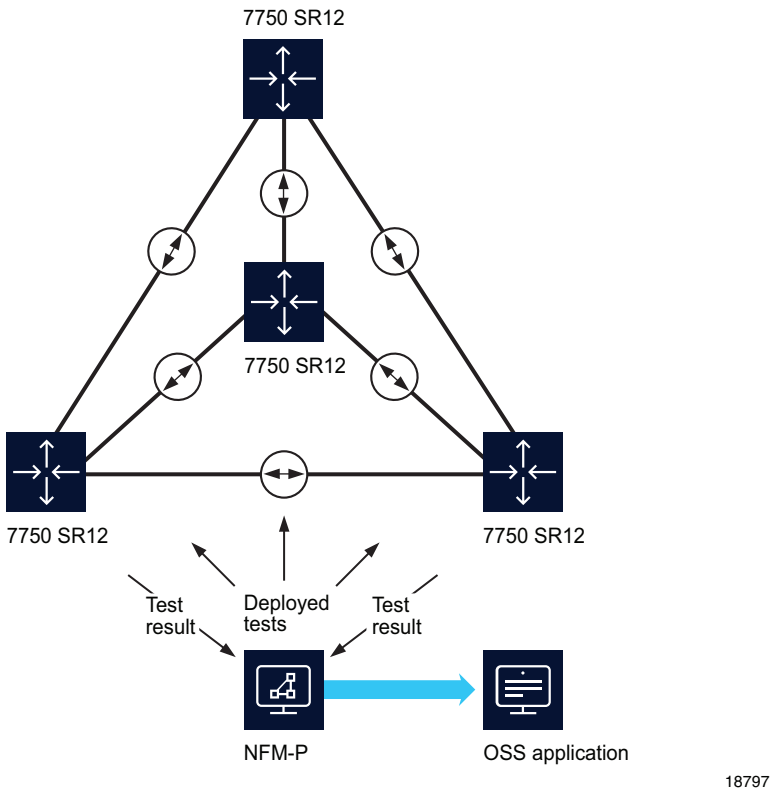

An NFM-P administrator creates and schedules an STM test suite to monitor the LSP mesh and identify potential overloading and reachability problems before they affect service traffic. The test suite contains ping and trace test definitions for three different classes of in-profile traffic and is scheduled to run every 15 min.

Note: This sample configuration uses an NFM-P schedule to create a scheduled task. See Chapter 5, “NFM-P-based schedules” for information about NFM-P schedules.

Because the same group of tests is to be run on multiple entities of the same type, the test suite consists of generated tests only.

The test policy that the NFM-P uses to generate the tests contains the following six test definitions:

- LSP ping - be in
- LSP ping - af in
- LSP ping - nc in
- LSP trace - be in
- LSP trace - af in
- LSP trace - nc in

The NFM-P generates 72 tests for the test suite—3 ping tests and 3 trace tests for each of the 12 LSPs and LSP paths. Every 15 min., the NFM-P scheduler executes the tests and stores the results. An OSS application periodically retrieves the test results, performs trend analysis, and provides the transport utilization information to NOC monitoring staff.

The NFM-P generates an alarm in the event of an LSP ping probe failure or an LSP trace path change, as specified in the test policy, so NFM-P operators become aware of new transport faults as they arise.

89.4.2 Sample network monitoring configuration steps

The following steps describe the configuration of the network monitoring sample shown in Figure 89-1, “Sample STM implementation” (p. 2909). For clarity, only the configuration steps required for the sample are described. See the procedures in “Procedures to use the STM” (p. 2945) for complete STM configuration information.

89.4.3 Stages

Create a test policy

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create and choose Test Policy. The Test Policy (Create) form opens.

3

Configure the parameters in Table 89-2, “Test policy General tab parameters” (p. 2913), using the supplied values.

Table 89-2 Test policy General tab parameters

Parameter	Value	Comment
Name	LSP Test Policy	—
NE Schedulable	disabled	NE-schedulable test results are not returned to the NFM-P.
Entity Type	LSP	—

4

Click on the Test Definitions tab.

5

To add LSP Ping test definitions to the test policy, click Add and choose MPLS→Add LSP Ping. The LSP Ping Definition (Create) form opens.

In the sample configuration, the operator creates a separate LSP Ping test definition for each of the following forwarding classes:

- be (best effort - low priority, delivery not guaranteed)
- af (assured forwarding - medium priority, delivery guaranteed)
- nc (network control - high priority, delivery guaranteed)

6

Configure the parameters for each test definition using the values supplied in Table 89-3, “LSP Ping Definition parameters” (p. 2913).

Table 89-3 LSP Ping Definition parameters

Parameter	Value	Comment
General tab		
Name	LSP Ping - be in LSP Ping - af in LSP Ping - nc in	Create three LSP Ping test definitions, one for each Name value.
Target Type	LSP	—
Test parameters tab		
Number of Test Probes	5	A value of 5 provides more analysis granularity than a single test probe, without affecting network performance.
Size (octets)	108	This is the minimum packet size for this type of test on the 7450 ESS and 7750 SR. The minimum packet size varies, depending on the type and release of device against which a test runs.
Forwarding Class	be af nc	For LSP Ping - be in test definition For LSP Ping - af in test definition For LSP Ping - nc in test definition To cover a range of traffic streams and to test different network queues, multiple test definitions are necessary. The sample configuration uses three LSP ping definitions that differ only by forwarding class.
Forwarding Profile	in	All test packets are marked in-profile. This means that they follow the same path as service traffic.
Results configuration tab		
Probe History Size (rows)	100	—
Trap Generation	Probe Failure	Test-probe failures provide early notification of LSP congestion or reachability issues to an NFM-P operator. OSS analysis of the results indicates trends for further investigation.

7

For each LSP Ping test Definition you create, click OK to save the changes and close the form. The Test Policy (Create) form is displayed, with the LSP Ping test definitions listed.

8

To add LSP Trace test definitions to the test policy, click Add and choose MPLS→Add LSP Trace. The LSP Trace Definition (Create) form opens.

In the sample configuration, the operator creates a separate LSP Trace test definition for each of the following forwarding classes:

- be (best effort - low priority, delivery not guaranteed)
- af (assured forwarding - medium priority, delivery guaranteed)
- nc (network control - high priority, delivery guaranteed)

9

Configure the parameters for each test definition using the values supplied in Table 89-4, “LSP Trace Definition parameters” (p. 2914).

Table 89-4 LSP Trace Definition parameters

Parameter	Value	Comment
General tab		
Name	LSP Trace - be in LSP Trace - af in LSP Trace - nc in	For first test definition For second test definition For third test definition
Target Type	LSP Path	—
Test parameters tab		
Number of Test Probes	5	A value of 5 provides more analysis granularity than a single test probe, without affecting network performance.
Size (octets)	108	This is the minimum packet size for this type of test on the 7450 ESS and 7750 SR. The minimum packet size varies, depending on the type and release of device against which a test runs.
Forwarding Class	be af nc	For LSP Trace - be in definition For LSP Trace - af in definition For LSP Trace - nc in definition To cover a range of traffic streams and to test different network queues, multiple test definitions are necessary. The sample configuration uses a different LSP Trace definition for each forwarding class.
Forwarding Profile	in	All LSP Trace test definitions in the sample involve in-profile traffic.
Results configuration tab		

Table 89-4 LSP Trace Definition parameters (continued)

Parameter	Value	Comment
Probe History Size (rows)	100	—
Trap Generation	Test Failure Path Change	A test failure trap may indicate a routing or congestion issue, but it also indicates test misconfiguration when the Maximum Time to Live parameter value is too low to allow a test probe to reach the destination LSP. A path change trap may indicate an intermittently congested or unresponsive LSP.

10

For each LSP Trace test Definition you create, click OK to save the changes and close the form. The Test Policy (Create) form is displayed, with the LSP Trace test definitions listed.

11

Click OK to save your changes and close the Test Policy (Create) form. The Service Test Manager (STM) form is displayed.

Create a schedule

12

Choose Tools→Schedules→Schedule from the NFM-P main menu. The Schedule form opens.

13

Click Create. The NFM-P Schedule (Create) form opens.

14

Configure the parameters in Table 89-5, “NFM-P Schedule parameters” (p. 2916) using the supplied values.

Table 89-5 NFM-P Schedule parameters

Parameter	Value	Comment
Name	LSP Assurance	—
Current Client Start Time	a future time	Specify the date and time at which the schedule is to be first triggered. This value is used to calculate the times of subsequent tests, based on the Frequency parameter value. Compare this value and the Frequency parameter value with those in other NFM-P schedules to ensure that the scheduled test times are staggered and the resulting NE task load is balanced.
Ongoing	enabled	This setting tells the NFM-P to repeatedly trigger execution of the task associated with the schedule.

Table 89-5 NFM-P Schedule parameters (continued)

Parameter	Value	Comment
Frequency	Per Minute / 15	Select the Run Every Minutes radio button and specify 15 as the number of minutes for the triggering frequency.

15

Save your changes and close the form.

16

Close the Schedule form.

Create an STM test suite

17

Click Create on the Service Test Manager form and choose Test Suite. The Test Suite (Create) form opens.

18

Configure the parameters in Table 89-6, “Test suite parameters” (p. 2917)using the supplied values.

Table 89-6 Test suite parameters

Parameter	Value	Comment
Name	LSP Test Suite	—
Entity Type	LSP	—
NE Schedulable	disabled	NE-schedulable test results are not returned to the NFM-P.

19

Choose a test policy to govern the generation of tests.

1. Click on the Test Policy tab and click Add.
2. Choose the LSP Test Policy entry you configured in Stage 3 and click OK. The test policy is listed on the Test Policy tab.

20

Choose the LSPs for the test suite to assess.

1. Click on the Tested Entities tab and click Add.
2. Choose the required LSPs and click OK. The LSPs are listed on the Tested Entities tab.

21

Click Apply to save the changes. Additional buttons appear and the form name changes to Test Suite (Edit).

22

Click on the Generated Tests tab.

23

Click Generate Tests. The NFM-P generates a test for each LSP listed on the Tested Entities tab. The tests are listed on the form as they are generated.

24

Apply a span of control to the test suite.

1. Click on the Spans tab and click Add.
2. Choose one or more configured spans to apply to the test suite, and click OK.

Create a scheduled task

25

Click Schedule at the bottom of the Test Suite form. The STM Scheduled Task (Create) form opens. LSP Test Suite is displayed in the Task panel.

26

Configure the parameters in Table 89-7, “NFM-P scheduled task parameters” (p. 2917) using the supplied values.

Table 89-7 NFM-P scheduled task parameters

Parameter	Value	Comment
Scheduled Task Name	Periodic LSP Check	—
Administrative State	Enabled	The NFM-P executes the scheduled task at regular intervals based on the specified Current Client Start Time and Frequency.

27

In the Schedule panel, click Select and choose the LSP Assurance schedule you configured in Stage 14.

28

Save your changes and close the forms.

89.5 Sample STM SAA accounting files configuration

89.5.1 SAA accounting file configuration scenario

This sample describes the configuration of SAA accounting files that are to be collected following the generation of an NE schedulable test of VLL services.

A network provider wants to use the NFM-P to view test result information for an NE schedulable test of VLL services, and wants the NFM-P to process the OAM results MIB entries and produce a compressed XML record that is stored in SAA accounting files on the nodes.

To perform this task, an NFM-P administrator creates an NE schedulable test for VLL services that records test result information in SAA accounting files. The administrator creates one accounting policy, one file policy, one test policy, and one test suite.

89.5.2 Sample SAA accounting files configuration steps

The following steps describe the end-to-end configuration for this sample. For clarity, only the configuration steps required for the sample are described. See the procedures in "Procedures to use the STM" (p. 2945) for complete STM configuration information.

Create an accounting policy

1

Choose Tools→Statistics→Accounting Policies from the NFM-P main menu. The Accounting Policies form opens.

2

Click Create. The Accounting Policy (Create) form opens.

3

Configure the parameters in Table 89-8, "Accounting policy parameters" (p. 2919) using the supplied values.

Table 89-8 Accounting policy parameters

Parameter	Value	Comment
Displayed Name	SAA Accounting Policy	—
Type	NE Schedulable Tests	You must choose this value to allow the configuration of SAA accounting files

4

Create a file policy.

1. Click Select in the File panel, The Select a File Policy form opens.
2. Click Create. The File Policy (Create) form opens.

3. Configure the required parameters. For this sample, enter SAA File Policy as the Displayed Name.
4. Click OK. The form closes and the policy is listed on the Select a File Policy form.

5

Select the newly-created SAA File Policy entry and click OK. The Accounting Policy (Create) form displays the new file policy in the File panel.

6

Click Apply, then click Switch Mode to change the configuration mode to Released.

7

Distribute the SAA Accounting policy to the NEs for which you require the SAA accounting statistics records. See 49.6 “To release and distribute a policy” (p. 1476).

8

Close the forms.

Create a test policy

9

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

10

Click Create and choose Test Policy. The Test Policy (Create) form opens.

11

Configure the parameters in Table 89-9, “Test policy parameters” (p. 2920) using the supplied values.

Table 89-9 Test policy parameters

Parameter	Value	Comment
Name	SAA Test Policy	—
NE Schedulable	enabled	You must enable this parameter to allow the configuration of SAA accounting files.
Entity Type	VLL Service	—
Accounting Files	enabled	You must enable this parameter to allow the configuration of SAA accounting files.

12

Add a test definition to the policy.

1. Click on the Test Definitions tab.
2. Click Add and choose Service→Add VCCV Ping. The VCCV Ping Definition (Create) form opens.
3. Configure the required parameters on the General, Test Parameters, and Results Configuration tabs.
4. Click OK to close the form. The VCCV Ping test definition is listed on the Test Policy (Create) form.

13

Click OK to close the Test Policy (Create) form. The Service Test Manager (STM) form is displayed.

Create a test suite

14

Click Create on the Service Test Manager (STM) form and choose Test Suite. The Test Suite (Create) form opens.

15

Configure the parameters in Table 89-10, “Test suite parameters” (p. 2920) using the supplied values.

Table 89-10 Test suite parameters

Parameter	Value	Comment
Name	SAA Test Suite	—
Entity Type	VLL Service	—
NE Schedulable	enabled	You must enable this parameter to allow the configuration of SAA accounting files.

16

Add the test policy you configured in Stage 11.

1. Click on the Test Policy tab, then click Add.
2. Choose the SAA Test Policy entry and click OK. The test policy is listed on the Test Policy tab.

17

Choose the NEs for which you require the SAA accounting statistics records.

1. Click on the Tested Entities tab, then click Add.
2. Choose the NEs that were selected for accounting policy distribution in Stage 7 and click OK. The NEs are listed on the Tested Entities tab.

18

Click Apply to save the changes. Additional buttons appear and the form name changes to Test Suite (Edit).

19

Click on the Generated Tests tab.

20

Click Generate Tests. The NFM-P generates a test for each NE listed on the Tested Entities tab. The tests are listed on the form as they are generated.

89.6 Sample STM threshold-crossing alarm configuration

89.6.1 Threshold-crossing alarm configuration scenario

The following sample configuration focuses on the configuration of a threshold-crossing alarm that is to be included in a test suite. It does not describe the configuration of the test suite, schedule or scheduled task associated with the sample. No test policy is required for the configuration, as it does not involve generated tests.

A network provider needs to use the NFM-P to monitor a customer VPLS named VPLS 17 that carries VoIP traffic. The goal is to have the NFM-P raise an alarm when the round-trip jitter value rises above a specified threshold.

Network jitter values typically range from 10 to 15 ms, but are occasionally as high as 25 ms. Jitter buffers in the end-user VoIP sets can accommodate up to 80 ms of jitter. A network engineer determines that a jitter test threshold of 30 ms is low enough to allow NOC operators sufficient time for investigating a jitter increase, yet high enough to exclude spurious jitter events.

An NFM-P administrator wants to create a scheduled CPE ping test suite for the VPLS that periodically measures the round-trip jitter between each VPLS site and an end device. The administrator creates one CPE ping test for each site in the VPLS.

89.6.2 Sample threshold-crossing alarm configuration steps

The following procedure defines the configuration steps required for this sample. See 89.4 “Sample STM network SLA monitoring configuration” (p. 2911) for information about sample test suite, test policy, or scheduling configurations. See the procedures in “Procedures to use the STM” (p. 2945) for complete STM configuration information. For conciseness, the procedure lists only the configuration steps that are specific to and necessary for the sample.

Create CPE ping test with threshold-crossing criteria

1

Click Create on the Service Test Manager (STM) form and choose L2 Service→CPE Ping. The CPE Ping (Create) form opens.

2

Configure the parameters in Table 89-11, “General CPE Ping test parameters” (p. 2922) using the supplied values.

Table 89-11 General CPE Ping test parameters

Parameter	Value	Comment
General tab		
Name	CPE Ping Test - VPLS 17, Site A	A descriptive name such as this is helpful for quickly identifying the entity under test when viewing the test results.
NE Schedulable	enabled	You must enable this parameter to allow the configuration of threshold-crossing criteria.
Select (in the Service panel)	VPLS 17	Click Search on the Select Service - CPE Ping form, then choose a service.
Select (in the Site panel)	Site A	Click Search on the Select Site - CPE Ping form, then choose a service.
Destination Path Address	*IP address of an end device connected to Site A*	Defines the IP address to which the CPE ping packets are sent.
Source IP Address	*IP address of the NE that originates the test packets*	This IP address must be in the same subnet as the Destination IP Address.
Source MAC Address	*MAC address of the NE that originates the test packets*	—
Test parameters tab		
Number of Test Probes	5	A value of 5 provides more analysis granularity than a single test probe, without affecting network performance.
Results configuration tab		
Probe History Size (rows)	100	This value provides a greater retained test-result history for occasional monitoring by an NFM-P operator.
Trap Generation	Probe Failure	Test-probe failures provide early notification of congestion or reachability issues to an NFM-P operator for troubleshooting purposes.

3

Click Apply. Additional buttons and tabs are displayed, and the form name changes to CPE Ping *Test Name* (Edit).

4

Click on the Threshold Alarms tab.

5

Click Add to add threshold criteria. The NE Threshold Event (Create) form opens.

6

Configure the parameters in Table 89-12, “Threshold-crossing CPE Ping test parameters” (p. 2923) using the supplied values.

Table 89-12 Threshold-crossing CPE Ping test parameters

Parameter	Value	Comment
Type	Round-Trip Jitter	—
Generate Alarm on Rising Threshold	enabled	This is the default setting. Alternatively or additionally, by enabling the Generate Alarm on Rising Threshold parameter, you can configure the NFM-P to generate an alarm when a test-result value falls below a specified threshold. You can also configure the NFM-P to clear a rising-threshold alarm when the test-result value falls below the value configured for the falling threshold.
Rising Threshold tab		
Threshold Value	30	This value specifies that the NFM-P is to generate an alarm when an individual test result contains a round-trip jitter value greater than 30 ms.

7

Save your changes and close the forms.

89.7 STM Y.1564 test configuration

89.7.1 ITU-T Y.1564 tests

The NFM-P supports the ITU-T Y.1564 test methodology, which is used to assess the proper configuration and performance of an Ethernet service before customer notification and delivery. The Y.1564 standard provides for measuring throughput, latency, frame loss, and jitter to assess if the service complies with an SLA. In the STM, Y.1564 tests are configured either as Y.1564 bidirectional tests, or Y.1564 service tests.

The Y.1564 bidirectional tests are supported on the 7705 SAR and 7210 SAS (except the 7210 SAS-E). These tests use a testhead profile for configuration of the test criteria. A variety of test types is supported, including bandwidth availability: see 89.7.2 “ITU-T Y.1564 bidirectional test” (p. 2925).

The Y.1564 service tests are supported on the 7210 SAS-K, supported FP4 or FP5 7x50, and 7250 IXR. These tests provide an enhanced configuration that differs from the bidirectional test configuration, using separate profiles or templates for acceptance criteria and payload, and allowing configuration of multiple test streams; see 89.7.8 “Y.1564 service tests” (p. 2935).

89.7.2 ITU-T Y.1564 bidirectional test

The Y.1564 bidirectional service tests are supported on the 7705 SAR and 7210 SAS (except the 7210 SAS-E). Support for specific functions varies depending on the device type and release; see the NE documentation for more information.

You can perform Y.1564 tests on VLL Epipe services configured on 7705 SAR NEs, and on VPLS and VLL Epipe services on 7210 SAS NEs, except for the following service types:

- PBB Epipe
- B-VPLS
- MVPLS
- I-VPLS

You can configure a supported NE as a testhead for performing Y.1564 tests. The testhead generates test data for various test types and sends it through the service from the source SAP to the remote destination SAP. The remote SAP performs a port loopback with MAC swap and returns the test data to the source SAP. The testhead analyzes the data and produces the test results.

For 7210 SAS NEs, you can assign a no-service access port or a virtual port as a testhead port. The source SAP for the service you are testing must reside on the same NE. For 7210 SAS-R devices, the testhead port and source SAP must reside on the same card.

For 7705 SAR NEs, you do not specify a testhead port; instead, test traffic is identified by a testhead MAC address specifically configured on the NE.

A testhead profile specifies the parameters for the test, including duration, line rate (CIR), burst size (PIR), frame payload details, and acceptance criteria. The values configured for acceptance criteria determine whether a test execution fails or succeeds. Profiles are created and distributed using the NFM-P policy framework. See Chapter 49, “Policies overview” for more information.

To prevent service disruptions when other services share the ports on which the test source SAP and remote SAP are configured, the NFM-P provides a function to move the SAPs to target ports where no services are configured. See Chapter 96, “Service throughput configuration” for more information.

You can use the NFM-P STM to create, save, and execute Y.1564 bidirectional tests for a configured service. You can add the Y.1564 bidirectional test to an NFM-P STM test suite as a first-run or last-run test, and you can use the test suite to schedule the test. Alternatively, you can use the Tests tab of the properties form for the service.

You can also use the NFM-P service topology map to configure or select Y.1564 bidirectional tests and view test results; see 4.2.2 “Managing OAM diagnostics from the topology view” (p. 176). Each test type displays results on a separate tab, except for step-load and bandwidth availability tests which display on one tab. A message is displayed when the number of results exceeds the available tabs; delete older results to allow new results to display.

Consider the following when you add the Y.1564 bidirectional test to an STM test suite.

- Test policies are not supported.
- Generated tests are not supported.
- The Y.1564 bidirectional test is not NE-schedulable.

- Parallel tests are not supported. You can run only one Y.1564 bidirectional test at a time. However, multiple test types can be run sequentially as a single test. See 89.7.5 “Y.1564 bidirectional test types” (p. 2926) in this section.
- You cannot include CIR step-load and CIR step-load color-aware tests as part of an STM test suite.

89.7.3 7210 SAS considerations

The system resource profile for the 7210 SAS testhead NE must be configured appropriately. You must allocate a minimum of six entries to the ingress internal ACL MAC resource pool, and a minimum of two entries to the egress internal resource pool (ACL MAC, IPv4 or IPv6 64-bit criteria). The testhead tool uses resources from these resource pools; if no resources are available, the testhead tool cannot function. You must not allocate resources to the egress internal ACL 128-bit IPv6 criteria. See 6.5.13 “System resource profile” (p. 220) in 6.5 “7210 SAS” (p. 216) and 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380) for more information about configuring the system resource profile.

You must configure a port loopback with MAC swap on the test source port before you initiate the test; see 16.24 “To configure Ethernet ports” (p. 599). Set the Type parameter to Internal.

When you configure a testhead profile for the 7210 SAS, the Profile ID value cannot be greater than 10.

You cannot choose an L2 uplink SAP as the source SAP.

If the service you are testing is a VPLS, you must assign a static MAC to the SAP or SDP (service circuit) on the source site; see 77.29 “To add or modify FIB entries associated with a VPLS” (p. 2290). Static MAC configuration is not required for Y.1564 tests on Epipe services.

For 7210 SAS-D NEs, if more than one mirror service is configured, you cannot run a Y.1564 bidirectional test.

The 7210 SAS-K supports the enhanced Y.1564 service test functionality in addition to the standard Y.1564 bidirectional test functionality. Configuration of the enhanced functionality (also called the service test testhead OAM tool) differs from the bidirectional test configuration; see 89.7.8 “Y.1564 service tests” (p. 2935).

89.7.4 7705 SAR considerations

You must configure a port loopback with MAC swap on the test destination SAP before you initiate the test; see 76.47 “To configure Ethernet loopback for a VLL Epipe L2 access interface on a 7705 SAR” (p. 2189). Set the Mode parameter to Internal, and enable the MAC Swap function.

Egress/Ingress scheduler mode must be set to 16-priority; see 76.44 “To configure scheduling on a VLL L2 access interface” (p. 2186).

For more information about Y.1564 bidirectional tests, including restrictions and configuration guidelines, see the NE documentation.

89.7.5 Y.1564 bidirectional test types

The following test types are supported:

- Single CIR

- Single CIR color-aware
- CIR step-load
- Traffic policing color-aware
- Single PIR
- Single PIR color-aware
- CIR step-load color-aware
- Bandwidth availability

When color-aware tests are selected, the NFM-P provides separate test results for in-profile and out-of-profile traffic. To execute color-aware tests, the Dot1P-In and Dot1P-Out parameters must be set to values other than the default values.

CIR step-load and CIR step-load color-aware tests are run in multiple iterations based on a user-defined percentage of the configured CIR value. With each iteration, the tested CIR value is increased. The percentage determines the number of steps to run. For example, if the configured CIR is 10000 kb/s, and the CIR Step (%) is 25, then the first iteration runs a test for a CIR of 2500 kb/s (25% of the 10000 kb/s line rate). The next iteration increases the test CIR by another 25% of the configured CIR, and so on for each iteration. In this example, the test runs four iterations, until the configured line rate is tested. Each iteration runs for 3 min, regardless of the setting of the test duration. If the first iteration fails, no further iterations are performed.

When a Traffic policing color-aware test is executed, the NFM-P uses a PIR that is calculated according to the formula:

PIR = CIR + 1.25 × (configured PIR value − configured CIR value).

The original configured PIR value is restored when the Traffic policing color-aware test is complete.

Traffic policing color-aware, Single PIR, and Single PIR color-aware tests will not execute if PIR (Admin) is disabled.

Configured test parameter values can change during test execution, depending on the type of test being executed. For example, when the Traffic policing color-aware test is executed, the PIR value is modified to suit the requirements of the test. When test execution is complete, the NFM-P restores the test parameter values to their original configuration.

You can select multiple test types to run in a single test execution, except for the Bandwidth availability test. Bandwidth availability tests must be run separately.

Bandwidth availability tests

The Bandwidth availability test measures the line rate for the service. The NFM-P manages the rate values and accuracy for the test, instructs the testhead to perform test iterations, and ends the test.

The first iteration of the test uses a rate of 960,000 kb/s by default. If the first iteration succeeds, the test ends. If the first iteration does not succeed, the test runs additional iterations to narrow in on a bandwidth that approaches the maximum available rate.

When acceptance criteria are at default values, a tested rate is successful when no packets are dropped. Otherwise, acceptance criteria settings influence the success or failure of an iteration.

A user-configured accuracy value determines the end of the test. When the difference between the highest successful test rate and the lowest failed rate is within the accuracy value, the test ends.

The last successful rate is displayed as the Available Bandwidth on the Y1564 Bi-Directional Test form. Each iteration of the Bandwidth availability test runs for three min, regardless of the test duration configured in the test profile. Lower values for the Accuracy parameter typically produce more iterations of the test. The NFM-P automatically ends a Bandwidth availability test if no SNMP traps are received from the testhead within a calculated time-out period.

89.7.6 Workflow to configure and run Y.1564 bidirectional tests

This workflow describes the high-level steps required to configure Y.1564 bidirectional tests on the NFM-P. For detailed Y.1564 bidirectional test configuration steps, see the sample procedure that follows this workflow. The service to be tested must be fully configured before running the test.

1. For a 7210 SAS testhead NE, configure the system resource profile. Allocate a minimum of six entries to the ingress internal ACL MAC resource pool, and a minimum of two entries to the egress internal resource pool (ACL MAC, IPv4 or IPv6 64-bit criteria). See 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380).
2. For 7210 SAS NEs, configure a port loopback with MAC swap on the test source port; see 16.24 “To configure Ethernet ports” (p. 599). Set the Type parameter to Internal.

   If you are configuring a test on a 7210 SAS-K site, enable the MAC-Swap-Enable parameter on the L2 access interface. For a VLL service, see 76.40 “To create a VLL L2 access interface on a terminating site” (p. 2174) ; for a VPLS, see 77.67 “To create a VPLS or MVPLS L2 access interface” (p. 2332).

   For configuration of enhanced Y.1564 service tests (also called the service test testhead OAM tool) on the 7210 SAS-K, see 89.7.8 “Y.1564 service tests” (p. 2935).

   For 7705 SAR NEs, configure an Ethernet loopback on the destination L2 access interface; see 76.47 “To configure Ethernet loopback for a VLL Epipe L2 access interface on a 7705 SAR” (p. 2189).
3. If the service you are testing is a VPLS, assign a static MAC to the SAP or SDP (service circuit) on the source site; see 77.29 “To add or modify FIB entries associated with a VPLS” (p. 2290).
4. Set the OLC state of the service to Maintenance. See the *NSP System Administrator Guide* for more information.
5. Configure the testhead.

   For the 7210 SAS, select a testhead port on the test source NE.

   For the 7705 SAR, configure the Marker Source MAC Address parameter on the test source NE.
6. Create a testhead profile and distribute the profile to the NE that contains the test-head port and the source SAP of the service to be tested.

   For 7210 SAS NEs, the Profile ID value cannot be greater than 10.
7. Configure the Y.1564 bidirectional test in the NFM-P STM.
8. Move SAPs as required; see Chapter 96, “Service throughput configuration” for more information.
9. Add the Y.1564 bidirectional test to an STM test suite, if required; see 89.12 “To create an STM test suite” (p. 2951) and the overview to 89.7 “STM Y.1564 test configuration” (p. 2924).
10. Execute the test and view the results.

11. Restore the original SAP configuration, if required.

89.7.7 Sample Y.1564 bidirectional test configuration steps

The following procedure lists the configuration steps that are required for this sample. See the procedures in "Procedures to use the STM" (p. 2945) for complete STM configuration information. For configuration of enhanced Y.1564 service tests on the 7210 SAS-K, see 89.7.8 "Y.1564 service tests" (p. 2935).

Prepare the service for the test

1

If the test source NE is a 7705 SAR, configure Ethernet loopback on the destination SAP; see 76.47 "To configure Ethernet loopback for a VLL Epipe L2 access interface on a 7705 SAR" (p. 2189).

If the test source NE is a 7210 SAS, configure a port loopback with MAC swap on the test source port; see 16.24 "To configure Ethernet ports" (p. 599). Set the Type parameter to Internal.

When you configure a test on a 7210 SAS-K, MAC swap address configuration is not required. However, if you are configuring a test on a 7210 SAS-K site, enable the MAC-Swap-Enable parameter on the L2 access interface. For a VLL service, see 76.40 "To create a VLL L2 access interface on a terminating site" (p. 2174) ; for VPLS, see 77.67 "To create a VPLS or MVPLS L2 access interface" (p. 2332).

2

For 7705 SAR NEs, ensure that the egress and ingress scheduler mode for the test SAPs is set to 16-priority; see 76.44 "To configure scheduling on a VLL L2 access interface" (p. 2186).

3

If the service you are testing is a VPLS, assign a static MAC to the SAP or SDP (service circuit) on the source site; see 77.29 "To add or modify FIB entries associated with a VPLS" (p. 2290).

When you configure the test on a 7210 SAS-K, static MAC configuration is not required for VPLS.

4

Set the OLC state of the service to Maintenance. See the *NSP System Administrator Guide* for more information.

Select a testhead port (7210 SAS only)

5

Note: You must configure the testhead port on the 7210 SAS NE that contains the source SAP of the service you need to test. For 7210 SAS-R NEs, the test-head port and source SAP must reside on the same card.

Open the properties form for the 7210 SAS NE that you need to configure.

6

Click on the Globals tab, then click on the Service tab.

7

Note: If you select a physical port as a no-service port, the port cannot contain any SAPs.

Note: Not all chassis types support virtual no-service ports. The number of available virtual ports and the port names vary depending on the chassis type. Virtual ports are not displayed on the navigation tree.

Note: You cannot select the same no-service port for more than one function.

Select a loopback no-service testhead port. Perform one of the following:

a.

Select a physical test-head port.

1. In the TestHead Port panel, click Clear if required, then click Select. The Select TestHead Port form opens.
2. Choose a port and click OK.

b.

Select a virtual testhead port.

1. Enable the Use Virtual TestHead Port parameter.
2. In the Virtual TestHead Port panel, click Clear if required, then click Select. The Select Virtual TestHead Port form opens.
3. Choose a port and click OK.

   The available virtual ports vary depending on the chassis type and card type.

8

Save your changes and close the form.

Configure a testhead MAC address (7705 SAR only)

9

Open the properties form for the NE that you need to configure; choose the 7705 SAR NE that contains the source SAP of the service you need to test.

10

Click on the Globals tab, then click on the OAM tab.

11

On the General tab, configure the Marker Source MAC Address parameter.

Nokia recommends that you use a non-zero MAC address that is unique in the tested network.

Create a testhead profile

12

Note: The testhead profile is distributed to the NE using the NFM-P policy distribution framework; see Chapter 49, “Policies overview”. You must distribute the testhead profile for the Y.1564 bidirectional test to the source site for the service that you are testing.

Choose Tools→Y1564 Profiles→Y1654 Test Head Profile from the NFM-P main menu. The Y1564 Test-Head Profiles form opens.

13

Click Create. The Y1564 Test-Head Profile (Create) form opens.

14

Configure the required general parameters.

For 7210 SAS NEs, the Profile ID value cannot be greater than 10.

The configured PIR value must be greater than the configured CIR value.

To execute color-aware tests, the Dot1P-In and Dot1P-Out parameters must be set to values other than default.

Traffic Policing Color-Aware, Single PIR, and Single PIR Color-Aware tests will not execute if the Disable PIR check box is selected.

15

Configure frame payload details. You can create up to eight frame payload configurations for each testhead profile.

1. Click on the Frame Payload tab, then click Create. The Y1564 Test-Head Payload (Create) form opens.
2. Configure the required parameters.

   The available parameters vary depending on the option you choose for the Type parameter.

   Additional parameters are available when you configure the VLAN Tag 1 and VLAN Tag 2 parameters.
3. Save your changes and close the form. The Y1564 Test-Head Profile (Create) form is displayed.

16

Configure acceptance criteria. You can create up to eight acceptance criteria configurations for each testhead profile.

1. Click on the Acceptance Criteria tab, then click Create. The Y1564 Test Head Acceptance Criteria (Create) form opens.
2. Configure the required parameters.

   Disable the Default check boxes to configure values other than the default values.
3. Save your changes and close the form. The Y1564 Test-Head Profile (Create) form is displayed.

17

Click Apply to save the configured profile. The Y1564 Test-Head Profile (Create) form is displayed with additional buttons.

18

Click on the General tab, then click Switch Mode to change the configuration mode to released.

19

Distribute the profile to the NE that is the testhead and source site for the service to be tested. See 49.6 “To release and distribute a policy” (p. 1476).

You cannot distribute a testhead profile to 7210 SAS NEs if the Profile ID value is greater than 10.

Configure and execute the test using the NFM-P STM

20

Note: You can only execute tests for VPLS or VLL Epipe services that are fully configured in the NFM-P.

Configure initial test settings.

1. Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

   Alternatively, you can configure tests from the service topology map, or from the Tests tab of the properties form for the service you need to test.
2. Click Create and choose Service →Y1564 BiDirectional Test from the drop-down menu. The Y1564 Bi-Directional Test (Create) form opens.
3. Configure the parameters on the General tab.

   You can select and execute multiple test types in a single test, except for the Bandwidth availability test. When the Bandwidth availability test type is selected, you cannot select other test types. See 89.7.5 “Y.1564 bidirectional test types” (p. 2926).

21

Choose a service for the test.

1. Click Select in the Service panel. The Select Service - Y1564 Bi-Directional Test form opens.
2. Choose the service that you need to test and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the service information and additional panels.

22

Configure the Source Site and SAP.

1. Click Select in the Source Site panel. The Select Source Site - Y1564 Bi-Directional Test form opens.
2. Choose the 7210 SAS site that contains the source SAP for the service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Source Site information and additional panels.
3. Click Select in the Source Test SAP section of the Source SAPs panel. The Select Source Test SAP - Y1564 Bi-Directional Test form opens.
4. Choose the port that contains the source SAP for the service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Source Test SAP information displayed.

23

Note: Configuration of the Remote Site and SAP is optional in an Epipe service because the second terminating site becomes the remote site and SAP by default. However, you must perform this step if you need to select a Remote SAP Target Port in order to move SAPs; see Stage 25.

Configure the Remote Site and SAP, if required.

1. Click Select in the Remote Site panel. The Select Remote Site - Y1564 Bi-Directional Test form opens.
2. Choose the site that contains the remote SAP for the service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with an additional panel and the Remote Site information displayed.
3. Click Select in the Remote Test SAP section of the Remote SAPs panel. The Select Remote Test SAP - Y1564 Bi-Directional Test form opens.
4. Choose the port that contains the remote SAP for the service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Remote SAP information displayed.

24

Configure the Forwarding Class parameter, if required.

25

Note: When the ports that contain the source SAP or remote SAP also contain SAPs for services other than the service you are testing, the test may have an impact on those other services. If service impact is a concern, you can choose to move the test source SAP or remote SAP to other ports that will not have an impact on services. When you move SAPs, leave the Ignore Service Impact Validations check box unselected. See Chapter 96, “Service throughput configuration” for more information.

If required, move the SAPs to prevent service disruptions.

1. If service impact is not a concern, enable the Ignore Service Impact Validations parameter and go to Stage 26.
2. Click Select in the Target Port section of the Source SAPs panel. The Select Target Port - Y1564 Bi-Directional Test form opens.
3. Choose a port that does not have an associated service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Target Port name displayed.
4. Click Select in the Target Port for Remote SAP section of the Remote SAPs panel. The Select Target Port for Remote SAP- Y1564 Bi-Directional Test form opens.
5. Choose a port that does not have an associated service and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Target Port for Remote SAP name displayed.

26

Choose and configure a testhead profile for the test.

1. Click on the Test Parameters tab.
2. Click Select in the Profile panel. The Select Profile - Y1564 Bi-Directional Test form opens.
3. Choose a profile and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the profile information and an additional panel.
4. If the site for the service to be tested service is a 7705 SAR, configure the Performance Monitoring parameter.

   When Performance Monitoring is enabled, time-stamped test packets are generated and the test measures latency and jitter. Performance monitoring can affect throughput results for some tests.
5. Click Select in the Payload section of the Payload and Acceptance Criteria panel. The Select Payload - Y1564 Bi-Directional Test form opens.
6. Choose a Frame Payload configuration and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Payload information displayed.

   For services configured on 7705 SAR NEs, you can choose up to four frame payload configurations per test.
7. If required, click Select in the Acceptance Criteria section of the Payload and Acceptance Criteria panel. The select Acceptance Criteria - Y1564 Bi-Directional Test form opens.
8. Choose an Acceptance Criteria configuration and click OK. The Y1564 Bi-Directional Test (Create) form reappears with the Acceptance Criteria ID displayed.

27

Click Apply. The Prepare and Execute buttons are activated.

28

If you configured a target port to perform a SAP move in Stage 25, click on the General tab, then click Prepare. The Source Test SAP and Remote Test SAP are moved to the Target ports; see Chapter 96, “Service throughput configuration” for more information.

29

Click Execute to run the test. The Results tab of the Y1564 Bi-Directional test is displayed. The test appears in the list.

30

Choose the test from the list to view the results. The Y1564 Bi-Directional Test - Results (View) form opens. The Result Status field displays Pending until the test is complete. The test duration is configured in the test profile and is displayed in the Test Duration panel. When the test is complete, the test information is displayed.

Note: When multiple tests are selected, the Current Test Type in Execution is displayed on the General tab of the Y1564 Bi-Directional Test form. When tests are complete, the current test type displayed is None.

Note: You can stop the test execution at any time by clicking Stop Execution. The NFM-P stops the current test type and runs the next selected test type, if any.

Note: Configured test parameter values can change during test execution, depending on the type of test being executed. For example, when the Traffic policing color-aware test is executed, the PIR value is modified to suit the requirements of the test. When test execution is complete, the NFM-P restores the test parameter values to their original configuration.

31

To restore the moved SAPs to their original ports, click on the General tab, then click Restore; see 96.6 “To restore a service after a throughput configuration” (p. 3255).

89.7.8 Y.1564 service tests

The NFM-P supports enhanced Y.1564 service test functionality on the 7210 SAS-K, supported FP4 or FP5 7x50, and 7250 IXR. For the 7210 SAS-K, the enhanced functionality is also called the service test testhead OAM tool. Y.1564 service tests are configured using separate profiles or templates for acceptance criteria, frame mix and payload. The tests allow for the configuration of multiple streams (also called flows), for which service performance metrics can be obtained.

Each stream can contain the following test types:

- CIR
- CIR-PIR

- Policing
- Performance
- Bandwidth Availability (NFM-P-only test, not present in CLI)

Two options are available for selecting the stream order, which determines how the tests in each stream are executed. The options are:

- Ordered
- Parallel

When Ordered is selected, tests are executed in sequence by test type, but within each test type the tests are executed in numerical order by stream ID. For example, all CIR tests are executed, beginning with the test configured in stream 1, then stream 2, and so on. Then all CIR-PIR tests are executed, again beginning with the lowest-numbered stream ID.

When Parallel is selected, tests are executed in sequence by test type, but within each test type, the tests are executed concurrently for all streams. For example, CIR tests are executed on both stream 1 and stream 2 at the same time, followed by CIR-PIR tests on both streams.

The Performance test type always executes in parallel, regardless of the stream order option selected.

If any of the test types fails, further execution of the remaining test types is halted, and a failed result is shown for the service test.

The NFM-P provides easily-viewed test results for the overall test, and for test types and streams. You can choose to collect and manage test results as accounting statistics, using an accounting policy.

The Y.1564 service test is configured, saved, and executed in the NFM-P STM using menu options and forms that are different from the Y.1564 bidirectional tests; see 89.7.10 “Sample Y.1564 service test configuration” (p. 2937).

See the NE OAM guides for configuration guidelines and more information about the service test testhead OAM tool.

89.7.9 Workflow to implement Y.1564 service tests

The following workflow outlines the basic steps required to perform Y.1564 service tests using the STM.

1

If required, configure an accounting policy for test results collection.

2

Configure profiles or templates for acceptance criteria, frame mix, and payload.

3

Configure the service test general parameters, to select the testhead NE stream order and test durations. Assign an accounting policy if required.

4

Configure service test streams to define test types and to select the required profiles or templates.

5

Execute the test.

6

View the results.

89.7.10 Sample Y.1564 service test configuration

The following sample procedure lists the detailed configuration steps required to configure a Y.1564 service test using the STM. See the procedures in “Procedures to use the STM” (p. 2945) for complete STM configuration information.

1

If required, configure an accounting policy to collect and manage test results, and distribute the policy to the testhead NE; see “To configure an accounting policy” in the *NSP NFM-P Statistics Management Guide*. You must choose Service TestHead for the Type parameter.

Configure test profiles or templates

2

Note: You must distribute the test profiles or templates for the Y.1564 service test to the source site for the services that you are testing. Test profiles and templates are distributed to the NE using the NFM-P policy distribution framework; see Chapter 49, “Policies overview”.

If you are configuring tests for a supported FP4 or FP5 7x50 site or a 7250 IXR site, go to Stage 5

For 7210 SAS-K NEs, configure an acceptance criteria profile.

1. Choose Tools→Y1564 Profiles→Y1564 Acceptance Criteria Profile from the NFM-P main menu. The Y1564 Acceptance Criteria Profiles form opens.
2. Choose a profile from the list, or click Create. The Y1564 Acceptance Criteria Profile (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the forms.

3

For 7210 SAS-K NEs, configure a frame mix profile.

1. Choose Tools→Y1564 Profiles→Y1564 Frame Mix Profile from the NFM-P main menu. The Y1564 Frame Mix Profiles form opens.
2. Choose a profile from the list, or click Create. The Y1564 Frame Mix Profile (Create|Edit) form opens.

3. Configure the required parameters. The default size values are defined by ITU-T Y.1564.
4. Save your changes and close the forms.

4

For 7210 SAS-K NEs, configure a payload profile.

1. Choose Tools→Y1564 Profiles→Y1564 Payload Profile from the NFM-P main menu. The Y1564 Payload Profiles form opens.
2. Choose a profile in the list, or click Create. The Y1564 Payload Profile (Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the forms.

5

For supported FP4 or FP5 7x50 NEs, or 7250 IXR NEs, configure an acceptance criteria template.

1. Choose Tools→Y1564 Profiles→Acceptance Criteria Template from the NFM-P main menu. The Acceptance Criteria Template form opens.
2. Choose a template from the list, or click Create. The Acceptance Criteria Template, Global Policy(Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the forms.

6

For supported FP4 or FP5 7x50 NEs, or 7250 IXR NEs, configure a frame size template.

1. Choose Tools→Y1564 Profiles→ Frame Size Template from the NFM-P main menu. The Frame Size Template form opens.
2. Choose a template from the list, or click Create. The Frame Size Template, Global Policy(Create|Edit) form opens.
3. Configure the required parameters.
4. Save your changes and close the forms.

Configure general service test settings

7

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

8

Choose a test from the list, or click Create and choose Service →Y1564 Service Test from the drop-down menu. The Y1564 Service Test (Create) form opens.

9

Configure the parameters on the General tab.

Select the site that contains the service SAPs to be tested, and click Apply. The selected site becomes the testhead for generating test traffic.

If required, select the accounting policy configured in Stage 1.

The parameters and configurations available for a test vary depending on the type of NE that is configured for the Site.

Configure service test streams

10

Click on the Service Test Streams tab, where a list of service test streams is displayed.

11

Choose a stream from the list, or click Create. The Y1564 Service Test Stream (Create|Edit) form opens.

12

Configure the required parameters, and assign the required profiles (for 7210 SAS) or templates (for supported FP4 or FP5 7x50 or 7250 IXR).

13

Save your changes and close the form.

Execute the test and view the results

14

Click on the Execute button. The test runs for the duration configured in “Configure general service test settings” (p. 2938), unless a failed result is returned for a stream. Test execution stops when a test fails.

15

When the test is complete, click on the Results tab. A list of test results is displayed.

16

Choose a result from the list, and click Properties. The Y1564 Service Test Result (View) form opens.

17

View the overall test results.

18

To view the service test stream results, click on the Streams tab. A list of streams is displayed.

19

Choose a result from the list, and click Properties. The Y1564 Service Test Stream Result (View) form opens.

20

View the stream results.

21

Close the form.

89.8 Sample OmniSwitch device SLA testing

89.8.1 CPE SLA Test-Head OAM test

You can use the NFM-P STM to perform a CPE SLA Test-Head OAM test to validate and test customer SLAs used on select OmniSwitch devices in Metro Ethernet networks. The test is critical for provisioning or troubleshooting network services between customer endpoints. The CPE SLA Test-Head OAM test supports unidirectional traffic and IPv4.

You can perform the following tasks on the NFM-P without an external test-head device:

- Generate specific flow-based traffic across the network to help identify flow-based issues.
- Identify the impact of QoS settings (for example, SAP profile or QoS policies) on the overall traffic.
- Check the throughput across a provider network. The throughput is displayed on the NFM-P GUI for each iteration of the test.
- Debug flow-specific traffic forwarding across the provider network.
- Analyze the behavior of various user-defined traffic patterns across the provider network.
- Perform the handover testing after the initial deployment.
- Perform on-demand testing and results monitoring using a central entity.

You can perform a CPE SLA Test-Head OAM test using one CPE test-head profile or a group of profiles. You can also run several tests for one CPE Test-Head profile or group of profiles.

The following conditions are applicable to the creation of a CPE SLA Test-Head profile or a group of profiles and their execution:

- A CPE SLA Test-Head Group profile can contain up to eight stream/flows.
- A CPE SLA Test-Head Group profile can contain up to eight analyzers.
- When you configure the STM test execution details, the analyzer test name and the generator test name must be the same.
- If the analyzer has a test (flow) which is not in the generator group profile, the test is not valid.
- When you run a CPE SLA test or test group, only one test can be run at a time. You cannot run

concurrent tests on the same device.

89.8.2 Sample CPE SLA Test-Head profile configuration

The following configuration steps describe how to create a CPE SLA Test-Head profile or a group of profiles on supporting OmniSwitch devices. For information about device support, see the NE documentation.

Note: For OS 6400, OS 6850E, OS 6855 U24X, and OS 9000E NEs running AOS Release 6.4.5, CPE SLA Test-Head profiles and CPE SLA Test-Head Group profiles can only be distributed if the profile has a Tx Rate of 64 kbps or above. If running AOS Release 6.4.6 or later, the Tx Rate must be 88 kbps or above.

89.8.3 CPE SLA Test-Head profile configuration steps

The procedure lists only the configuration steps that are required for this sample. See the procedures in "Procedures to use the STM" (p. 2945) for complete STM configuration information.

Create a CPE SLA Test-Head profile

1

Choose Tools→Y1564 Profiles→AOS CPE Test-Head Profile from the NFM-P main menu. The AOS CPE Test-Head Profiles form opens.

2

Click Create. The AOS CPE Test-Head Profile (Create) form opens.

3

Configure the required parameters.

Additional panels and parameters are available when you configure the Frame Type parameter or set the L2 Saa parameter to true.

4

Click Apply, then click Switch Mode to change the configuration mode to Released.

5

Distribute the CPE Test-Head profile to the required NEs. See Chapter 49, "Policies overview" for more information about policies and policy distribution.

6

Configure the local VLAN and port settings for the CPE Test-Head profile.

1. Click on the Local Definitions tab, then click Search. A list of the NEs to which you distributed the profile in Stage 5 is displayed.
2. Choose an NE and click Properties. The CPE Test-Head Profile form opens.

3. In the VLAN panel, click on the Select button for the Service ID parameter. The Select vlanSitePointer - CPE Test-Head Profile form opens.
4. Click Search and choose a VLAN site.
5. Click Local Audit On to change the mode to Local Edit only.
6. Click on the Select button for the Port parameter. The Select vlan port - CPE Test-Head Profile form opens.
7. Click Search and choose a port.
8. Click Apply and close the form.

7

As required, repeat Stage 1 to Stage 6 to create additional CPE SLA Test-Head profiles.

The CPE SLA test OAM diagnostic test on the STM (see Stage 10), can utilize the following profile types:

- Generator profile
- Analyzer profile
- Loopback profile

You can also group profiles as described in Stage 8 ; otherwise, go to Stage 10.

Create an AOS CPE SLA Test-Head Group profile

8

If required, create an AOS CPE SLA Test-Head Group profile which can be use to bind individual CPE Test-Head profiles to form a group of tests.

Perform the following:

1. Choose Tools→Y1564 Profiles→AOS CPE Test-Head Group Profile from the NFM-P main menu. The AOS CPE Test-Head Group Profile form opens.
2. Click Create. The AOS CPE Test-Head Group Profile (Create) form opens.
3. Configure the required parameters.

   If the Role parameter is set to the Generator option, perform Stage 9. Otherwise, continue with Stage 8.
4. Click on the CPE Test Group Member tab, then click Create. The Select Port - AOS CPE Test Group Profile form opens.
5. Choose the required CPE Test-Head profiles and click OK. The AOS CPE Test-Head Group Profile (Create) form opens displaying the selected profiles.
6. Click OK to save the AOS CPE Test-Head Group profile.
7. Click Search to display the AOS CPE Test-Head Group profile.
8. Choose the AOS CPE Test-Head Group profile and perform Stage 4, Stage 5, and Stage 6.
9. Go to Stage 10.

9

Configure the required feeder port for the generator.

1. Open the Properties form for the required NE.
2. Click on the Globals tab, then on the AOS OAM tab.
3. Click on the Select button for the Feeder Port Name parameter. The Select Feeder Port form opens.
4. Choose the required feeder port and click OK. The NE Properties form reappears.
5. Save your changes and close the form.
6. Perform the rest of the substeps in Stage 8.

Configure the CPE SLA test execution details on the STM

10

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

11

Click Create and choose one of the following options:

a. AOS CPE→CPE SLA Test. This option allows you to run a test on a CPE test-head profile or a group of profiles. The CPE SLA Test (Create) form opens. Perform Stage 12 to Stage 16, and Stage 18 to configure this option.

b. AOS CPE→CPE SLA Test Group. This option allows you to run several tests on a on a CPE test-head profile or group of profiles. The CPE SLA Test Group (Create) form opens with the General tab displayed. Complete Stage 12, Stage 13, Stage 16, Stage 17, and Stage 18 to configure this option.

12

Configure the required parameters.

13

Click Select in the Generator panel and choose a generator profile.

14

Click Select in the Analyzer panel and choose an analyzer profile.

15

Click Select in the Loopback panel and choose an loopback profile.

16

Click on the Test Parameters tab and configure the required parameters.

17

Click on the CPE SLA Test Group Members tab, then click Create. The CPE Test-Head Group Profile form opens.

1. Click Search to display the CPE Test-Head Group profiles.
2. Choose the required CPE Test-Head Group profiles and click OK.

18

Save your changes and close the forms.

Run the STM test and analyze the results

19

On the Service Test Manager (STM) form, choose one of the following options from the object drop-down menu, then click Search.

a. CPE SLA Test (AOS SAS). This option displays all single CPE test-head profile tests that have been created.

b. CPE SLA Test Group (AOS SAS). This option displays all group CPE test-head profile tests that have been created.

20

Choose the appropriate CPE SLA test and click Execute to run the test.

21

To view the test results, click on the Results tab.

Procedures to use the STM

89.9 STM workflow

89.9.1 Stages

1

As required, change the STM default settings to meet your operational requirements when creating or running OAM diagnostics tests with the STM.

a. Enable the debug STM mode parameter on the NFM-P user preferences form to access additional forms on the STM that can be used to configure OAM diagnostic test limits and view additional test configuration information. See 1.23 “To configure NFM-P user preferences” (p. 116) for more information.

b. Customize the NFM-P system preferences such as the default retention time for db test results and log file. See the *NSP System Administrator Guide* for more information.

2

As required, review the appropriate workflow for the OAM diagnostic test you need to run with the STM to determine if any prerequisite steps are required.

a. Review “Procedures to configure and perform OAM diagnostic tests” (p. 3006)in Chapter 90, “OAM diagnostic tests”.

b. Review 89.7.6 “Workflow to configure and run Y.1564 bidirectional tests” (p. 2928) in this chapter.

c. Review “Procedures to configure Ethernet CFM” (p. 3107)in Chapter 91, “Ethernet CFM”.

3

Create the individual OAM diagnostic tests to be run as a discrete test or to be used in an STM test policy or STM test suite. See Table 90-1, “NFM-P supported OAM diagnostic tests and configurations” (p. 2981) for a list of all NFM-P supported OAM diagnostic tests that can be configured using the STM and their applicable procedures.

4

Create or modify an STM test policy to specify which OAM test definitions are to be included in the STM test suite; see 89.10 “To configure an STM test policy” (p. 2947).

5

As required, configure threshold-crossing alarms or NM threshold-crossing alarms as a post-STM OAM diagnostic test or STM test policy creation exercise; see 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949).

6

Create or modify an STM test suite; see 89.12 “To create an STM test suite” (p. 2951) or 89.13 “To modify an STM test suite and view additional information” (p. 2955).

7

As required, configure test limits before running the OAM diagnostic test or test suite to view additional test configuration information after performing the test; see 89.14 “To configure OAM diagnostic test limits on the STM and view additional test configuration information” (p. 2957).

8

Run individual or multiple OAM diagnostic tests from the STM and immediately view the test results; see 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959).

9

As required, view and compare OAM diagnostic test results on the STM as a post-analysis exercise; see 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960).

10

As required, schedule the execution of an STM test suite as a one-time event or as an ongoing task.

1. Create an NFM-P-based schedule before you schedule the execution of an STM test suite; see 5.6 “To configure an NFM-P-based schedule” (p. 193).
2. Associate an STM test suite with an NFM-P-based schedule; see 5.7 “To associate a task with an NFM-P-based schedule” (p. 194).
3. As required, view which STM test suites are associated with an NFM-P-based schedule; see 5.8 “To view scheduled tasks associated with an NFM-P-based schedule” (p. 195).
4. Turn up the STM test suite scheduled task to allow the test suite to execute as per the schedule details; see 5.10 “To turn up or shut down an NFM-P-based scheduled task” (p. 196).

11

Run an STM test suite from the STM; see 89.17 “To execute an STM test suite” (p. 2961). Alternatively, perform 5.11 “To immediately execute an NFM-P-based scheduled task” (p. 197) to immediately execute the STM test suite from the Scheduled Task form.

12

Monitor the diagnostic results from the scheduled test suite and the alarm list for indications of network or service faults and threshold-crossing alarms.

a. View the STM test suite results; see 89.18 “To view STM test suite results” (p. 2962).

b. View and compare the STM test suite results for a tested entity; see 89.19 “To view and compare STM test suite results for a tested entity” (p. 2963).

13

Interpret OAM diagnostic test results on the STM; see 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964).

14

As required, edit or delete an OAM diagnostic test or test suite; see 89.21 “To edit an OAM diagnostic test” (p. 2975), 89.22 “To delete an OAM diagnostic test” (p. 2975), and 89.23 “To delete an STM test suite” (p. 2976).

89.10 To configure an STM test policy

89.10.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Perform one of the following:

a. Create an STM policy. Click Create and choose Test Policy. The Test Policy (Create) form opens.

b. Modify an exiting STM policy. Perform the following:

1. Choose Test Policy (Assurance) from the object menu, then click Search.
2. Choose a test policy and click Properties. The Test Policy (Edit) form opens.

Note: If you modify a test policy that is associated with a test suite, you must click Generate to regenerate the tests in the test suite.

Note: You can apply changes made to an STM test policy to all associated STM test suites by clicking Update Test Suites at the bottom of the Test Policy (Create|Edit) form.

3

Configure the required parameters.

Some parameters are configurable only during test policy creation, and cannot be modified for existing test policies.

The NE Schedulable and MEF35 parameters are mutually exclusive. The NE Schedulable parameter must be selected for SAA-type STM test policies. The MEF35 parameter must be selected for PM Session-type STM test policies.

The Test Family parameter is available when the MEF35 Mode parameter is selected. Configure the Test Family parameter as follows:

- Choose the Ethernet option to enable configuration of DMM, SLM, and LMM PM Session tests for Entity Types of VLL, VPLS, or Composite services.
- Choose the IP option to enable configuration of a TWAMP Light PM Session test for Entity Types of VPRN Service or Router.
- Choose the MPLS option to enable configuration of an MPLS DM Session Test for Entity Types of LSP or LSP (MPLS-TP). The LSP entity type includes RSVP and RSVP Auto LSPs.

You must select the NE Schedulable parameter if the test you are creating will be used in a test suite that requires an NE-scheduled policy.

The Strategy parameter in the Test Generation panel is available when the Entity Type parameter is set to VLL Service or VPRN Service.

The Types parameter in the Address Types panel is configurable when the Entity Type parameter is set to VPRN Service.

The Accounting Files parameter is configurable when the NE Schedulable parameter is selected.

The Continuously Executed and Test Result Storage parameters are configurable when the Accounting Files parameter is selected.

4

Click on one of the following:

a. Test Definitions tab. This tab is available when the MEF35 Mode parameter is not selected. Use this tab to add all OAM test definitions to a test policy except for PM Session OAM test types.

b. PM Session Test Definitions tab. This tab is available when the MEF35 Mode parameter is selected. Use this tab to add all PM Session OAM test definitions to a test policy.

5

Add one or more test definitions to the policy.

1. Click Add and choose an option from the cascading menu, for example, Services→Add Site Ping or PM Session Test→Add CFM SLM Session Test. The selected test definition form opens.

   The available options and test categories vary, depending on the settings specified for the Entity Type, NE Schedulable, or other parameters in Step 3.

2. Configure the required parameters.

   The available parameters depend on the type of test you chose in 1. See the XML API Reference to view information about configurable parameters and their applicability.

3. Click on each of the remaining tabs on the selected test definition creation form, for

instance: Test Parameters, Results Configuration, CFM Details, Bin Group Details, Measurement Interval Details, and so on. Configure the required parameters on each tab. The available tabs vary, depending on the type of test you chose in 1.

4. As required, configure threshold-crossing alarms for the test definition by performing Step 4 to Step 9 of 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949).

   Note: Configuration of threshold-crossing alarms is optional, and can be performed either during test definition creation, or later; see 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949).

6

Save your changes and close the form. The Test Policy (Create|Edit) form is displayed, with test definitions listed on the Test Definitions tab.

7

CAUTION

Service Disruption

Do not modify the test policy while a scheduled task for the suite is enabled

The operation of scheduled test suites that use the test policy may be adversely affected if you modify the test policy while a scheduled task for the suite is enabled.

If you are configuring an existing test policy, click Update Test Suites to apply the test policy changes to all associated test suites.

You can click on the Usages tab to view a list of the test suites that use the test policy.

8

Save your changes and close the forms.

END OF STEPS

89.11 To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy

89.11.1 Purpose

Perform this procedure to configure threshold-crossing alarms or NM threshold-crossing alarms as a post-STM OAM diagnostic test or STM test policy creation exercise.

You can configure threshold-crossing alarms or NM threshold-crossing alarms for OAM diagnostic tests based on the following requirements:

- within an discrete OAM diagnostic test or for a test definition within an STM test policy
- for tests that are NE-schedulable or non-NE-schedulable

An alarm is raised when a threshold is crossed, either because the value rose above or fell below the configured level.

Note: The NM threshold-crossing alarms are not configurable for service site ping, VPRN trace, MTU ping, Mrinfo and Mtrace tests. The type of NM threshold-crossing event available for selection depends on the test type.

Note: Configuring threshold-crossing alarms on a test definition within the test policy creates a threshold definition that applies to generated tests. Configuring threshold-crossing alarms directly on a test that is NE-schedulable applies just to the test.

Note: If you modify a test policy associated with a test suite, you must regenerate the generated tests in the test suite.

89.11.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Perform one of the following:

a. Choose Test (Assurance) from the object menu to configure threshold alarms for a discrete OAM diagnostic test.

b. Choose Test Policy (Assurance) from the object menu to configure threshold alarms in a test definition within a test policy.

3

Click Search and perform one of the following:

a. Choose an NE-schedulable or non-NE-schedulable test and click Properties. The Test (Edit) form opens.

b. Choose an NE-schedulable or non-NE-schedulable policy and click Properties. The Test Policy (Edit) form opens.

4

Perform one of the following:

a. Click on the Thresholds tab for a discrete OAM diagnostic test.

b. Click on the Test Definition tab for a test definition within a test policy.

5

Perform one of the following:

a. Click Create for a discrete OAM diagnostic test. The NE Threshold Event (Create) form opens.

b. Click Add for a test definition within a test policy. The NE Threshold Event (Create) form opens.

6

Configure the required parameters.

The Clear Alarm on Falling Threshold parameter is configurable when the Include Falling Threshold is enabled.

7

Click on the Rising Threshold tab and configure the Threshold Value parameter.

8

Click on the Falling Threshold tab and configure the Threshold Value parameter.

Note: The Falling Threshold tab can be accessed when the Include Falling Threshold parameter is enabled on the General tab of the NM Threshold Event (Create) form.

9

Click OK to save your changes and close the form.

10

Repeat Step 4 to Step 9 to configure threshold events on additional tests.

11

Select a test and click Execute. A threshold-crossing alarm appears on the dynamic alarm list if the threshold rises above or falls below the configured level.

12

Save your changes and close the forms.

END OF STEPS

89.12 To create an STM test suite

89.12.1 Before you begin

As required, review the 89.2.4 “STM test suite design considerations” (p. 2907) information in this chapter before starting this procedure.

Before performing this procedure, you must configure the STM test policy and test definitions to be included in the STM test suite.

89.12.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create and choose Test Suite. The Test Suite (Create) form opens.

3

Configure the required parameters.

The MEF35 and NE Schedulable parameters are mutually exclusive. The MEF35 parameter must be selected if the test suite will contain a test policy that uses PM Session Test Definitions. The NE Schedulable, Validation Test Suite, and Timeout parameters are configurable when the MEF35 Mode parameter is not selected.

You must select the NE Schedulable parameter if the test suite you are creating will use an NE-schedulable test policy.

The Validation Test Suite parameter is configurable when the Entity Type parameter is set to a type that supports validation, for example a VLL service.

The First-run Execution Sequence and Last-run Execution Sequence parameters are configurable when the MEF35 Mode parameter is not selected.

The Accounting Files parameter is configurable when the NE Schedulable parameter is selected.

The Continuously Executed and Test Result Storage parameters are configurable when the Accounting Files parameter is selected.

4

If you chose None as the option for the Entity Type parameter in Step 3, go to Step 8. Otherwise go to Step 5.

5

Add a test policy to the test suite.

1. Click on the Test Policy tab, then click Add.
2. Choose a test policy and click OK. The policy is listed on the Test Policy tab.

   The available policies depend on the option specified for the Entity Type parameter in Step 3.

6

Define the network entities for the test suite.

1. Click on the Tested Entities tab, then click Add.

2. Choose one or more entities and click OK. The selected entities are listed on the Tested Entities tab.

   The available entities depend on the option specified for the Entity Type parameter in Step 3.

   The Test Suite Count column in the Search form indicates the number of configured test suites that a specific test entity is currently associated with. The counter includes test suites that do not contain generated tests.

3. Select any required entities and click Lock Test Generation.

   Locking the test generation means that the initial settings you specify when creating the test suite will be retained, even if later on you change some attributes (for instance, if you are testing SDP tunnels and you extend the tested entities by selecting another SDP tunnel). All customized test attributes are retained when re-generating the tests.

   The NFM-P auto-locks tested entities when generated tests are modified or deleted.

   The Test Re-generation locked field in the list displays a check mark to indicate that the lock is enabled, and the name of the button changes to Unlock Test Generation. If you do not want to lock the test, click Unlock Test Generation. The Test Re-generation locked field in the list will not display a check mark.

7

If you selected the MEF35 Mode parameter in Step 3, go to Step 12.

8

Define first-run tests.

1. Click on the First-run Tests tab, then click Add.
2. Choose one or more tests and click OK.
3. Reorder the first-run test execution sequence, if required. When there is more than one test in the first-run list, the Move Up and Move Down buttons are available. Select a test entry and use the buttons to reorder the tests.

9

Define last-run tests.

1. Click on the Last-run Tests tab, then click Add.
2. Choose one or more tests and click OK.
3. Reorder the last-run test execution sequence, if required. When there is more than one test in the last-run list, the Move Up and Move Down buttons are available. Select a test entry and use the buttons to reorder the tests.

10

Click Apply to save the changes. Additional buttons become available at the bottom of the form.

11

If you chose None as the option for the Entity Type parameter in Step 3, go to Step 16. Otherwise, go to Step 12.

12

Click on the Generated Tests tab or the Generated PM Sessions Tests tab, as applicable.

13

Click Generate Tests.

The Generate Tests button is available only when you have a policy and at least one entity associated with the STM test suite.

The NFM-P begins generating tests for the test suite based on the test policy. The tests appear in a list on the Generated Tests form as they are generated. The name of the originating tested entity for each test is also displayed on the list, for quick reference.

Note: When the test policy or the configuration of an entity on the Tested Entities tab changes, you must regenerate the tests in the test suite by clicking Generate Tests.

Note: When tests have been generated in a test suite which has a test policy already associated with it, these tests will not be altered (deleted and recreated, or set to default tests), when Generate Tests is clicked again.

Note: Any of the tests generated by a test suite can be deleted directly from the list on the Generated Tests tab. Select the tests that you want to delete and click Delete.

Note: For a test suite containing a P2MP LSP Ping Test Policy, the test suite generator will generate one test per S2L path (for P2MP LSP) or one test per leaf node address (for P2MP LDP).

14

Click on the Generation Error Logs tab to view the log files that the NFM-P creates before and during test generation. The Stage drop-down allows you to select the type of logs to examine, including pre-generation and deployment information.

Note: Generation Error Logs messages appear only when failures occur during the test generation process.

15

Click on the Generated Artifacts tab to view a list of all objects created at the time of Test Suite creation or while generating the contained tests. Objects include: MDs, MAs, Global MEGs, MEPs, Bin Groups, tests (including SAA and PM Session), and accounting policies. The Stage drop-down allows you to select the type of information to examine.

Note: The Generated Artifacts tab will not show objects that already existed prior to generating the tests and are just being reused by those tests. Only newly-created objects are displayed.

Note: If objects are manually deleted, they will also be removed from the Generated Artifacts list. For example, if a MEP is deleted, the MEP object and the tests originating from the MEP will all be deleted from the list. If a GMEG is deleted, then all MAs, MEPs, and tests associated with the GMEG will also be deleted from the list.

16

Close the forms. See 89.17 “To execute an STM test suite” (p. 2961) for information about executing a test suite.

END OF STEPS

89.13 To modify an STM test suite and view additional information

89.13.1 Steps

1

You can modify an STM test suite using two methods. Perform one of the following:

a. Modify the test suite using the STM. Go to Step 2.

b. Add an entity to the test suite from the entity to be tested. Go to Step 9.

2

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

3

Choose Test Suite (Assurance) from the object drop-down menu and click Search. A list of test suites appears.

4

Choose a test suite and click Properties. The Test Suite (Edit) form opens.

5

Configure the required parameters on the General tab.

Some parameters are configurable only during test suite creation, and cannot be modified for existing test suites.

6

Modify the components of the test suite.

The following tabs list components of a test suite. Add, remove, or modify components as required.

- Test Policy tab — displays the test policy that is associated with the test suite
- Tested Entities tab — lists the entities that are the subjects of the tests in the test suite

- First-run Tests tab — lists the group of tests in the test suite that run first
- Generated Tests tab — lists the generated tests that are included in the test suite
- Last-run Tests tab — lists the group of tests in the test suite that run last

To configure a test suite component list item, select the item and click Properties.

7

View additional information.

The following tabs display information about the creation and execution of the test suite:

- Results tab — lists the per-run results that are returned by each execution of the test suite
- Individual Test Results tab — lists all individual test results. This is similar to what is displayed when a Tested Entity is chosen and you click on Show Results.
- Generation Error Logs tab — lists the log entries that are created before and during test generation for the test suite. The Stage drop-down menu allows you to select the type of logs to examine, including pre-generation and deployment information.
- Generated Artifacts tab - Lists all objects created at the time of Test Suite creation or while generating the contained tests. Objects include: MDs, MAs, Global MEGs, MEPs, Bin Groups, tests (including SAA and PM Session), and accounting policies. The Stage drop-down menu allows you to select the type of information to examine.
- Faults tab — displays the faults associated with the test suite

Note: Generation Error Logs messages only appear if any failures occur during the test generation process.

Note: The Generated Artifacts tab will not show objects that already existed prior to generating the tests and are just being reused by those tests. Only newly-created objects are displayed.

Note: If objects are manually deleted, they will also be removed from the Generated Artifacts list. For example, if a MEP is deleted, the MEP object and the tests originating from the MEP will all be deleted from the list. If a GMEG is deleted, then all MAs, MEPs, and tests associated with the GMEG will also be deleted from the list.

Note: You can use the Quick Filter on the test result tab pages to narrow down and limit the displayed results.

8

Save your changes and close the forms.

9

Open the Properties form of a supported entity that you want to add to a test suite; for example, a VLL or VPRN service or VPLS.

10

Click on the Tests tab. The Test Suite tab is displayed.

11

Click Add. The Add form appears.

12

Click Search to display a list of test suites. Only test suites applicable to this type of entity will be displayed.

13

Choose a test suite and click OK. The Add form closes and the selected test suite appears in the Test Suite tab.

The entity has now been added to this test suite. You can confirm this by clicking Properties to display the Test Suite (Edit) form. The added entity will appear on the list in the Tested Entities tab. You can also modify any other parameters of the test suite here, as detailed in Step 2 to Step 6 above.

Note: Once you have added the entity, you can also run the test suite from the Test Suite sub-tab by clicking Execute. See 89.17 “To execute an STM test suite” (p. 2961) for information about executing a test suite.

END OF STEPS

89.14 To configure OAM diagnostic test limits on the STM and view additional test configuration information

89.14.1 Purpose

Using the STM, you can optionally configure some limits for OAM diagnostic tests. For example, you can specify the maximum number of pings and traces performed on managed devices during the execution of a test, or you can limit the number of tests that can be performed. The default is to allow an unlimited number of tests.

You can also perform this procedure to see additional information about OAM tests configured on managed devices, such as deployment and schedule information.

Note: Limiting the number of tests does not raise an alarm on the NFM-P, or an SNMP trap on the managed device. The indication that the limit is reached is that no further test results are returned.

89.14.2 Steps

1

Ensure that the Debug STM Mode parameter on the NFM-P User Preferences form is selected. Choose Application→User Preferences from the NFM-P main menu, or see 1.23 “To configure NFM-P user preferences” (p. 116) for more information.

The drop-down options specified in this procedure are available only when Debug STM Mode is enabled.

2

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

3

Choose one of the following options from the object menu.

a. NE Test Agent (Assurance). You can specify the maximum number of pings and traces allowed to be performed on individual managed devices during the execution of a test, or limit the number of tests that can be performed on all managed devices. This option also provides information about the node ID, concurrent ping and trace information, LTT concurrent requests, deployed test counts, and alarm status.

b. Deployed Test (Assurance). You can view additional deployment information about each configured OAM diagnostic test on managed devices such as the node ID, name and description, management ownership, deployment and execution state, schedule status and start time, and alarm status.

c. NE Schedulable Test (Assurance). You can view additional information about NE schedulable tests on managed devices, such as the node ID, name and description, management ownership, test deployed, mode, state, runs and failures, accounting policy ID, accounting suppression, card TCA profile, and alarm status.

4

Click Search. A list of policies or tests appears, depending on the option chosen in Step 3.

5

Choose an item and click Properties. A form opens showing the current configuration and information.

6

As required, view the deployment information, or change the test limits for the managed device to meet your operational requirements.

7

Save your changes and close the forms.

END OF STEPS

89.15 To run one or more OAM diagnostic tests from the STM and view the test results

89.15.1 Purpose

You can run a single OAM diagnostic test from the properties form associated with the test, or from the Service Test Manager (STM) form. To run multiple tests simultaneously, or to add them to an STM test suite, you must use the Service Test Manager (STM) form.

You can view the test results immediately after the test is complete, as outlined in this procedure, or you can view and compare the results of multiple tests, as outlined in 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960).

89.15.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click on the object drop-down and expand the Test (Assurance) menu item until the required test type appears.

3

Choose the test type from the menu and click Search. A list of pre-configured tests is displayed.

4

Run the tests. Choose one or more tests in the list and click Execute.

Note: You can select and execute multiple tests from the test list, except for CCM Tests. If multiple tests including CCM tests are selected, the Execute button is not available.

5

View the test results.

1. Select a test in the list and click Properties. The Test (Edit) form opens.
2. Click on the Results tab. A list of test results is displayed.

6

Archive the results, if required. Select one or more of the test results and click Archive.

Archived test results are stored indefinitely. They can only be removed by manually selecting and deleting them.

7

Close the forms.

END OF STEPS

89.16 To view and compare OAM diagnostic test results on the STM

89.16.1 Steps

Perform this procedure to view the results of multiple OAM diagnostic tests accessible on the STM. You can also use this procedure to compare two test results from the same type of test.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose the required test result type. Click on the object drop-down and expand one of the following menu item until the required test results type appears.

a. Test Result (Assurance) — contains the test results from almost all test types that have been run.

b. Test Suite Result (Assurance) — contains the test results from all STM test suites that have been run.

c. Archived Test Results (Assurance) — contains the test results from all individual tests that have been run and archived.

d. One Time Validation Results (Assurance) — contains the test results from all One Time Validation tests that have been run.

3

Choose the required test results type and click Search. The list of test results is displayed.

4

Choose a test result and click Properties. The test results form opens.

5

Click on the tabs to view the test results. You can click View beside an object to open the properties form for that object. Alternatively, click View Test, or View Test Object, to view the required information.

6

To compare results from two tests of the same type, perform one of the following:

a. Choose two test results and click Compare. The Compare form opens. Go to Step 8.

b. Choose one test result and click Compare. This allows you compare the selected test result with a test result from a test results archive. The Compare form opens.

c. Click Compare without choosing any test results. This allows you to select two test results for comparison from a test results archive. The Compare form opens.

The Compare form displays two object fields at the top. If you performed Step 6 a, then these two fields will be populated with the names of the chosen tests. If you performed either of Step 6 b or Step 6 c, then one or both of the fields will be blank.

7

Select the test results to assign as Object A and Object B, as required.

1. Click Select beside the required field. The Select Object form opens.
2. Choose the Test Result type from the drop-down selector and click Search. You can choose a regular or archived test result. A list of test results is displayed.
3. Choose a test result and click OK. The Select Object form closes and the Compare form displays the selection.

   When one of the Object fields is populated and you want to populate the other, then the list displays only the required type of test results, either regular or archived.

8

Click Compare. The comparison is performed and an entry appears in the list.

The Swap button allows you to interchange the names in the Object fields without having to reselect the test results. This is useful when organizing the order in which results are compared in the list.

9

Choose the entry and click Properties. The Difference form is displayed, with the differences between specific properties listed. The Value A column shows the property value in Object A and the Value B column shows the value of the same property in Object B.

The names of the test (Class Name) and test objects are displayed in the field on the right.

10

View the information and close the forms.

END OF STEPS

89.17 To execute an STM test suite

89.17.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test Suite (Assurance) from the object menu.

3

Click Search, then choose a test suite in the list and click Properties. The Test Suite (Edit) form opens with the General tab displayed.

4

Click Execute. The test suite begins to run. You can view the test results when the test is complete.

Note: You can run an OAM validation test from the properties form of a service tunnel or an applicable service. See the applicable service management chapter for more information.

Note: When you execute a test suite from the properties form of an applicable service, only the tests belonging to that service will be executed. Tests within the suite, but not associated with the current tested entity, will not be executed.

Note: When you execute a test suite under the tested entity (by clicking Execute under the entity's Tests>Test Suite sub-tab), NFM-P will create a Tested Entity Result under the entity's Tests>Tested Entity Result sub-tab.

Note: You can click Delete All Result in either the Results tab or the Individual Test Results tabs to remove all listed test results. Alternatively, you can use the Quick Filter on the test result tab pages to narrow down and limit the displayed results. Clicking Delete All Result will then only delete the filtered results. If no filter condition is specified, then clicking Delete All Result will delete all the pages of results, not just the currently displayed page.

Note: If the list under either of these tabs contains more than 1000 entries, only 1000 entries are displayed at a time, with paging available for the remainder. However, clicking Delete All Result deletes all pages of results, not just the currently displayed page.

5

Close the Test Suite (Edit) form.

END OF STEPS

89.18 To view STM test suite results

89.18.1 Purpose

Perform this procedure to view STM test suite results from the Test Suite Result form.

Note: You can also view the test suite results for an object from the Tests tab of the configuration form for the object.

89.18.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test Suite Result (Assurance) from the object menu.

3

Click Search, then select a test suite result and click Properties. The Test Suite Result form opens.

4

View the information and statistics on the General tab.

5

To view test results, click on the Results tab and perform one of the following:

a. Use the horizontal scroll bar at the bottom of the list to view the tabular information about the listed tests.

b. Choose an entry from the list and click Properties. The results form for the selected test opens.

6

Close the forms.

END OF STEPS

89.19 To view and compare STM test suite results for a tested entity

89.19.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test Suite (Assurance) from the object menu and click Search. A list of test suites is displayed.

3

Choose a test suite in the list and click Properties. The Test Suite (Edit) form opens.

4

Click on the Tested Entities tab to view a list of entities tested by the test suite.

Note: The Test Generation State column displays the overall test state of each entity. The possible values are:

- Not Generating — No test generation is applicable; for example, the test suite type is set to none.
- No Test Generated — Generation did not produce any tests.
- Policy Distribution — Required policies are being distributed in preparation for test generation.
- Pre-Generation — Prerequisite tasks and object configurations are being performed.
- Test Generation — Tests are being generated.
- Generation Complete — Test suite has successfully generated.
- Generation Failed — Test generation failed; test suite configuration should be reviewed.

5

Choose a tested entity and click Show Results. The Tested Entity (Edit) form opens.

6

Click Search, then choose a test result and click Properties. The results form for the selected test opens.

7

To compare two test results from the same type of test, choose the two test results and click Compare. The Difference form opens.

8

Close the forms.

END OF STEPS

89.20 To interpret OAM diagnostic test results on the STM

89.20.1 Steps

1

Perform the appropriate OAM diagnostic test. See Table 90-1, “NFM-P supported OAM diagnostic tests and configurations” (p. 2981) for a list of supported OAM diagnostic tests and their applicable procedures.

2

Perform 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960) to view OAM diagnostic test results.

3

The General and Response Packets tabs display the key information about OAM diagnostic results.

When no packet information is required, as for an OAM ping, the information appears on the Packets Results form without any tabs. For example, status and return code information appears directly on the Packets Results form.

You can view test packet information by performing the following:

1. Click on the Response Packets tab. A list of test objects appears.
2. Choose an object and click Properties. The Packets Results form opens.

4

Interpret the results, based on the status and return code information.

a. For MTU OAM diagnostics, the key information is how many frames were sent and incrementally increased in size before the frames could not be sent.

   When the frame cannot be sent because it is too large, that results in a request timeout message. The largest frame that was sent is the last frame size with an associated success response.

b. For tunnel OAM diagnostics, the key information is the result of the diagnostic, displayed in the status message. The following table lists the displayed messages and descriptions.

Table 89-13 Tunnel OAM diagnostics results

Displayed message	Description
Request Timeout	The request timed out with a reply.
Orig-SDP Non-Existent	The request was not sent because the originating SDP does not exist.
Orig-SDP Admin-Down	The request was not sent because the originating SDP administrative state is operationally down.
Orig-SDP Oper-Down	The request was not sent because the originating SDP operational state is down.
Request Terminated	The operator terminated the request before a reply could be received or before the timeout of the request could occur.
Far End: Originator-ID Invalid	The request was received by the far end, but the far end indicates that the originating SDP ID is invalid.
Far End: Responder-ID Invalid	The request was received by the far end, but the responder ID is not the same destination SDP ID that was specified.
Far End:Resp-SDP Non-Existent	The reply was received, but the return SDP ID used to respond to the request does not exist

Table 89-13 Tunnel OAM diagnostics results (continued)

Displayed message	Description
Far End:Resp-SDP Invalid	The reply was received, but the return SDP ID used to respond to the request is invalid.
Far End:Resp-SDP Down	The reply was received, but the return SDP ID indicates that the administrative or operational state of the SDP is down.
Success	The tunnel is in service and working as expected. A reply was received without any errors.

c. For service site OAM diagnostics, the key information is the result of the diagnostic, which is displayed in the status message and with the other records from the General tab.

As the diagnostic traverses the service across the originating and destination IP addresses, the service tunnels, and the used VCs, the status of each portion of the service is displayed. The following table lists the displayed messages and descriptions.

Table 89-14 Service site OAM diagnostics results

Displayed message	Description
Sent - Request Timeout	The request timed out with a reply.
Sent - Request Terminated	The request was not sent because the diagnostic was terminated by the operator.
Sent - Reply Received	The request was sent and a successful reply message was received.
Not Sent - Non-Existent Service-ID	The configured service ID does not exist.
Not Sent - Non-Existent SDP for Service	There is no SDP for the service being tested.
Not Sent - SDP For Service Down	The SDP for the service is down.
Not Sent - Non-Existent Service Egress Label	There is a service label mismatch between the originator and responder.

d. For MAC, VPRN, and multicast FIB ping OAM diagnostics, the key information is the result of the diagnostic. The following table lists the return codes and descriptions.

Table 89-15 MAC, VPRN, multicast FIB ping OAM diagnostics results

Displayed message	Description
notApplicable (0)	The OAM diagnostic message does not apply to the OAM diagnostic performed.
fecEgress (1)	The replying router is an egress for the FEC.
fecNoMap (2)	The replying router has no mapping for the FEC.
notDownstream (3)	The replying router is not a downstream router.
downstream (4)	The replying router is a downstream router, and the mapping for this FEC on the router interface is the specified label.

Table 89-15 MAC, VPRN, multicast FIB ping OAM diagnostics results (continued)

Displayed message	Description
downstreamNotLabel (5)	The replying router is a downstream router, and the mapping for this FEC on the router interface is not the specified label.
downstreamNotMac (6)	The replying router is a downstream router, but it does not have the specified MAC address.
downstreamNotMacFlood (7)	The replying router is a downstream router, but it does not have the specified MAC address and cannot flood the request to other routers.
malformedEchoRequest (8)	The received echo request is malformed.
tlvNotUnderstood (9)	One or more TLVs were not understood.
downstreamNotInMFib (10)	The replying router is a downstream router, but it is not part of the MFIB.
downstreamMismatched (11)	The downstream mapping is mismatched.
upstreamIfIdUnkn (12)	The upstream interface index is unknown.
noMplsFwd (13)	The label switched successfully but MPLS forwarding did not occur at stack-depth.
noLabelAtStackDepthh (14)	The label entry at stack-depth did not occur.
protoIntfMismatched (15)	The protocol is not associated with the interface at FEC stack-depth.
terminatedByOneLabel (16)	The ping was terminated prematurely due to the label stack shrinking to a single label.
seeDDMapForRetCodeSubCode (17)	See DDMap TLV for the return code and sub-return code.
fecStackChange (18)	The label switched with an FEC stack change.

e. For multicast router OAM diagnostics, the key information includes information that is related to adjacent routers, supported protocols, traffic metrics, and time-to-live thresholds. Administrators can use this information to identify bidirectional adjacency relationships.

f. For multicast trace OAM diagnostics, the key information is the result of the diagnostic, displayed in the status message. The following table lists the displayed messages and descriptions.

Table 89-16 Multicast trace OAM diagnostics results

Displayed message	Description
noError (0)	No error.
wrongIf (1)	The router is not forwarding the multicast source or group traffic because the router is not part of the multicast path.
pruneSent (2)	The router sent a prune request upstream that impacts the multicast source and group in the trace request.

Table 89-16 Multicast trace OAM diagnostics results (continued)

Displayed message	Description
pruneRecvd (3)	The router stopped forwarding traffic for the multicast source and group in response to a request from the next hop router.
scoped (4)	The multicast group is subject to administrative scoping.
noRoute (5)	The router has no route for the multicast source or group and no way to determine a potential route.
wrongLastHop (6)	The router is not the correct last-hop router.
notForwarding (7)	The router is not forwarding the multicast source or group traffic for an unspecified reason.
reachedRP (8)	Request arrived on the rendezvous point or core.
rpfIf (9)	The request arrived on the expected RPF interface.
noMulticast (10)	The request arrived on an interface that is not enabled for multicast traffic.
infoHidden (11)	One or more hops are hidden from the trace.
fatalError (12)	Fatal error.
noSpace (129)	There is insufficient room to insert another response data block in the packet.
oldRouter (130)	The previous router hop cannot process the multicast trace request.
adminProhib (131)	The multicast trace was administratively prohibited.
unknown (132)	Unknown error.

g. For most diagnostics, common return codes are used. The return codes indicate the status of OAM tests, usually when there was a problem performing the test. The following table lists return codes and descriptions.

Table 89-17 OAM diagnostics return codes

Return code	Description
notApplicable (0)	A response was received on the device which is not applicable to the OAM diagnostic performed.
responseReceived (1)	A response to the OAM diagnostic was received on the device.
unknown (2)	The OAM diagnostic failed for an unknown reason.
internalError (3)	An internal error on the device caused the diagnostic to not be performed.
maxConcurrentLimitReached (4)	The device cannot perform the OAM diagnostics because there are too many OAM diagnostic operations already running.
requestTimedOut (5)	The OAM diagnostic could not be completed because no reply was received within the allocated timeout period.
unknownOrigSdpId (6)	Indicates an invalid or non-existent originating service tunnel.
downOrigSdpId (7)	The originating service tunnel is operationally down.
requestTerminated (8)	The OAM diagnostic was canceled before the timeout or reply period was reached.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
invalidOriginatorId (9)	The far-end device replied indicating that the originating ID was invalid.
invalidResponderId (10)	The far-end device replied with an invalid responding ID.
unknownRespSdpId (11)	The far-end device replied with an invalid response service tunnel ID.
downRespSdpId (12)	The responding service tunnel with the given ID is operationally or administratively down.
invalidServiceId (13)	An invalid or non-existent service ID.
invalidSdp (14)	An invalid or non-existent service tunnel for the service.
downServiceSdp (15)	The service tunnel bound to the service is down.
noServiceEgressLabel (16)	The egress label for the service does not exist.
invalidHostAddress (17)	The IP address for the host is invalid, for example, in the case of a broadcast or multicast IP address.
invalidMacAddress (18)	The MAC address specified in the OAM diagnostic is invalid.
invalidLspName (19)	The LSP name specified in the OAM diagnostic is invalid.
macIsLocal (20)	The MAC address is the local SAP or device MAC address, not the MAC address of the downstream SAP or device, therefore the MAC ping or trace cannot be sent.
farEndUnreachable (21)	No route is available to the far-end GRE service tunnel.
downOriginatorId (22)	The originating ping device is operationally down.
downResponderId (23)	The device responding to the ping is operationally down.
changedResponderId (24)	The ID of the device responding to the ping has changed.
downOrigSvcId (25)	The originating service identified by the ID is operationally down.
downRespSvcId (26)	The service responding to the ping identified by the ID is operationally down.
noServiceIngressLabel (27)	The ingress label for the service does not exist.
mismatchCustId (28)	The subscriber ID identified with the service differs between the originating device and the responding device.
mismatchSvcType (29)	The service type identified with the service differs from one device to another.
mismatchSvcMtu (30)	The service MTU size associated with the service differs between the originating device and the responding device.
mismatchSvcLabel (31)	The service label identified with the service differs between the originating device and the responding device.
noSdpBoundToSvc (32)	There is no service tunnel bound to the service.
downOrigSdpBinding (33)	The service tunnel associated with the originating device's service is operationally down.
invalidLspPathName (34)	The LSP path name specified in the OAM diagnostic is invalid.
noLspEndpointAddr (35)	There is no LSP endpoint address specified in the OAM diagnostic.
noActiveLspPath (36)	There is no active LSP path.
downLspPath (37)	The far end of the LSP is operationally down.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
invalidLspProtocol (38)	The LSP protocol is not valid or is not supported.
invalidLspLabel (39)	The LSP label is invalid.
routeIsLocal (40)	For a VPRN ping, the route is a local route.
noRouteToDest (41)	For a VPRN ping, there is no route available to the destination of the OAM diagnostic.
localExtranetRoute (42)	For a VPRN ping, the route is a local extranet route.
srcIpInBgpVpnRoute (43)	For a VPRN ping, the source IP address belongs to a BGP VPN route.
srcIpInvalid (44)	For a VPRN ping, the source IP address is invalid or no route is available to the source IP address.
bgpDaemonBusy (45)	For a VPRN trace, the BGP routing process is busy on the device, and VPRN route target information cannot be retrieved.
mcastNotEnabled (46)	Multicast is not enabled on the device, so the diagnostic cannot be performed.
mTraceNoSGFlow (47)	No (*,G)/(S,G) flow on the router. The multicast trace cannot be initiated.
mTraceSysIpNotCfg (48)	The system IP address is not configured. The address is required for a response to a multicast trace.
noFwdEntryInMfib (49)	No forwarding entry could be found for the specified source and destination addresses in the MFIB.
dnsNameNotFound (50)	The domain name specified in the dns query does not exist.
noSocket (51)	For icmp-ping, unable to get socket.
socketOptVprnIdFail (52)	For icmp-ping, unable to set SO_VPRNID for socket.
socketOptIfIndexFail (53)	For icmp-ping, unable to set IP_IFINDEX for socket.
socketOptNextHopFail (54)	For icmp-ping, unable to set IP_NEXT_HOP for socket.
socketOptMtuDiscFail (55)	For icmp-ping, unable to set IP_MTU_DISC for socket.
socketOptSndbufFail (56)	For icmp-ping, unable to set SO_SNDBUF for socket.
socketOptHdrincFail (57)	For icmp-ping, unable to set IP_HDRINCL for socket.
socketOptTosFail (58)	For icmp-ping, unable to set IP_TOS for socket.
socketOptTtlFail (59)	For icmp-ping, unable to set IP_TTL for socket.
bindSocketFail (60)	For icmp-ping, unable to bind socket.
noRouteByIntf (61)	For icmp-ping, no route to destination via the specified interface.
noIntf (62)	For icmp-ping, no interface specified.
noLocalIp (63)	For icmp-ping, unable to find local IP address.
sendtoFail (64)	For icmp-ping, the send to function failed.
rcvdWrongType (65)	For icmp-ping, received packet of wrong ICMP type.
noDirectInterface (66)	For icmp-ping, no direct interface to reach destination.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
nexthopUnreachable (67)	For icmp-ping, unable to reach the next-hop.
socketOptHwTimeStampFail (68)	For icmp-ping, unable to set IP_TIM_TIME for socket.
noSpokeSdpInVll (69)	For vccv-ping, unable to find spoke-sdp given SdpId:vc-id.
farEndVccvNotSupported (70)	For vccv-ping, far end does not support the VCCV options.
noVcEgressLabel (71)	For vccv-ping, no VC egress label to send vccv-ping
socketOptIpSessionFail (72)	For icmp-ping, unable to set IP_SESSION for socket.
rcvdWrongSize (73)	For icmp-ping, received packet of wrong size.
dnsLookupFail (74)	For icmp-ping, dns lookup failed.
noIpv6SrcAddrOnIntf (75)	For icmp-ping, no ipv6 source on the interface.
multipathNotSupported (76)	For lsp-trace, downstream NE does not support multipath.
nhIntfNameNotFound (77)	For lsp-ping/trace, given next-hop interface name not found.
msPwInvalidReplyMode (78)	For vccv-ping, MS-PW switching NE supports ip-routed reply mode only.
ancpNoAncpString (79)	ANCP string unknown to the system.
ancpNoSubscriber (80)	Subscriber unknown to the system.
ancpNoAncpStringForSubscriber (81)	Subscriber has no associated ANCP string.
ancpNoAccessNodeforAncpString (82)	No access NE is found for the given ANCP string.
ancpNoAncpCapabilityNegotiated (83)	ANCP capability not negotiated with the involved DSLAM.
ancpOtherTestInProgress (84)	Another ANCP test is running for this ANCP string.
ancpMaxNbrAncpTestsInProgress (85)	Maximum number of concurrent ANCP tests reached.
spokeSdpOperDown (86)	For vccv-ping, spoke-sdp is operationally down.
noMsPwVccvInReplyDir (87)	Switching NE in MS-PW with no VCCV support in echo reply direction.
p2mpLspNameOrInstInvalid (88)	P2MP LSP name or instance provided is not valid.
p2mpLspS2LPathDown (89)	LSP path to S2L is down.
p2mpLspS2LAddressInvalid (90)	One or more S2L addresses is not valid.
p2mpLspNotOperational (91)	P2MP LSP is operationally down.
p2mpLspTrMultipleReplies (92)	Probe returned multiple responses. Result may be inconsistent.
invalidMepId (93)	The user-configured MEP identifier is not valid.
multipleReplies (94)	More than one reply received, when only one was expected.
packetSizeTooBig (95)	Packet size is too big.
gtpPingError (96)	General GTP Ping error.
gtpPingRsrcUnavailable (97)	GTP Path management resource unavailable.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
gtpPingDupRequest (98)	Duplicate request for the same peer.
gtpPingCleanUpInProg (99)	GTP Path management clean up in progress.
invalidInterface (100)	The specified egress interface does not exist.
p2mpLspNotFound (101)	The P2MP-LSP not found given LDP-ID or sender-address.
ethCfmSlmInLoss (102)	Synthetic Loss Measurement (SLM) probe lost in transit from far-end node to local agent.
ethCfmSlmOutLoss (103)	SLM probe lost on transmit from local agent to far-end node.
ethCfmSlmUnacknowledged (104)	SLM probe lost but unable to identify loss reason.
spokeSdpFecNoBndFound (105)	Spoke-sdp-fec is invalid or has no associated SDP binding yet.
mtraceNotSupportedP2mp (106)	Mtrace not supported in base routing context when incoming interface is P2MP.
useFec129Parameters (107)	Specify FEC129 parameters for FEC129 PW instead of sdp-id:vc-id.
dnsServerUnexpectedResponse (108)	The nameserver received an unexpected response.
dnsServerResponseFormErr (109)	The nameserver responded with FORMERR.
dnsServerResponseServFail (110)	The nameserver responded with SERVFAIL.
dnsServerResponseNotImp (111)	The nameserver responded with NOTIMP.
dnsServerResponseRefused (112)	The nameserver responded with REFUSED.
sendFailUndefinedServiceId (113)	The OAM test cannot be performed because the associated service does not exist.
sendFailWrongServiceType (114)	The OAM test cannot be performed because the associated service is of the wrong type.
sendFailSubnettedService (115)	The OAM test cannot be performed on a service with a specified mac subnet length.
invalidRespServiceId (116)	Invalid or non-existent responder Service-ID.
adminDownOrigSdpBind (117)	SDP binding is administratively down on the originator side.
operDownRespSdpBind (118)	SDP binding is operationally down on the responder side.
adminDownRespSdpBind (119)	SDP binding is administratively down on the responder side.
sdpBindVcidMismatch (120)	SDP binding VC ID mismatch between originator and responder.
sdpBindTypeMismatch (121)	SDP binding type mismatch between originator and responder.
sdpBindVcTypeMismatch (122)	SDP binding VC type mismatch between originator and responder.
sdpBindVlanVcTagMismatch (123)	SDP binding VLAN VC tag mismatch between originator and responder.
adminDownOrigSvc (124)	Service on the originator side is administratively down.
adminDownRespSvc (125)	Service on the responder side is administratively down.
adminDownOrigSdpId (126)	The originating SDP-ID is administratively down.
adminDownRespSdpId (127)	The responding SDP-ID is administratively down.
mTraceSourceIsNotRemote (128)	The multicast trace route test cannot be performed because the source address is not remote.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
invalidVirtualRouterId (129)	The OAM test cannot be performed because the associated virtual router ID is invalid.
ldpPrefixIsLocal (130)	The OAM test cannot be performed because the associated LDP prefix is local to the system.
sourceIpIsNotLocal (131)	The OAM test cannot be performed because the associated source address is not local to the system.
nextHopIpIsLocal (132)	The OAM test cannot be performed because the associated next hop address is local to the system.
targetIpIsLocal (133)	The OAM test cannot be performed because the associated target address is local to the system.
invalidControlPlaneOption (134)	The OAM test cannot be performed because the control plane send or receive option is not allowed with the specified service.
iomRevisionNotSupported (135)	The OAM test cannot be performed due to a mismatch in supported revision with the provisioned IOMs.
invalidSourceMacOption (136)	The OAM test cannot be performed because the source MAC option is not allowed with the specified service.
sendFailSpbMgdService (137)	The OAM test cannot be performed on a service which is managed by SPB.
useStaticPwParameters (138)	The spoke SDP is configured for static PW.
type1Fec129PwNotSupported (139)	The OAM test is not supported on type 1 FEC129 pseudowires.
mplsTpLspPathNotOperational (140)	The OAM test cannot be performed because the MPLS-TP LSP path is not operational.
invalidStaticMplsTpLsp (141)	The OAM test only supports static MPLS-TP LSPs.
controlWordNotValid (142)	The control-word for the spoke-SDP is not valid for this OAM test.
pwPathIdNotConfigured (143)	The pw-path-id provisioning is not complete.
notSupportedOnVcSwitchService (144)	The OAM test is not supported on a VC-switching service.
sdpFarEndNotSupported (145)	The OAM test does not support the SDP far-end value.
mplsTpLspPathShutdown (146)	The MPLS-TP path is currently shut down.
forceOptionIsBlocked (147)	The force option is currently blocked because of the configuration of a related entity.
intfForLspPathIsNotOperational (148)	The interface for the LSP path is not operational.
ttlExpired (149)	The destination could not be reached because the time-to-live (IPv4) or hop limit (IPv6) was too small. This results from an ICMPv4 type 11 code 0 or ICMPv6 type 3 code 0 message.
networkUnreachable (150)	The network specified by the destination address is unreachable. This results from an ICMPv4 type 3 code 0 or ICMPv6 type 1 code 0 message.
hostUnreachable (151)	The host specified by the destination address is unreachable. This results from an ICMPv4 type 3 code 1 or ICMPv6 type 1 code 3 message.
bgpLabelPrefixIsLocal (152)	The OAM test cannot be performed because the associated BGP Label Route prefix is local to the system.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
bgpLabelPrefixUnknown (153)	The OAM test cannot be performed because the BGP target FEC prefix entry is not found in the Routing Table.
ldpPrefixUnknown (154)	The OAM test cannot be performed because the LDP target FEC prefix entry is not found in the routing table.
l2tpv3DeliveryTypeUnsupported (155)	The OAM test cannot be performed because the L2TPv3 delivery type is unsupported.
vPingPeerCvNoLspPing (156)	The OAM test cannot be performed because the peer CV bits do not support LSP ping.
vPingPeerCcNoCtrlWord (157)	The OAM test cannot be performed because the peer CC bits do not support control word.
sendFailEvpnCfgdServiceX (158)	The OAM test cannot be performed on a VPLS service that has EVPN configured.
sendFailEvpnCfgdServiceX (158)	The OAM test cannot be performed on a VPLS service that has EVPN configured.
sendFailed (159)	The OAM test cannot send the test packet.
minimumPacketSizeNotMet (160)	The OAM test cannot send the test packet as the minimum required packet size exceeds the user supplied packet size.
invalidTargetFecType (161)	The OAM test does not support this FEC type.
p2mpLspPingNotSupportedOnMgmtRtr (162)	The OAM p2mpLspPing test is not supported on the management router.
ipv4SdpFarEndsOnly (163)	The OAM test only supports bindings and SDPs using an IPv4 far end.
vxlanEgrBndSvcMismatch (164)	The OAM test cannot send packets to a VXLAN egress binding not owned by the same service as the one sending.
vxlanNoMatchingTep (165)	The OAM test can only send to VTEPs that have been configured.
vxlanEvpnUnconfigured (166)	The OAM test can only send on a service with EVPN configured.
ipv6SdpFarEndsNotSupported (167)	The OAM test does not support bindings and SDPs using an IPv6 far end.
oamTestOverSRTunNotSupported (168)	The OAM test is not supported on SR tunnels.
sendFailEvpnCfgdPbbService (169)	The OAM test cannot be performed on a VPLS/Epipe service that is associated with a b-vpls that has EVPN configured.
txPortDown (170)	The transmit port is operationally down.
noTxPort (171)	No transmit port.
parentAdminDown (172)	The parent (SPOKE, SAP, SVC or PORT) of the MEP is admin shutdown.
destMacResolveFail (173)	Unable to resolve the remote-mepid to a unicast layer2 MAC address.
vxlanIpV6TermUnsupported (174)	The OAM test does not support IPv6 tunnel termination points.
ipPrefixIsLocal (175)	The OAM test cannot be performed because the associated IP prefix is local to the system.
ipPrefixUnknown (176)	The OAM test cannot be performed because the Target FEC prefix entry is not found in the Routing table.
greEthBrdgdDelvryTypeUnsupported (177)	The OAM test cannot be performed because the GRE Ethernet Bridged delivery type is unsupported.

Table 89-17 OAM diagnostics return codes (continued)

Return code	Description
mtrace2Disabled (178)	The mtrace2 test cannot be performed because mtrace2 is disabled at the system level.

END OF STEPS

89.21 To edit an OAM diagnostic test

89.21.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test (Assurance) or PM Session Test (Assurance) from the object menu and click Search. A list of tests is displayed.

3

Choose a test and click Properties. The properties form for the test opens.

4

Configure the required parameters.

5

Save your changes and close the forms.

END OF STEPS

89.22 To delete an OAM diagnostic test

89.22.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test (Assurance) or PM Session Test (Assurance) from the object menu and click Search. A list of tests is displayed.

3

Choose one or more tests and click Delete.

4

Confirm the action. The test or tests are deleted from the list of available OAM diagnostic tests.

5

Close the form.

END OF STEPS

89.23 To delete an STM test suite

89.23.1 Before you begin

If the STM test suite is scheduled to run, you must remove the scheduled task associated with the test suite before you can delete the test suite. See Chapter 5, “NFM-P-based schedules” for more information.

89.23.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Test Suite (Assurance) from the object menu and click Search. A list of test suites is displayed.

3

Select a test suite and click Delete.

4

Confirm the action. The test suite is deleted and removed from the list.

5

Close the form.

END OF STEPS

90 OAM diagnostic tests

90.1 Overview

90.1.1 Purpose

This chapter describes the NFM-P OAM diagnostic tests, and provides information about configuring and performing OAM diagnostic tests.

90.1.2 Contents

OAM diagnostic tests

90.2 OAM diagnostic test overview

90.2.1 Introduction to OAM diagnostic tests

The NFM-P provides a set of configurable in-band or out-of-band, packet-based OAM diagnostic tests that allows you to pro-actively manage networks and customer SLAs including:

- on-demand or scheduled service performance monitoring, and SLA verification to ensure that a service meets its guaranteed performance settings in a controlled test time
- networking troubleshooting activities including fault detection, fault verification, and fault isolation

This chapter provides information about each OAM diagnostic test that can be configured using the NFM-P. See Chapter 89, “Service Test Manager” for information about higher-level functions that can be performed with each OAM diagnostic test using the STM, such as creating STM test policies and STM test suites.

The correct delivery of services requires that a number of operations must occur correctly at different levels in the service. For example, operations such as the association of packets to a service, VC labels to a service, and each service to a service tunnel, must be performed successfully for the service to pass traffic to subscribers according to their SLAs. Even when tunnels are operating correctly and are correctly bound to services, incorrect information may cause connectivity issues.

For in-band, packet-based testing, the OAM packets closely resemble customer packets to effectively test the forwarding path. However, these packets are distinguishable from customer packets, so they are kept within the service provider network and not forwarded to the customer. For out-of-band testing, OAM packets are sent across a portion of the transport network; for example, across LSPs to test reachability.

90.2.2 On-demand OAM diagnostic test configuration

You can create, manage, and run most OAM diagnostic tests on-demand from the NFM-P STM as a single test entity.

Some OAM diagnostic tests can be configured from alternate launch points on the NFM-P GUI, including the following:

- from a network object or NE Properties form; for example, LSP pings from the LSP Properties form
- from a service tunnel or service form
- from the service topology or composite service flat topology view on service flat topology maps; see Chapter 4, “Topology map management” for information.
- using a user-defined CLI script; for example, creating a device-specific ping and traceroute script for troubleshooting issues on an OmniSwitch

The following table lists the OAM diagnostic tests that are accessible from the NFM-P GUI and provides a link to the appropriate test description and applicable procedure.

Table 90-1 NFM-P supported OAM diagnostic tests and configurations

OAM test category / test type	Test network object or service level	Network level	See Procedure
Ethernet CFM OAM diagnostic tests			
91.2.2 “Global MEG check” (p. 3089)	Ethernet	Port	91.19 “To create and run a Global MEG OAM diagnostic test from the STM” (p. 3120)
91.2.3 “Continuity check” (p. 3089) (includes 91.2.4 “CFM dual-ended loss test” (p. 3089))			91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122)
91.2.5 “CFM loopback” (p. 3090)			91.21 “To create and run a CFM loopback OAM diagnostic test from the STM” (p. 3123)
91.2.6 “CFM link trace” (p. 3090)			91.22 “To create and run a CFM link trace OAM diagnostic test from the STM” (p. 3125)
91.2.7 “CFM Eth test” (p. 3090)			91.23 “To create and run a CFM Eth OAM diagnostic test from the STM” (p. 3127)
91.2.8 “CFM two way delay test” (p. 3090)			91.24 “To create and run a CFM two way delay OAM diagnostic test from the STM” (p. 3128)
91.2.9 “CFM one way delay test” (p. 3091)			91.25 “To create and run a CFM one-way delay OAM diagnostic test from the STM” (p. 3130)
91.2.10 “CFM single ended loss test (7705 SAR only)” (p. 3091)			91.26 “To create and run a CFM single-ended loss OAM diagnostic test from the STM” (p. 3131)
91.2.11 “CFM two way SLM” (p. 3091)			91.27 “To create and run a CFM two way SLM OAM diagnostic test from the STM” (p. 3132)
91.2.12 “CFM LM test” (p. 3091)			91.28 “To create and run a CFM LM OAM diagnostic test from the STM” (p. 3134)
Performance Monitoring OAM diagnostic tests			
92.2.2 “PM session test support and configuration” (p. 3138)	Ethernet	Network element	92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146)
92.2.3 “CFM DMM session test” (p. 3139)			92.7 “To configure a CFM DMM session OAM diagnostic test from the STM” (p. 3149)
92.2.4 “CFM SLM session test” (p. 3139)			92.8 “To configure a CFM SLM session OAM diagnostic test from the STM” (p. 3150)
92.2.5 “CFM LMM session test” (p. 3139)			92.9 “To configure a CFM LMM session OAM diagnostic test from the STM” (p. 3151)

Table 90-1 NFM-P supported OAM diagnostic tests and configurations (continued)

OAM test category / test type	Test network object or service level	Network level	See Procedure
92.2.6 “TWAMP Light session test” (p. 3140)	IP Layer 3 interface	Network element	92.11 “To configure a TWAMP Light reflector” (p. 3153)
92.2.6 “TWAMP Light session test” (p. 3140)			92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154)
92.2.7 “TCC test” (p. 3140)			92.13 “To configure a TCC OAM diagnostic test from the STM” (p. 3155)
92.2.8 “MPLS DM session test” (p. 3140)	MPLS	LSP	92.10 “To configure an MPLS DM session OAM diagnostic test from the STM” (p. 3152)
Service OAM diagnostic tests			
“Service site ping” (p. 2986)	Service transport binding (level 5)	Network element	90.7 “To create and run a service site ping OAM diagnostic test from the STM” (p. 3007)
“VCCV ping” (p. 2986)	Service (level 6)	VPRN site	90.8 “To create and run a VCCV ping OAM diagnostic test from the STM” (p. 3008)
“VCCV trace” (p. 2988)			90.9 “To create and run VCCV trace OAM diagnostic test from the STM” (p. 3009) 90.10 “To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM” (p. 3010) (static PW to a dynamic PW segment) 90.11 “To create and run a VCCV trace OAM diagnostic from a dynamic PW to a static PW segment from the STM” (p. 3012) (dynamic PW to a static PW segment)
“Y.1564 bidirectional tests” (p. 2988)		Network element	89.7 “STM Y.1564 test configuration” (p. 2924)
L2 service OAM diagnostic tests			

Table 90-1 NFM-P supported OAM diagnostic tests and configurations (continued)

OAM test category / test type	Test network object or service level	Network level	See Procedure
"MAC populate" (p. 2989)	Service (level 6)	VPLS site Epipe VLL site VLL	90.12 "To create and run a MAC populate OAM diagnostic test from the STM" (p. 3013)
"MAC purge" (p. 2989)			90.13 "To create and run a MAC purge OAM diagnostic test from the STM" (p. 3014)
"MAC ping" (p. 2989)			90.14 "To create and run a MAC ping OAM diagnostic test from the STM" (p. 3015)
"MAC trace" (p. 2990)			90.15 "To create and run a MAC trace OAM diagnostic test from the STM" (p. 3016)
"CPE ping" (p. 2991)		VPLS site Epipe VLL site	90.16 "To create and run a CPE ping OAM diagnostic test from the STM" (p. 3017)
"ANCP loopback" (p. 2991)		Network element	90.17 "To create and run an ANCP loopback OAM diagnostic test from the STM" (p. 3018)
"VXLAN ping" (p. 2992)		VPLS site	90.18 "To create and run a VXLAN ping OAM diagnostic test from the STM" (p. 3019)
L3 service OAM diagnostic tests			
"VPRN ping" (p. 2992)	Service (level 6)	VPRN site	90.19 "To create and run a VPRN ping or VPRN trace OAM diagnostic test from the STM" (p. 3022) (STM method) 90.20 "To create and run a VPRN Ping, VPRN Trace, ICMP Ping, or ICMP Trace OAM diagnostic test from a service manager form" (p. 3023) (from a service)
"VPRN trace" (p. 2992)			
Service transport OAM diagnostic tests			
"Tunnel ping" (p. 2993)	Service transport (level 4)	Service tunnel (SDP)	90.21 "To create a tunnel ping OAM diagnostic test from the STM" (p. 3026) (STM method) 90.22 "To create and run a tunnel ping OAM diagnostic test from a service tunnel" (p. 3026) (service tunnel method)
"MTU ping" (p. 2994)			90.23 "To create and run an MTU ping OAM diagnostic test from the STM" (p. 3028) (STM method) 90.24 "To create and run an MTU ping OAM diagnostic test from a service tunnel" (p. 3028) (service tunnel method)
MPLS OAM diagnostic tests			

Table 90-1 NFM-P supported OAM diagnostic tests and configurations (continued)

OAM test category / test type	Test network object or service level	Network level	See Procedure
"LSP ping" (p. 2995)	Transport (level 3)	LSP LSP path MPLS-TP LSP	90.25 "To create and run a MPLS LSP ping OAM diagnostic test from the STM" (p. 3030)
"LSP trace" (p. 2996)			90.26 "To create and run a MPLS LSP trace OAM diagnostic test from the STM" (p. 3031)
"LDP tree trace" (p. 2996)		MPLS site LDP site	90.27 "To create and run a MPLS LDP tree trace OAM diagnostic test from the STM" (p. 3032)
"P2MP LSP ping" (p. 2997)		P2MP LSP, P2MP mLDP P2MP path	90.28 "To create and run a MPLS P2MP LSP ping OAM diagnostic test from the STM" (p. 3033)
"P2MP LSP trace" (p. 2997)			90.29 "To create and run a MPLS P2MP LSP trace OAM diagnostic test from the STM" (p. 3037)
L1/L2 OAM diagnostic tests			
"ATM ping" (p. 2998)	Layer 1 or Layer 2 (level 1)	ATM PVC connection	90.30 "To create and run an ATM ping OAM diagnostic test from the STM" (p. 3039) (STM method) 90.31 "To configure an ATM OAM loopback from a device Properties form" (p. 3040) (Properties form method)
Multicast OAM diagnostic tests			
"BIER ping" (p. 2998)	Routed network (level 2)	IP multicast traffic	90.32 "To create and run a BIER ping OAM diagnostic test from the STM" (p. 3040)
"BIER trace" (p. 2998)	Routed network (level 2)	IP multicast traffic	90.33 "To create and run a BIER trace OAM diagnostic test from the STM" (p. 3042)
"MFIB ping" (p. 2999)	Service (level 6)	VPLS site Epipe VLL site	90.34 "To create and run an MFIB ping OAM diagnostic test from the STM" (p. 3043)
"Mrinfo" (p. 2999)	Routed network (level 2)	IP multicast traffic	90.35 "To create and run an Mrinfo OAM diagnostic test from the STM" (p. 3044)
"Mtrace" (p. 2999)			90.36 "To create and run an Mtrace OAM diagnostic test from the STM" (p. 3044)
"Mtrace2" (p. 2999)			90.37 "To create and run an Mtrace2 OAM diagnostic test from the STM" (p. 3045)
ICMP OAM diagnostic tests			

Table 90-1 NFM-P supported OAM diagnostic tests and configurations (continued)

OAM test category / test type	Test network object or service level	Network level	See Procedure
"ICMP ping" (p. 3000)	Service (level 6)	VPRN site	90.38 "To create and run an ICMP ping OAM diagnostic test from the STM" (p. 3046)
"ICMP trace" (p. 3000)			90.39 "To create and run an ICMP trace OAM diagnostic test from the STM" (p. 3047)
"DNS ping" (p. 3000)			90.40 "To create and run an ICMP DNS ping OAM diagnostic test from the STM" (p. 3048)
VMN OAM diagnostic tests			
"PRBS test" (p. 3000)	1830 VWM TLU/ITP	Port	90.41 "To create and run a PRBS test" (p. 3048)
AOS CPE OAM diagnostic tests			
"CPE SLA test" (p. 3001)	Ethernet	Port or NE	89.8 "Sample OmniSwitch device SLA testing" (p. 2940)
"CPE SLA test group" (p. 3001)			89.8 "Sample OmniSwitch device SLA testing" (p. 2940) 90.42 "To create and run a OmniSwitch CPE SLA diagnostic test from the STM" (p. 3051)
Non-STM OAM diagnostic tests			
"OmniSwitch ping and traceroute OAM test" (p. 3001)	Ethernet	Port or NE	90.44 "To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script" (p. 3058) 90.45 "To configure and run an OmniSwitch OAM diagnostic ping test CLI script" (p. 3062) 90.46 "To configure and run an OmniSwitch OAM traceroute test CLI script" (p. 3064)
"OmniSwitch advanced loopback test" (p. 3002)	Ethernet	Port	90.47 "To configure an advanced loopback test on an OmniSwitch from a device Properties form" (p. 3067)
"F5 OAM loopback test (7705 SAR-M/ME)" (p. 3002)	Ethernet	Port	90.48 "To run the F5 OAM loopback diagnostic test from a 7705 SAR-M/ME Properties form" (p. 3068)
"802.3ah EFM OAM diagnostic test" (p. 3002)	Ethernet	Port	90.49 "To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form" (p. 3069) (Properties form method) 90.50 "To configure an 802.3ah EFM OAM diagnostic test on an OmniSwitch Properties form" (p. 3074) (OmniSwitch specific)

Table 90-1 NFM-P supported OAM diagnostic tests and configurations (continued)

OAM test category / test type	Test network object or service level	Network level	See Procedure
"One-time service validation test" (p. 3003)	Service (level 6)	VLL, VPLS, and VPRN services Composite services Dynamic LSPs Service tunnels	90.55 "To run a one-time validation test on a service" (p. 3084)

90.2.3 Multiple OAM diagnostic test configuration

You can also create, schedule, and run test suites that contain groups of OAM diagnostic tests using the NFM-P STM. The use of test suites is especially valuable when multiple objects of the same type require testing. Test suites can be scheduled to run on a regular basis to provide continual network performance feedback.

See Chapter 89, "Service Test Manager" for information about configuring multiple OAM diagnostic tests as part of a test suite using the STM.

See Chapter 5, "NFM-P-based schedules" for information about scheduled tasks.

90.3 OAM diagnostic test descriptions

90.3.1 Introduction

This section provides a description of supported NFM-P OAM diagnostic test functionality grouped by their functional area on the STM. Unless where noted, all test are accessible from the STM. See Table 90-1, "NFM-P supported OAM diagnostic tests and configurations" (p. 2981) for a list of all supported OAM diagnostic tests and their applicable procedures.

90.3.2 Ethernet CFM OAM diagnostic tests

Use the Ethernet CFM (Connectivity Fault Management) OAM diagnostic tests to detect, isolate, and report connectivity faults for L2 objects in Ethernet networks. The NFM-P Ethernet CFM function is implemented based on the IEEE 802.1ag OAM standard.

See Chapter 91, "Ethernet CFM" for a description of each Ethernet CFM OAM diagnostic test and applicable configuration procedure.

90.3.3 PM OAM diagnostic tests

Use the Performance Monitoring (PM) OAM diagnostic tests to assess the configuration and performance of Carrier Ethernet services such as Ethernet frame delay, frame delay variation, frame loss, and frame throughput measurements as specified by the ITU-T Y-1731 standard.

See Chapter 92, "Performance Monitoring tests" for a description of each PM OAM diagnostic test and applicable configuration procedure.

90.3.4 Service OAM diagnostic tests

Use the Service OAM diagnostic tests to assess the configuration and performance of an Ethernet service before customer notification and delivery, and to troubleshoot and resolve problems that are related to an individual service within the provider network.

Service site ping

The service site ping OAM diagnostic test, which is called svc-ping in CLI, provides end-to-end connectivity testing for an individual service within the provider network.

This diagnostic test operates at a higher level than the tunnel ping OAM because it verifies connectivity for an individual service rather than connectivity across the service tunnel. This allows you to isolate a problem within the service rather than at the port, which is the endpoint of the service tunnel.

The test verifies a service ID for correct and consistent provisioning between two service endpoints.

The following information can be verified from a service site ping OAM:

- local and remote service sites exists
- current state of the local and remote service sites
- local and remote service types are correlated
- same customer is associated with the local and remote service sites
- service-to-circuit association at both the local and remote service sites using the Use Local Tunnel and Use Remote Tunnel options, to check the circuit between service sites
- local and remote ingress and egress service labels match

See 90.7 “To create and run a service site ping OAM diagnostic test from the STM” (p. 3007) for information about how to create and run a service site ping OAM diagnostic test from the STM.

Note: The OAM messages, which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address that is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the service site ping OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously, the NE’s system interface address was used for this purpose.

VCCV ping

The VCCV ping OAM diagnostic test, which is called vccv-ping in the CLI, performs in-band connectivity tests. It can be used for VPLS, MVPLS, HVPLS, and all types of VLLs.

When applied to a VLL service, the test supports cross-circuit tests as long as the circuit types match; for example, an Epipe-to-Epipe connector, or an Epipe-to-VPLS/MVPLS site connector. The purpose of the ping is to determine that the destination PE device is the egress for the L2 FEC. The following figure shows a sample VCCV ping OAM diagnostic test.

Figure 90-1 VCCV ping OAM diagnostic

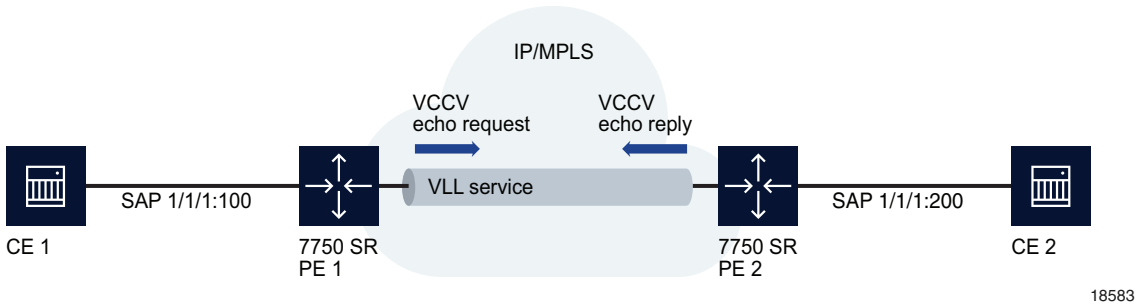

In Figure 90-1, “VCCV ping OAM diagnostic” (p. 2988) , the ping test packet or packets are sent with the destination IP address of PE 2. The request is encapsulated in a VLL packet and is forwarded to PE 2. PE 2 replies to the source device IP address. The packets are sent using the same encapsulation and along the same path as user packets in the VLL. This test provides a check of both the control plane and the data plane.

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the VCCV ping OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously the NE’s system interface address was used for this purpose.

When this test is applied to a VPLS, MVPLS or HVPLS, downstream SDP bindings are not configurable as they are in VLL services. Downstream bindings are typically related to switching sites and these are not relevant to VPLS. In VPLS, only the first SDP binding is configured, and this can be either a mesh or spoke binding.

Also, since an HVPLS comprises a mixture of active and standby spokes, the VCCV ping test execution will only be successful on active spokes that have associated active return spokes. A test suite will generate all the required VCCV ping tests, but only active spokes with active return spokes are counted in the statistics results. This also includes any service connector used in a composite service, provided that the connector is a spoke connector between a VPLS and an Epipe service.

See 90.8 “To create and run a VCCV ping OAM diagnostic test from the STM” (p. 3008) for information about how to create and run a VCCV ping OAM diagnostic test from the STM.

VCCV trace

The VCCV trace OAM diagnostic test, which is called vccv-trace in the CLI, displays the hop-by-hop path used by a VLL. VCCV trace can trace the entire path of a PW with a single command issued at the T-PE or at an S-PE. It is an iterative process by which the ingress T-PE or T-PE sends successive VCCV-Ping messages with incrementing the TTL value, starting from TTL=1. The process is terminated when the reply is from the egress T-PE or when a timeout occurs.

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the VCCV Trace OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously the NE’s system interface address was used for this purpose.

See 90.9 “To create and run VCCV trace OAM diagnostic test from the STM” (p. 3009) for information about how to create and run a VCCV trace OAM diagnostic test from the STM. To create a VCCV trace OAM diagnostic test from a static PW to a dynamic PW segment, perform 90.10 “To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM” (p. 3010) . To create a VCCV trace OAM diagnostic test from a dynamic PW to a static PW segment, perform 90.11 “To create and run a VCCV trace OAM diagnostic from a dynamic PW to a static PW segment from the STM” (p. 3012) .

Y.1564 bidirectional tests

Y.1564 bidirectional tests allow you to assess the configuration and performance of an Ethernet service before customer notification and delivery. Y.1564 provides a standard for measuring throughput, latency, frame loss, and jitter to assess if the service complies with an SLA. Multiple test types are supported. See 89.7 “STM Y.1564 test configuration” (p. 2924) in Chapter 89, “Service Test Manager” for more information.

90.3.5 L2 service OAM diagnostic tests

Use the L2 service OAM diagnostic tests to troubleshoot and resolve problems that are related to FIBs, MAC addressing, or hostname in a VLAN. These tests are generally used in combination; for example, MAC populate to inject a MAC address into the network, MAC ping to determine where the address was learned, MAC trace to determine the path, CPE ping to test the VPLS, and MAC purge to remove the injected MAC address.

MAC populate

The MAC populate OAM diagnostic test, which is called mac-populate in CLI, is used to:

- determine if the FIB table is accurate by testing forwarding plan correctness. This is done by populating a service FIB with an OAM-tagged MAC entry. This MAC entry indicates that the NE is the egress NE for the MAC address of a service. You can then use the FIB manager to see the OAM-tagged MAC entry.

- send a message through the flooding domain to learn a MAC address, as if a customer packet with that source MAC address had flooded the domain from that ingress point of the service.

You can:

- force an existing MAC address to become OAM-tagged
- distinguish, in the FIB manager, MAC addresses that are OAM-tagged
- age an OAM-tagged MAC address

In a MAC populate, the OAM-tagged MAC address is populated on the egress point of the service. You can specify whether to flood this OAM-tagged MAC address to other devices so that the same OAM-tagged entry is added to the FIB tables of other devices.

See 90.12 “To create and run a MAC populate OAM diagnostic test from the STM” (p. 3013) for information about how to create and run a MAC populate OAM diagnostic test from the STM.

MAC purge

The MAC purge OAM diagnostic test, which is called mac-purge in CLI, is used to delete an OAM-tagged entry from a FIB, which was generated using the MAC populate OAM diagnostic test. This clears the FIB of any learned information for a specific MAC address, allows the FIB to be populated only by a MAC populate request, and can be used to flush all devices in a service domain.

See 90.13 “To create and run a MAC purge OAM diagnostic test from the STM” (p. 3014) for information about how to create and run a MAC purge OAM diagnostic test from the STM.

MAC ping

The MAC ping OAM diagnostic test, which is called mac-ping in CLI, is used to test connectivity in a VLL or VPLS by verifying a remote MAC address at the far end of a service. The MAC ping determines the existence of the far-end egress point of the service. MAC pings can be sent in-band or out-of-band.

You must specify either:

- the target (far-end) MAC address
  OR
- the broadcast address

In a MAC ping that is out-of-band, the ping is forwarded along the flooding domain when no MAC address bindings exist or is sent along the bindings if MAC address bindings exist. A response ping is sent from the far-end device when there is an egress binding for the service.

In a MAC ping that is in-band, the ping is sent with a VC label TTL of 255. The ping packet goes across each hop, and when it reaches the egress router, it is identified by the OAM label and the response is sent back along the management plane.

Figure 90-2, “Sample MAC ping OAM diagnostic test” (p. 2990) shows a sample MAC ping OAM diagnostic test from one end of a service to the far-end MAC address of the service.

See 90.14 “To create and run a MAC ping OAM diagnostic test from the STM” (p. 3015) for information about how to create and run a MAC ping OAM diagnostic test from the STM.

Figure 90-2 Sample MAC ping OAM diagnostic test

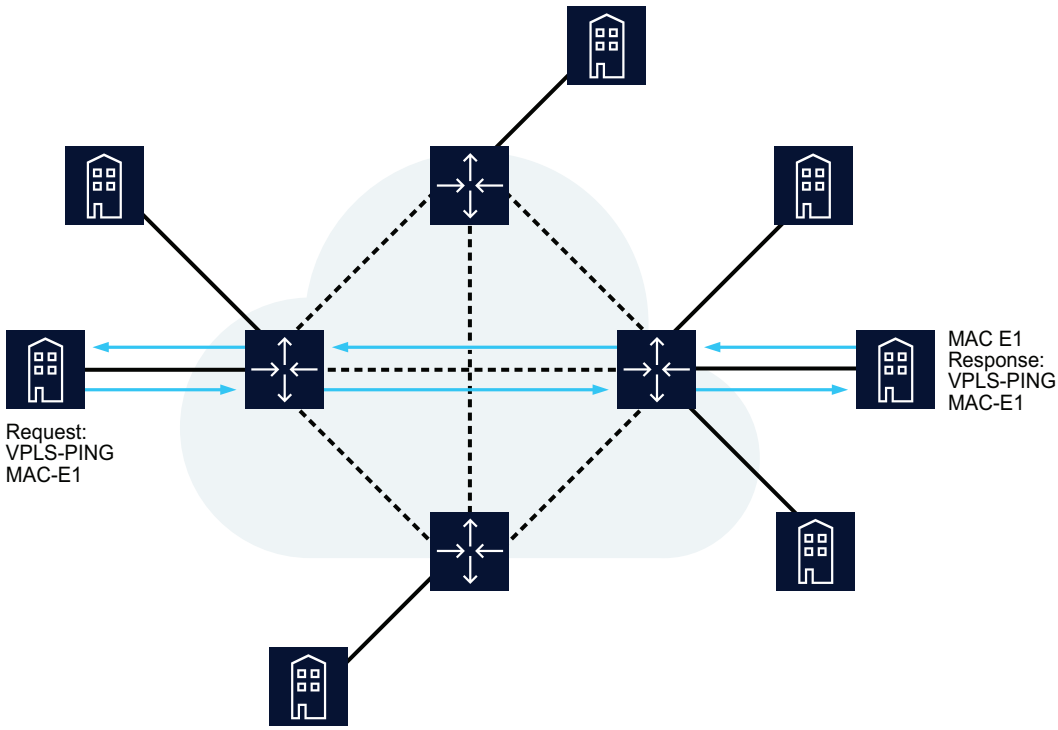

MAC trace

The MAC trace OAM diagnostic test, which is called mac-trace in CLI, displays the hop-by-hop route of MAC addresses used to reach the target MAC address at the far end of a service. MAC traces can be sent in-band or out-of-band.

You must specify either:

- the target (far-end) MAC address
  OR
- the broadcast address

In a MAC trace that is out-of-band, the destination IP address is specified by mapping the destination MAC address. If the destination MAC address is known to be a specific site, the far-end IP address of the service tunnel is used. If the destination MAC address is not known, the packet is sent to all service tunnels in the service.

In a MAC trace that is in-band, the trace request contains tunnel encapsulation, VC label, OAM, and other information. If the destination MAC address is known, the appropriate tunnel encapsulation and VC label is used. If the destination MAC address is not known, the packet is sent to all service tunnels, including all necessary tunnel encapsulation and egress VC labels for each bound service tunnel.

The following figure shows a sample MAC trace OAM diagnostic test. See Figure 90-3, “MAC trace OAM diagnostic test” (p. 2991) for information about how to create and run a MAC trace OAM diagnostic test from the STM.

Figure 90-3 MAC trace OAM diagnostic test

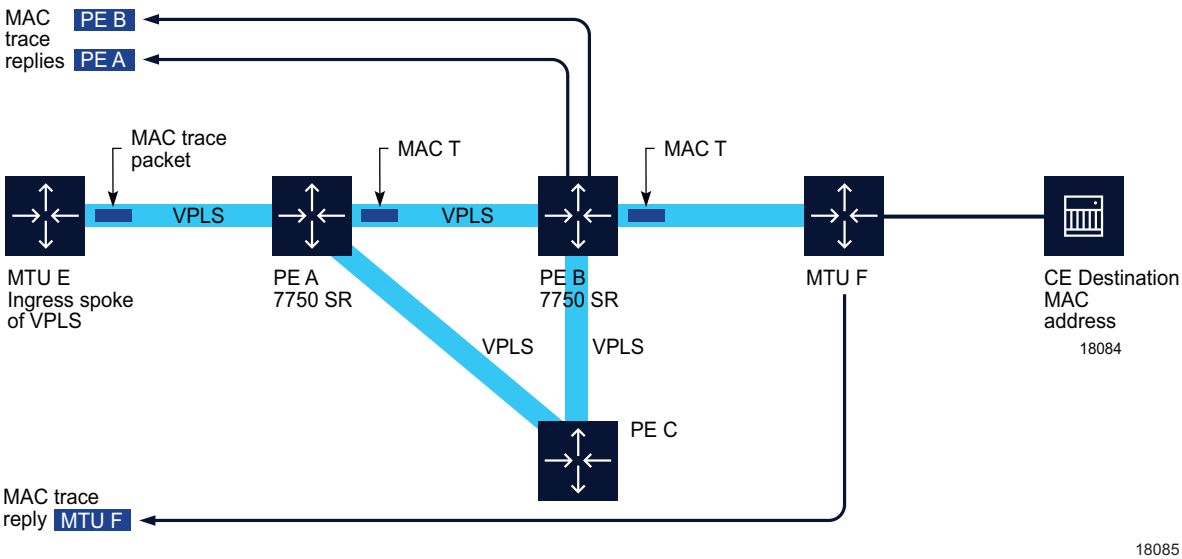

CPE ping

The CPE ping OAM diagnostic test, which is called cpe-ping in the CLI, is used to trace the end-to-end switching of specified MAC addresses of customer premises equipment. This ping extends the functionality of the MAC ping beyond the egress (customer-facing) port by allowing a ping to the SAP of a VPLS or a VLL Epipe service over a VPLS PBB backbone.

See 90.16 “To create and run a CPE ping OAM diagnostic test from the STM” (p. 3017) for information about how to create and run a CPE ping OAM diagnostic test from the STM.

ANCP loopback

The ANCP loopback OAM diagnostic test, which is called oam ancp in the CLI, is used to send DSL OAM commands to complete an OAM test to the access NE from a centralized point or when operational boundaries prevent direct access to the DSLAM. The ANCP loopback test raises an alarm that generates a log event displaying both successful and failed results.

See 90.17 “To create and run an ANCP loopback OAM diagnostic test from the STM” (p. 3018) for information about how to create and run an ANCP loopback OAM diagnostic test from the STM.

VXLAN ping

VXLAN is the overall Layer 2 tunneling mechanism used to connect services in data center architectures. The VXLAN connection is a point-to-point Layer 2 connection that has only two Virtual Tunnel Endpoints (VTEPs). From a transport perspective, the VXLAN tunnels are similar in

nature to a GRE SDP binding. The VXLAN connections are part of the subscriber's VPLS services connecting the Virtual Machine to the data center. Each subscriber has their own VPLS instance.

The VXLAN ping test is used to confirm that the Virtual Tunnel is operational by using Echo Request/Reply functions, as well as query the status of the Virtual Network Instance (VNI).

See 90.18 “To create and run a VXLAN ping OAM diagnostic test from the STM” (p. 3019) for information about how to create and run a VXLAN ping OAM diagnostic test from the STM.

90.3.6 L3 service OAM diagnostic tests

Use the L3 service OAM diagnostic tests to troubleshoot and resolve problems related to VPRN services provisioned on the NFM-P.

VPRN ping

The VPRN ping OAM diagnostic test determines the existence of the far-end egress point of the service. This allows testing of whether a specific destination can be reached. VPRN pings can be sent in-band or out-of-band. When a VPRN ping test packet is sent, a reply is generated if the targeted prefix is reachable over a VPRN SAP or VPRN spoke interface; otherwise the test packet is dropped in CPM. This also applies in the case of a routed VPLS interface.

You can perform a VPRN ping OAM diagnostic test from the STM, as described in 90.19 “To create and run a VPRN ping or VPRN trace OAM diagnostic test from the STM” (p. 3022) . You can also perform this test from a VPLS/MVPLS or VLL service form, as described in 90.20 “To create and run a VPRN Ping, VPRN Trace, ICMP Ping, or ICMP Trace OAM diagnostic test from a service manager form” (p. 3023) , or from a service flat topology map.

VPRN trace

The VPRN trace OAM diagnostic test displays the hop-by-hop path for a destination IP address within a VPRN service. This allows operators to know the destination path of customer traffic. VPRN traces can be sent in-band or out-of-band. When a VPRN trace test packet is sent, a reply is generated if the targeted prefix is reachable over a VPRN SAP or VPRN spoke interface. Otherwise the test packet is dropped in CPM. This also applies in the case of a routed VPLS interface.

You can perform a VPRN trace OAM diagnostic test from the STM, as described in 90.19 “To create and run a VPRN ping or VPRN trace OAM diagnostic test from the STM” (p. 3022) . You can also perform this test from a VPLS/MVPLS or VLL service form, as described in 90.20 “To create and run a VPRN Ping, VPRN Trace, ICMP Ping, or ICMP Trace OAM diagnostic test from a service manager form” (p. 3023) , or from a service flat topology map.

90.3.7 Service transport OAM diagnostic tests

Use the service transport OAM diagnostic tests to troubleshoot and resolve problems related to VPRN services provisioned on the NFM-P.

Tunnel ping

The tunnel ping OAM diagnostic test, which is called sdp-ping in the CLI, performs in-band unidirectional or bidirectional connectivity tests on service tunnels (also called an SDP). Use tunnel ping to troubleshoot and resolve tunnel and service problems that are related to issues that circuits may have transmitting traffic across the GRE or MPLS network.

The OAM packets are sent in-band in the tunnel encapsulation, so they follow the same path as the service traffic. The response can be received out-of-band in the control plane or in-band using the data plane for a bidirectional test.

For a unidirectional test, tunnel ping OAM tests:

- egress service tunnel ID encapsulation
- whether the packet can reach the far-end IP address destination of the service tunnel ID within its encapsulation
- whether a packet of the specified size goes to the far-end IP address of the service tunnel ID within its encapsulation
- forwarding class mapping to ensure that the test packet is treated the same as the customer traffic
- determine whether SLA delay metrics are met

For a bidirectional test, tunnel OAM uses a local egress service tunnel ID and an expected remote service tunnel ID, so the user can specify where the returned messages should be sent from based on the far-end tunnel ID.

Figure 90-4, “Sample tunnel ping OAM diagnostic test” (p. 2993) shows how a tunnel OAM packet can be inserted to test the connectivity between two customer LANs across the IP/MPLS core.

See 90.21 “To create a tunnel ping OAM diagnostic test from the STM” (p. 3026) for information about how to create and run an tunnel ping OAM diagnostic test from the STM.

Figure 90-4 Sample tunnel ping OAM diagnostic test

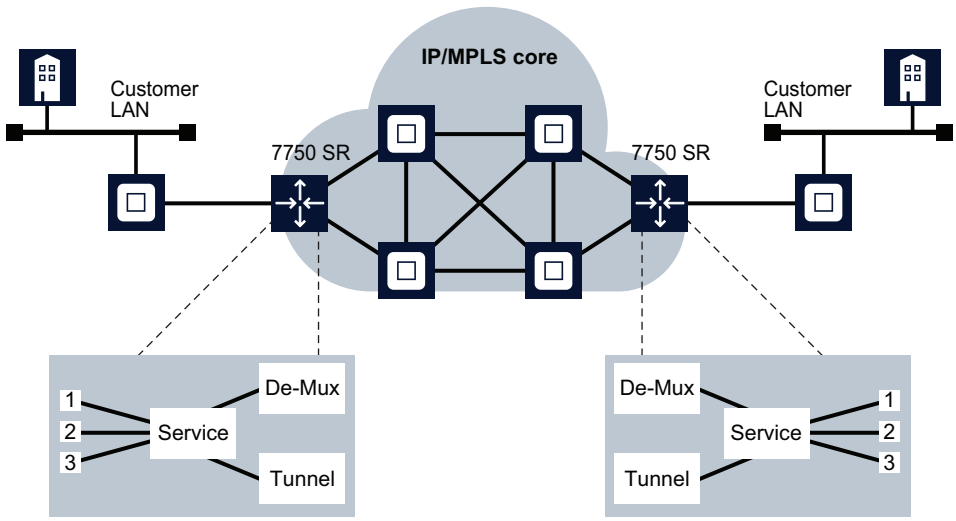

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP

session. The messages of the tunnel ping OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously the NE's system interface address was used for this purpose.

MTU ping

The MTU ping OAM diagnostic test, which is called sdp-mtu in CLI, provides a tool for service providers to:

- determine the maximum frame (MTU) size supported between the service ingress and the service termination point on a service tunnel (also called an SDP), to within one byte
- solve troubleshooting issues that are related to equipment used across the network core that may not be able to handle large frame sizes

In a large network, network devices can support a variety of packet sizes, up to a limit, that are transmitted across its interfaces. This size limit is referred to as the MTU of network interfaces. You must consider the MTU of the entire service tunnel end-to-end when you provision services, especially for VLL services in which the service must support the ability to transmit the largest customer packet.

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the MTU ping OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously the NE's system interface address was used for this purpose.

See 90.23 "To create and run an MTU ping OAM diagnostic test from the STM" (p. 3028) for information about how to create and run an MTU ping OAM diagnostic test from the STM. See 90.24 "To create and run an MTU ping OAM diagnostic test from a service tunnel" (p. 3028) for creating and running an MTU ping OAM diagnostic test from a service tunnel.

90.3.8 MPLS OAM diagnostic tests

Use the MPLS OAM diagnostic tests to test OAM functionality for MPLS provisioned networks based on RFC 4379.

LSP ping

The LSP ping OAM diagnostic test, which is called lsp-ping in CLI, performs in-band MPLS LSPs, MPLS-TP LSPs, LDPs, and MPLS path connectivity tests.

You can use an LSP ping to:

- detect data plane failures in LSPs or MPLS-TP LSP and with LSP or MPLS-TP LSP connectivity
- test whether the LSP or MPLS-TP LSP tunnels are working in both directions

In an LSP or MPLS-TP LSP ping, the originating router creates an MPLS echo request packet for the LSP and MPLS path to be tested. The MPLS echo request packet is sent and awaits an MPLS echo reply packet from the router that terminates the LSP. The status of the LSP is displayed when the MPLS echo reply packet is received. Figure 90-5, "LSP ping OAM diagnostic test" (p. 2995) shows an example of an LSP ping OAM diagnostic test.

See 90.25 “To create and run a MPLS LSP ping OAM diagnostic test from the STM” (p. 3030) for information about how to create and run an LSP ping OAM diagnostic test from the STM.

Figure 90-5 LSP ping OAM diagnostic test

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the LSP Ping use the T-LDP local LSR ID as the source IP address. Previously the NEs system interface address was used for this purpose.

LSP trace

The LSP trace OAM diagnostic test, which is called lsp-trace in CLI, displays the hop-by-hop route used by MPLS LSPs, MPLS-TP LSPs, LDPs, and the MPLS path. You can use an LSP trace to isolate a data plane failure to a particular router and to provide LSP path tracing.

The following information can be determined from the test:

- hop-by-hop path
- destination path of the packets

When performing an MPLS LSP or MPLS-TP LSP trace, the originating router creates an MPLS echo request packet for the LSP to be tested. The packet contains increasing TTL values. The MPLS echo request packet is sent and awaits a TTL exceeded response or the MPLS echo reply packet from the router that terminates the LSP. The devices along the hop-by-hop route reply to the MPLS echo request packets with TTL and MPLS echo reply information.

Figure 90-6, “LSP Trace diagnostic test” (p. 2996) shows an example of an LSP trace OAM diagnostic test.

See 90.26 “To create and run a MPLS LSP trace OAM diagnostic test from the STM” (p. 3031) for information about how to create and run an LSP trace OAM diagnostic test from the STM.

Figure 90-6 LSP Trace diagnostic test

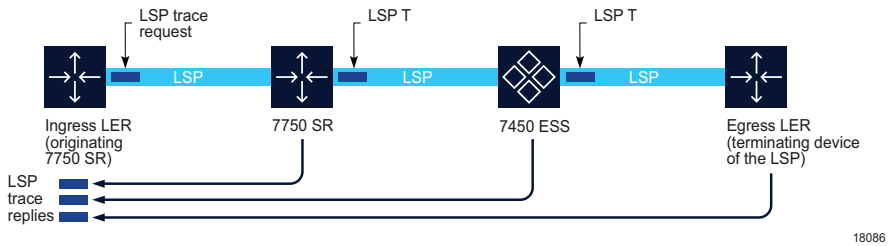

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the LSP trace OAM diagnostic test use the T-LDP local LSR ID as the source IP address. Previously the NE's system interface address was used for this purpose.

LDP tree trace

The LDP tree trace OAM diagnostic test, which is called ldp-treetrace in CLI, is used to detect and discover the ECMP routing paths for an LSP between egress and ingress routers.

The following information is determined from the test:

- number of ECMP paths
- number of failed hops

In an LDP tree trace, the originating router creates an MPLS echo request packet. The packet contains a set of IP header destination addresses. Routers along the path reply to the request with information about themselves and neighboring routers in the downstream path. The originating router uses this information to probe the downstream routers until it discovers a bit map setting common to all routers along the path. The result is a tree of available routers that the originating router can use for next hops. After discovery, the paths are tested using LSP ping and LSP trace.

Note: The OAM messages which operate over an LDP LSP, and/or over a PW signaled by T-LDP, and which require a response using the IP path, must use a source IP address which is reachable within the same routing domain as that of the LSR ID of the LDP or T-LDP session. The messages of the LDP tree trace use the T-LDP local LSR ID as the source IP address. Previously the NE's system interface address was used for this purpose.

See 90.27 "To create and run a MPLS LDP tree trace OAM diagnostic test from the STM" (p. 3032) for information about how to create and run an LDP tree trace OAM diagnostic test from the STM.

P2MP LSP ping

The PSMP LSP ping OAM diagnostic test, which is called p2mp-lsp-ping in CLI, performs in-band LSP connectivity tests.

The echo request message is sent on the active P2MP instance and is replicated in the data path over all branches of the P2MP LSP instance. By default, all egress LER devices that are leaves of the P2MP LSP instance will reply to the echo request message.

You can reduce the scope of the echo reply messages by entering a list of addresses for the egress LER devices that are required to reply. A maximum of 5 addresses can be specified in a single run of the p2mp-lsp-ping command. If all 5 egress LER devices are 7x50 NEs, they will be able to parse the list of egress LER addresses and will reply. The p2mp-lsp-ping command specifies that only the top address in the P2MP Egress Identifier TLV must be inspected by an egress LER. When interoperating with other implementations, a 7x50 egress LER will respond if its address is anywhere in the list. Furthermore, if another vendor implementation is the egress LER, only the egress LER matching the top address in the TLV may respond.

If you enter the same egress LER address more than once in a single p2mp-lsp-ping command, the head-end NE displays a response to a single one and displays a single error warning message for the duplicate ones. When queried over SNMP, the head-end NE issues a single response trap and issues no trap for the duplicates.

See 90.28 “To create and run a MPLS P2MP LSP ping OAM diagnostic test from the STM” (p. 3033) for information about how to create and run a P2MP LSP ping OAM diagnostic test from the STM.

P2MP LSP trace

The PSMP LSP trace OAM diagnostic test, which is called p2mp-lsp-trace in CLI, performs in-band LSP connectivity tests.

The LSP trace capability allows you to trace the path of a single S2L path of a P2MP LSP. Its operation is similar to that of the p2mp-lsp-ping command, but the sender of the echo reply request message includes the Downstream Detailed Mapping TLV to request the downstream branch information from a branch LSR or BUD LSR. The branch LSR or BUD LSR will then also include the Downstream Detailed Mapping TLV to report the information about the downstream branches of the P2MP LSP. An egress LER must not include this TLV in the echo response message.

The parameter probe-count operates in the same way as in the LSP trace on a P2P LSP. It represents the maximum number of probes sent per TTL value before giving up on receiving the echo reply message. If a response is received from the traced NE before reaching maximum number of probes, then no more probes are sent for the same TTL. The sender of the echo request then increments the TTL and uses the information it received in the Downstream Detailed Mapping TLV to start sending probes to the NE downstream of the last NE that replied. This continues until the egress LER for the traced S2L path replies.

See 90.29 “To create and run a MPLS P2MP LSP trace OAM diagnostic test from the STM” (p. 3037) for information about how to create and run a P2MP LSP trace OAM diagnostic test from the STM.

90.3.9 L1/L2 OAM diagnostic tests

Use the L1/L2 OAM diagnostic tests to test ATM PVC connections.

ATM ping

The ATM ping OAM diagnostic test, which is called atmoam-ping in CLI, performs an ATM ping on an ATM PVC from the PVC endpoint using ATM OAM loopback cells. An ATM ping test verifies VC integrity and endpoint connectivity for PVCs using OAM loopback capabilities.

See 90.30 “To create and run an ATM ping OAM diagnostic test from the STM” (p. 3039) for information about how to create and run an ATM ping OAM diagnostic test from the STM. See 90.31 “To configure an ATM OAM loopback from a device Properties form” (p. 3040) for creating and running a ATM ping OAM diagnostic test from an NE Properties form.

90.3.10 Multicast OAM diagnostic tests

Use the multicast OAM diagnostic tests to troubleshoot and resolve problems with the multicast component of a VPLS or a VPRN service.

BIER ping

The BIER ping OAM diagnostic test, which is called bier-ping in CLI, confirms connectivity within a BIER-enabled multicast network. Test packets are generated and sent in a specified sub-domain. Bier ping tests can be configured to specify a BFR ID, a range of BFR IDs, or a BFR prefix.

See 90.32 “To create and run a BIER ping OAM diagnostic test from the STM” (p. 3040) for information about how to create and run a BIER ping OAM diagnostic test from the STM.

BIER trace

The BIER trace OAM diagnostic test, which is called bier-trace in CLI, identifies the hop-by-hop route within a BIER-enabled multicast network. Test packets are generated and sent in a specified sub-domain from any BFR to any other BFR. BIER trace tests can be configured to specify a BFR ID or a BFR prefix.

See 90.33 “To create and run a BIER trace OAM diagnostic test from the STM” (p. 3042) for information about how to create and run a BIER trace OAM diagnostic test from the STM.

MFIB ping

The multicast FIB (MFIB) ping OAM diagnostic test, which is called mfib-ping in CLI, identifies the SAPs that egress an IP multicast stream within a multicast component of a VPLS. This diagnostic test can also be used to display the SAPs that are operationally up in the VPLS.

See 90.34 “To create and run an MFIB ping OAM diagnostic test from the STM” (p. 3043) for information about how to create and run an MFIB ping OAM diagnostic test from the STM.

Mrinfo

The Mrinfo (multicast router information) OAM diagnostic test, which is called mrinfo in CLI, identifies VPRN multicast information for the target router. The information includes details that are related to adjacent routers, supported protocols, traffic metrics, and time-to-live thresholds. Administrators can use this information to identify bidirectional adjacency relationships.

See 90.35 “To create and run an Mrinfo OAM diagnostic test from the STM” (p. 3044) for information about how to create and run a Mrinfo OAM diagnostic test from the STM.

Mtrace

The Mtrace (multicast trace) OAM diagnostic test, which is called mtrace in CLI, identifies the hop-by-hop route used by VPRN multicast traffic to reach the target router. This diagnostic gathers the hop address, routing error conditions, and packet statistics at each hop. The NFM-P attempts to trace the receiver-to-sender route for the traffic. The destination of the diagnostic can be any PIM-enabled interface in the routing instance.

In this test, the command sends a ping to the head of the specified multicast group (which can be one or more hops away from the launch NE). If that ping fails, the command will ping the neighboring NE, then the neighbor's neighbor, and so on up the multicast group chain to determine which NE fails to respond.

See 90.36 “To create and run an Mtrace OAM diagnostic test from the STM” (p. 3044) for information about how to create and run an Mtrace OAM diagnostic test from the STM.

Mtrace2

The Mtrace2 (multicast trace version 2) OAM diagnostic test, which is called mtrace2 in CLI, is similar to the Mtrace OAM test, but with added functionality for isolating packet loss problems and configuration problems. Tests are initiated from an Mtrace2 client toward a specified source, or a Rendezvous Point (RP) if no source address is specified. In the STM configuration, there is an option for StarG, which allows initiation from any source. Mtrace2 uses UDP-based messages, and supports base router and VPRN instances. Both IPv4 and IPv6 are supported.

See 90.37 “To create and run an Mtrace2 OAM diagnostic test from the STM” (p. 3045) for information about how to create and run an Mtrace2 OAM diagnostic test from the STM.

90.3.11 ICMP

Use the ICMP OAM diagnostic tests to troubleshoot and resolve IP reachability problems; for example, to send error messages indicating that a requested VPRN service is not available or that a host or router could not be reached.

ICMP ping

The ICMP ping OAM diagnostic test, which is called icmp-ping in CLI, identifies the reachability of a remote host across the IP network. The test is used with ICMP trace to detect and localize faults in IP networks.

You can also enable the ICMP ping OAM diagnostic test from the VRF site of a subscriber VPRN service. An ICMP ping determines the existence of the far-end egress point of the service. This allows testing of whether a specific destination can be reached. ICMP pings can be sent in-band or out-of-band.

ICMP ping tests can be used to control the operational state of L3 service SAPS. On supporting NEs, users can configure a policy-based ICMP ping template and apply it to all required IES and VPRN L3 interfaces. The interfaces can use the template to populate ICMP ping test configuration values. For more information about ICMP ping templates, see 90.51 “To configure an ICMP Ping template” (p. 3079).

See 90.38 “To create and run an ICMP ping OAM diagnostic test from the STM” (p. 3046) for information about how to create and run an ICMP ping OAM diagnostic test from the STM.

ICMP trace

The ICMP trace OAM diagnostic test, which is called icmp-trace in CLI, identifies the diagnostic used to trace the ICMP traceroute control table. The tool is used with ICMP ping to detect and localize faults in IP networks.

ICMP trace displays the hop-by-hop path for a destination IP address within a VPRN service. This allows operators to know the destination path of subscriber traffic. ICMP traces can be sent in-band or out-of-band.

See 90.39 “To create and run an ICMP trace OAM diagnostic test from the STM” (p. 3047) for information about how to create and run an ICMP trace OAM diagnostic test from the STM.

DNS ping

The DNS ping OAM diagnostic test, which is called dns-ping in CLI, identifies the diagnostic used to ping the DNS name, if a DNS name resolution is configured. See 90.40 “To create and run an ICMP DNS ping OAM diagnostic test from the STM” (p. 3048) for information about how to create and run a DNS ping OAM diagnostic test from the STM.

90.3.12 VWM diagnostic test

Use the VWM OAM diagnostic tests to detect signal failures between two 1830 VWM TLUs or two 1830 VWM ITPs.

PRBS test

You can verify the data path between two 1830 VWM TLUs or two 1830 VWM ITPs by performing the PRBS test. The test is performed by generating PRBS signals, transmitting the signals through a daisy chain of interfaces, and looping them back to detect signal failures.

The NFM-P allows you to run the PRBS test and view the results. The PRBS measurements are performed on client or line ports of the 1830 VWM TLU and 1830 VWM ITP devices and loopbacks are configured on client or line ports of the 1830 VWM TLU and 1830 VWM ITP devices. The administrative state of both ports is set to Maintenance before the PRBS test is performed.

You can stop the test manually at any time or automatically after the configured duration has elapsed. On the same port, a new measurement can only be started after the previous measurement is stopped. The result of a measurement on a specific port is lost when a new measurement is started on the same port.

The result of the test is obtained by reading the number of errors detected in the received PRBS signal in the 1830 VWM TLU or 1830 VWM ITP.

See 90.41 “To create and run a PRBS test” (p. 3048) for more information about performing the PRBS test.

90.3.13 AOS CPE OAM diagnostic tests

Use the AOS CPE OAM diagnostic tests to perform L2 ping and link traces on select OmniSwitch devices in Metro Ethernet networks. See 89.8 “Sample OmniSwitch device SLA testing” (p. 2940) in Chapter 89, “Service Test Manager” for more information.

CPE SLA test

The NFM-P STM allows you to perform a CPE SLA OAM diagnostic test to validate and test customer SLAs used on select OmniSwitch devices in Metro Ethernet networks. The test is critical for provisioning or troubleshooting network services between customer endpoints. The NFM-P STM allows you to test supports unidirectional traffic and IPv4.

See 89.8 “Sample OmniSwitch device SLA testing” (p. 2940) for information about how to configure a CPE SLA OAM diagnostic test using the STM.

CPE SLA test group

The CPE SLA test group allows you to perform a CPE SLA test-head test using one CPE test-head profile or a group of profiles. You can also run several tests for one CPE test-head profile or group of profiles.

See 89.8 “Sample OmniSwitch device SLA testing” (p. 2940) for information about how to create a CPE SLA test-head Group profile using the STM, which can be use to bind individual CPE test-head profiles to form a group of tests.

90.3.14 Non-STM OAM diagnostic and validation tests

The following OAM diagnostic and validation tests are configured from alternate launch points on the NFM-P GUI other than the STM.

OmniSwitch ping and traceroute OAM test

The OmniSwitch ping and traceroute OAM diagnostic test uses user-defined CLI scripts to allow you to troubleshoot and resolve problems with IP reachability on an OmniSwitch.

See 90.44 “To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script” (p. 3058) for information about creating an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script.

See 90.45 “To configure and run an OmniSwitch OAM diagnostic ping test CLI script” (p. 3062) and 90.46 “To configure and run an OmniSwitch OAM traceroute test CLI script” (p. 3064) for information about configuring and running an OmniSwitch ping and traceroute CLI scripts.

OmniSwitch advanced loopback test

The OmniSwitch advanced loopback test allows you to troubleshoot and resolve network faults on OmniSwitch ports to isolate the network segments where errors have occurred. See 90.47 “To configure an advanced loopback test on an OmniSwitch from a device Properties form” (p. 3067) for information about configuring an OmniSwitch advanced loopback test.

F5 OAM loopback test (7705 SAR-M/ME)

The F5 OAM loopback test allows you to validate ATM bonding on DSL ports on a 7705 SAR-M/ME device equipped with a DSL module. The parameters for the F5 OAM loopback test are only visible when a DSL module is connected to an ISAM device with the bonding type of ATM and the device is operationally up.

See 90.48 “To run the F5 OAM loopback diagnostic test from a 7705 SAR-M/ME Properties form” (p. 3068) for information about configuring an F5 OAM loopback test.

802.3ah EFM OAM diagnostic test

An 802.3ah EFM OAM diagnostic test addresses three key operational issues when deploying Ethernet across geographically disparate locations: link monitoring, fault signaling, and remote loopback.

Link monitoring provides some basic error definitions for Ethernet so entities can detect failed and degraded connections. Fault signaling provides mechanisms for one entity to signal another that it has detected an error. Remote loopbacks allows one NE to put another NE into a state whereby all inbound traffic is immediately reflected back onto the link.

The 802.3ah EFM OAM test supports:

- EFM OAM capability discovery
- active and passive modes
- remote failure indication
- local and remote loopbacks
- EFM OAMPDU tunneling
- high-resolution EFM OAM timers

You can configure a 802.3ah EFM OAM diagnostic test on Ethernet ports in network mode on the 7210 SAS, 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS. See 90.49 “To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form” (p. 3069) for more information.

For OmniSwitch NEs, the 802.3ah EFM OAM diagnostic test is supported at the NE level or at the Ethernet port level on the Ethernet ports in network mode. See 90.50 “To configure an 802.3ah EFM OAM diagnostic test on an OmniSwitch Properties form” (p. 3074) for more information.

One-time service validation test

A one-time service validation test provides a mechanism to quickly test a service or composite service either at the time it is provisioned or when it is already in service.

A variety of services and transport components can be tested using this function including:

- VLL, VPLS, and VPRN services
- composite services
- dynamic LSPs
- service tunnels

A one-time service validation test operates like a Validator test suite however the temporary objects that are automatically created for this method are deleted upon its completion. Items deleted include the test suite, tests, MEGs, and MEPs. No user action is required to do this.

Another significant difference with the One-time service validation test is that no results file is created. Instead, all tests and results are stored as archived objects that remain after the test is finished. The results can be viewed in the NFM-P GUI. See the XML API Reference (oneTimeValidate under sas.TestManager) for the required syntax to view the results using the XML API.

Executing a one-time service validation test also updates the OAM Validation Failed state field as required, indicating whether or not the validation was successful.

See 90.55 “To run a one-time validation test on a service” (p. 3084) for information about configuring a one-time service validation test.

90.4 Sample OAM diagnostic test configuration

90.4.1 Sample OAM diagnostic sequence

The following figure shows a sample OAM diagnostic sequence illustrating how you can use multiple OAM tests to verify service creation and diagnose service problems. This sample shows a VPLS diagnostic sequence.

Figure 90-7 Sample OAM diagnostic sequence for a VPLS

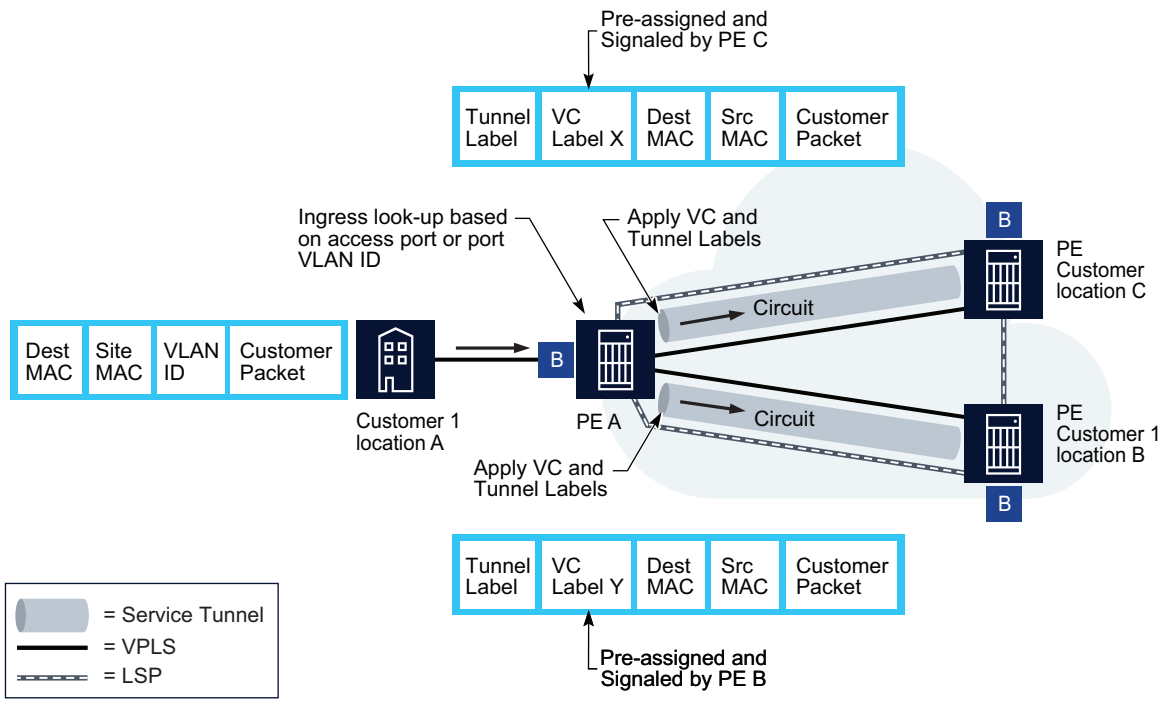

90.4.2 Sample OAM diagnostic configuration tasks for a service

The following table lists the high-level tasks necessary to configure the sample in Figure 90-7, “Sample OAM diagnostic sequence for a VPLS” (p. 3004).

Table 90-2 Sample OAM diagnostic configuration tasks for a service

Task	Description
Service creation and OAM validation	

Table 90-2 Sample OAM diagnostic configuration tasks for a service (continued)

Task	Description
1. MPLS and LSP creation	Create an MPLS network and LSPs for use by the service tunnels connecting the VPLS sites. Test the validity of the LSPs using LDP tree trace or LSP ping. If the results indicate a problem with the path, use LSP trace to check the specific MPLS path for the device causing the ping failure.
2. Service tunnel creation	Create a service tunnel that uses the MPLS LSP created earlier and perform a tunnel ping on the service tunnel to verify tunnel connectivity. Create all tunnels necessary to interconnect the VPLS sites, and repeat the tunnel ping to ensure tunnel connectivity. After all of the tunnels are created, use the tunnel ping remote tunnel option and specify the return tunnel path. Verify bidirectional tunnel connectivity.
3. Service creation	Create a service using the service tunnels to interconnect the VPLS sites, either using mesh or spoke service tunnel bindings. Use a service site ping between each VPLS site device and its neighboring sites to verify service configuration consistency.
4. MAC diagnostics	Connect the CPE devices to the VPLS and verify traffic. Use MAC trace from the edge devices to verify MAC address learning by the VPLS sites and to ensure that the correct associations are made between MAC addresses and the service tunnels or SAPs to which they are bound. Use MAC ping against an unknown MAC address to verify that no response is returned. Use MAC populate to create an OAM-specific MAC address. Use MAC ping and MAC trace against the created MAC address to verify that customer traffic is not affected by the additional MAC address. Use MAC purge to remove the created OAM MAC address.
Service OAM diagnostics	
5. Diagnose traffic flow problems at a specific MAC address	Use MAC ping against the MAC address to which traffic is not flowing. Use the source and destination MAC address to simulate customer traffic routes as closely as possible. Use MAC trace to pinpoint the location of the traffic failure. Check for MAC filter rules or MAC table sizes to identify possible causes of the failure; for example, incorrect configurations.
6. Diagnose the components of the service	Use service site ping to test the potential next hops to ensure consistent configuration. Use tunnel ping to the far end of the tunnel using the remote tunnel option and specify the return tunnel path. Verify bidirectional tunnel connectivity. Use LSP ping to determine if the tunnel is working. Use an LSP trace to determine if an intervening device is down.

90.5 OSS client OAM diagnostic test results file retrieval

90.5.1 Method to receive OAM logs

XML API clients can use the <registerSasLogToFile> method to receive OAM logs that are based on specific class types. Clients must register to receive the logs, and if no client registers, no logs are created. See the *NSP NFM-P XML API Developer Guide* for detailed information about XML API client methods.

The OAM data is exported to the file after the data is read from the NE. The file is saved to a specified directory on the NFM-P server.

Before you configure XML API clients to receive Test Results files, it is recommended that you configure the DB Test Result Retention Time (hours), Default Test Result Storage, and Log Retention Time (minutes) parameters as required. These are accessed by clicking Administration→System Preferences and then clicking the Test Manager tab.

Procedures to configure and perform OAM diagnostic tests

90.6 OAM diagnostic test workflow

90.6.1 Stages

Note: The tests mentioned in this workflow are generally used on an as-required (on-demand) basis. However, certain tests (for example, PM Session tests) can also be run on a scheduled (proactive) basis.

There are some user and system preferences that can change the operational behavior of OAM diagnostic tests such as specifying the default retention time for db test results or specifying OAM diagnostic test limits. See the STM chapter workflow in 89.9 “STM workflow” (p. 2945) for more information.

Not all service assurance tools are applicable to every NE managed by the NFM-P.

1

Create the transport network and customer services that you want to monitor or troubleshoot. See the appropriate service management chapter in this guide.

2

Review a typical OAM diagnostic test configuration that illustrates how you can use multiple OAM tests to verify service creation and diagnose service problems; see 90.4 “Sample OAM diagnostic test configuration” (p. 3004) .

3

As required, review both the generic and device-specific STM implementation samples in Chapter 89, “Service Test Manager” to help you understand how the various OAM diagnostics tests can be used to verify or monitor service operation, monitor and identify network problems.

4

If the OAM diagnostic test being configured is part of an STM test policy or STM test suite, review the STM chapter workflow to determine any prerequisite or post-requirement configuration requirements; see 89.9 “STM workflow” (p. 2945) . Also see Chapter 92, “Performance Monitoring tests” for information regarding the PM Session tests.

5

If a service needs to be tested or monitored for SLA compliance, or a customer service is compromised and corrective action is required, use the appropriate OAM diagnostic test to analyze, verify, or troubleshoot the problem. See Table 90-1, “NFM-P supported OAM diagnostic tests and configurations” (p. 2981) for a list of all NFM-P supported OAM diagnostic tests and their applicable procedures (covers 91.19 “To create and run a Global MEG OAM diagnostic test from the STM” (p. 3120) to 90.55 “To run a one-time validation test on a service” (p. 3084) in this chapter and 92.7 “To configure a CFM DMM session OAM diagnostic test from the STM” (p. 3149) to 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) in Chapter 92, “Performance Monitoring tests”).

6

If a running OAM diagnostic test needs to be halted, you can click Stop Execution available on the test form. The test's Result Status will indicate Skipped. The Stop Execution button is also available when running PM Session tests and on test suite forms.

90.7 To create and run a service site ping OAM diagnostic test from the STM

90.7.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service→Create Service Site Ping. The Service Site Ping (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

When you set the target IP address, then choose a service ID, only service IDs from the selected site are available. When you set the Use Local Tunnel and Use Remote Tunnel parameters, the test becomes a ping that tests the service tunnel bindings between the service sites.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 "To run one or more OAM diagnostic tests from the STM and view the test results" (p. 2959) . See 89.20 "To interpret OAM diagnostic test results on the STM" (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.8 To create and run a VCCV ping OAM diagnostic test from the STM

90.8.1 Steps

You can perform this test from a VPLS/MVPLS or VLL service form, or from the service flat topology map.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service→VCCV Ping. The VCCV Ping (Create) form opens.

3

Configure the required parameters.

4

Set the Test Type parameter in the Test Object panel.

To create a VCCV ping from a static PW to a dynamic PW segment, set the Test Type parameter to Static and configure the required parameters in the MPLS-TP VCCV Ping panel.

To create a VCCV ping from a dynamic PW to a static PW segment, set the Test Type parameter to Regular.

If you set the Test Type parameter to FEC 128, configure the required parameters in the FEC 128 Configurations panel.

5

Select the required spoke or mesh SDP binding in the First Spoke/Mesh SDP Binding panel.

Note: The selection of spoke bindings is applicable to both VPLS and VLL services; mesh bindings are only applicable to VPLS and MVPLS. However, when creating this test for MVPLS and VPLS, the concept of downstream spoke SDP bindings does not apply.

6

To configure this test for a VLL, select the required downstream spoke SDP binding in the Downstream SDP Binding panel.

7

Set the Target FEC Type parameter.

If creating a VCCV ping from a static PW to a dynamic PW segment, and the next segment is a dynamic PW, set the Target FEC Type parameter to FEC 129 and configure the required parameters in the Test Multi-Segment PW (FEC 129) panel for the dynamic PW. The PW ID parameter is the Service ID of the service that contains the dynamic PW between the source and destination NEs in this segment.

If creating a VCCV ping from a dynamic PW to a static PW segment, set the Target FEC Type parameter to Static and configure the required parameters in the Test Multi-Segment PW (Static) panel for the dynamic PW.

8

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

If you are configuring a VCCV ping test for a VLL with a switching site, you must set the Reply Type parameter to IP.

9

Click on the Results Configuration tab and configure the required parameters.

10

Save the changes and close the forms.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.9 To create and run VCCV trace OAM diagnostic test from the STM

90.9.1 Steps

You can also perform this diagnostic test from a VLL service form, or from the service flat topology map. To create a VCCV trace OAM diagnostic test from a static PW to a dynamic PW segment, perform 90.10 “To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM” (p. 3010) . To create a VCCV trace OAM diagnostic test from a dynamic PW to a static PW segment, perform 90.11 “To create and run a VCCV trace OAM diagnostic from a dynamic PW to a static PW segment from the STM” (p. 3012) .

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service→VCCV Trace. The VCCV Trace (Create) form opens.

3

Configure the required parameters.

4

Set the Test Type parameter in the Test Object panel to Regular.

5

Select a spoke SDP binding in the First Spoke SDP Binding panel.

6

Configure the required parameters.

7

Set the Target FEC Type parameter to None.

8

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

If you are configuring a VCCV trace test for a VLL with a switching site, you must set the Reply Type parameter to IP.

9

Click on the Results Configuration tab and configure the required parameters.

10

Save the changes and close the forms.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.10 To create and run a VCCV trace OAM diagnostic from a static PW to a dynamic PW segment from the STM

90.10.1 Steps

You can perform this test from a VPLS/MVPLS or VLL service form, or from the service flat topology map.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service→VCCV Trace. The VCCV Trace (Create) form opens.

3

Configure the required parameters.

4

Set the Test Type parameter in the Test Object panel to Static.

5

Configure the required parameters.

6

Configure the Send Mode parameter in the MPLS-TP VCCV Trace panel.

7

Select the required spoke or mesh SDP binding in the First Spoke/Mesh SDP Binding panel.

Note: The selection of spoke bindings is applicable to both VPLS and VLL services; mesh bindings are only applicable to VPLS and MVPLS. However, when creating this test for MVPLS and VPLS, the concept of downstream spoke SDP bindings does not apply.

8

To configure this test for a VLL, select the required downstream spoke SDP binding in the Downstream SDP Binding panel.

9

Set the FEC Type parameter to FEC 128.

10

Configure the required parameters in the Test Multi-Segment PW (FEC 128) panel for the dynamic PW.

The PW ID parameter is the Service ID of the service that contains the dynamic PW between the source and destination NEs in this segment.

11

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

If you are configuring a VCCV trace test for a VLL with a switching site, you must set the Reply Type parameter to IP.

12

Click on the Results Configuration tab and configure the required parameters.

13

Save the changes and close the forms.

14

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.11 To create and run a VCCV trace OAM diagnostic from a dynamic PW to a static PW segment from the STM

90.11.1 Steps

You can perform this test from a VPLS/MVPLS or VLL service form, or from the service flat topology map.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service→VCCV Trace. The VCCV Trace (Create) form opens.

3

Configure the required parameters.

4

Set the Test Type parameter in the Test Object panel to Regular.

5

Select the required spoke or mesh SDP binding in the First Spoke/Mesh SDP Binding panel.

Note: The selection of spoke bindings is applicable to both VPLS and VLL services; mesh bindings are only applicable to VPLS and MVPLS. However, when creating this test for MVPLS and VPLS, the concept of downstream spoke SDP bindings does not apply.

6

To configure this test for a VLL, select the required downstream spoke SDP binding in the Downstream SDP Binding panel.

7

Set the FEC Type parameter to Static.

8

Configure the required parameters in the Test Multi-Segment PW (Static) panel for the dynamic PW.

9

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

If you are configuring a VCCV trace test for a VLL with a switching site, you must set the Reply Type parameter to IP.

10

Click on the Results Configuration tab and configure the Required parameters.

11

Save the changes and close the forms.

12

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.12 To create and run a MAC populate OAM diagnostic test from the STM

90.12.1 Steps

You can also perform this diagnostic test from the Tests tab on a VPLS or Epipe VLL service form.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→MAC Populate. The MAC Populate (Create) form opens.

3

Configure the required parameters.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

4

Click on the Results Configuration tab and configure the required parameters.

5

Save the changes and close the forms.

6

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.13 To create and run a MAC purge OAM diagnostic test from the STM

90.13.1 Steps

You can also perform this diagnostic test from the Test tab of a VPLS or an Epipe VLL.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→MAC Purge. The MAC Purge (Create) form opens.

3

Configure the required parameters.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

4

Click on the Results Configuration tab and configure the required parameters.

5

Save the changes and close the forms.

6

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.14 To create and run a MAC ping OAM diagnostic test from the STM

90.14.1 Steps

You can also perform this diagnostic test from the Test tab of a VPLS or an Epipe VLL.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→MAC Ping. The MAC Ping (Create) form opens.

3

Configure the required parameters.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.15 To create and run a MAC trace OAM diagnostic test from the STM

90.15.1 Steps

You can also perform this diagnostic test from the Test tab of a VPLS or an Epipe VLL.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→MAC Trace. The MAC Trace (Create) form opens.

3

Configure the required parameters.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.16 To create and run a CPE ping OAM diagnostic test from the STM

90.16.1 Additional information

You can also perform this diagnostic test from:

- the Test tab→CPE Ping sub-tab of the required VPLS or VLL Epipe service provisioned over a VPLS PBB backbone
- the topology view of the required VPLS or VLL Epipe service

This test can also be added to the FirstRun/LastRun in a test suite.

If the Epipe or the B-VPLS (in the case of PBB Epipe) uses SDP bindings, the system configuration must be network chassis mode D compatible.

To run this test for a VLL Epipe over a VPLS PBB backbone, the NEs involved must be at IOM3 or above.

90.16.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→CPE Ping. The CPE Ping (Create) form opens.

3

Configure the required parameters.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

The Source MAC Address parameter is only displayed and configurable when you are creating this test for a VPLS.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

The Send Via Control Plane and Reply Via Control Plane parameters are only configurable for a VPLS.

5

Click on the Results Configuration tab and configure the required parameters.

6

Click on the NM Thresholds tab to configure threshold-crossing alarms.

Note: The NM Thresholds tab will only be displayed and configurable when the NE Schedulable parameter is not enabled. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

7

Save the changes and close the forms.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.17 To create and run an ANCP loopback OAM diagnostic test from the STM

90.17.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L2 Service→ANCP Loopback. The ANCP Loopback (Create) form opens.

3

Configure the required parameters.

4

Select the originating NE in the From IP Address panel.

5

Click on the Test Parameters tab.

6

To configure the ANCP loopback diagnostic for a subscriber ID:

1. Set the Target Type parameter to Subscriber Ident String.
2. Configure the Subscriber Ident String parameter in the ANCP Target Type Details panel.

7

To configure the ANCP loopback diagnostic for an ANCP string:

1. Set the Target Type parameter to ANCP String.
2. Configure the ANCP String parameter in the ANCP Target Type Details panel.

8

Configure the required parameters in the ANCP Details panel.

9

Click on the Results Configuration tab and configure the required parameters.

Note: The ANCP OAM test returns a result trap only if it can locate the subscriber line ANCP string. In this case, it correctly returns the failure or success message. If the ANCP string (i.e., the subscriber line) does not exist on the DSLAM, the network element cannot ping it and does not return a trap. The OAM test in NFM-P remains in the Pending state until the request times out. This behavior differs from the behavior of the other OAM tests. To correct the problem, check the Trap Generation: Probe Failure check box.

10

Save the changes and close the forms.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.18 To create and run a VXLAN ping OAM diagnostic test from the STM

90.18.1 Steps

Perform this test to run a VXLAN ping test to confirm that a virtual tunnel is operational using the echo request/reply functions, as well as query the status of the Virtual Network Instance (VNI). VXLAN exists only on VPLS and this test is used to ping from site to site in the service. The VXLAN ping test can also be created from a VPLS service and from VPLS sites under the OAM→Tests→VXLAN tab. This functionality is also available in the VPLS service topology map.

1

Choose Manage→Service→Services from the NFM-P main menu. The Service Manager form opens.

2

Choose the required service from the list and click Properties. The *Service* (Edit) form opens.

3

Choose a site from the navigation tree. The Site information is displayed.

4

Click on the BGP tab, then on the General sub-tab.

1. Enable the Enable BGP parameter.
2. Configure the Route Distinguisher parameter.
3. Click RT and PW Template Configuration to open the BGP Configuration form.
4. Click the VSI Import Policies tab and configure required Import Route Target.
5. Click the VSI Export Policies tab and configure required Export Route Target.

5

Click on the EVPN tab.

6

Configure the required parameters.

Note: The Enable BGP EVPN parameter must be enabled for this test.
The CFM MAC Advertisement parameter allows the advertisement through BGP of any configured Up MEP, vMEP and MIP MAC addresses on SAPs and SDP bindings. These are not advertised by default. When this parameter is enabled, locally-generated ETH-CFM PDUs utilize the TEP and VTEP connections as transport for EVPN services. CFM will inform and keep the STM up to date on MEP and MIP MAC addresses, including any additions, deletions, and changes. If the CFM MAC Advertisement parameter is disabled, the Up MEP, vMEP and MIP MAC addresses will be withdrawn in BGP.

7

Click Create in the VXLAN panel. The VXLAN (Create) form opens.

8

Configure the Network Identifier parameter and close the form.

Note: The Network Identifier parameter must have the same value for all sites to be included in the test.

9

Repeat Step 4 to Step 8 for all sites in the service which are to be included in the test.

10

Close the service form.

11

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

12

Click Create and choose L2 Service→VXLAN Ping. The VXLAN Ping (Create) form opens.

13

Configure the required parameters on the General tab.

Note: The Outer IP Destination parameter determines which site in the service to test.

14

Click on the Test Parameters tab and configure the required parameters.

Note: Enabling the I-Flag parameter specifies that the OAM PDUs contain a valid VNI. If it is not enabled, the OAM PDUs are prevented from being forwarded beyond the terminating Virtual Tunnel Endpoints (VTEPs).

Note: You must set the Relay Mode parameter to the Overlay option when the test uses non-system IP addressing.

15

Click on the Results Configuration tab and configure the required parameters.

16

Click Execute to run the test.

17

Click on the Results tab and select the test, then click Properties.

The VXLAN Ping Result form opens. Test results are displayed on the General, Response Probes, and Details tabs.

18

See 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) for additional information on running OAM diagnostic tests and 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

19

Save the changes and close the forms.

END OF STEPS

90.19 To create and run a VPRN ping or VPRN trace OAM diagnostic test from the STM

90.19.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L3 Service→VPRN Ping or VPRN Trace. The VPRN Ping (Create) form or VPRN Trace (Create) opens.

3

Configure the required parameters.

4

Select a VPRN site In the Test Object Panel.

5

Click on the Test Parameters tab.

6

Configure the parameters in the Execution Details panel.

7

Configure the parameters in the Test Probe panel.

The Time To Live parameter only applies to VPRN ping tests.

The Initial Time to Live, DiffServ Field and Maximum Time to Live parameters only apply to VPRN trace tests.

8

Click on the NM Thresholds tab to configure threshold-crossing alarms. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

9

Save the changes and close the forms.

10

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.20 To create and run a VPRN Ping, VPRN Trace, ICMP Ping, or ICMP Trace OAM diagnostic test from a service manager form

90.20.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Service Manager form opens.

2

Choose a service from the list and click Properties. The *Service* (Edit) form opens.

3

Click on the Sites tab.

4

Choose a site or sites from the list and click Properties. The Site (Edit) form opens.

5

Click on the OAM tab, then the Tests tab.

6

Configure and run a VPRN ping or VPRN trace OAM diagnostic.

a. To create and run a VPRN ping:

1. Click on the VPRN Ping tab. A list of VPRN diagnostics appears.
2. Double-click on a row in the list to edit an existing test, or click Create to create a new test. The VPRN ping form opens.
3. Configure the required parameters.
4. Select a VPRN site.
5. Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

6. Click on the Results Configuration tab and configure the required parameters.
7. Click Apply to save the changes and confirm the action.
8. Click Execute. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

b. To create and run a VPRN trace:

1. Click on the VPRN Trace tab. A list of VPRN trace diagnostics appears.
2. Double-click on a row in the list to edit an existing test, or click Create to create a new test. The VPRN Trace (Create) form opens.
3. Configure the required parameters.
4. Select a VPRN site.
5. Click on the Test Parameters tab and configure the required parameters.

   The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.
6. Click on the Results Configuration tab and configure the required parameters.
7. Click Apply to save the changes and confirm the action.
8. Choose the diagnostic from the list.
9. Click Execute. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

7

Configure and run an ICMP ping or ICMP trace OAM diagnostic.

a. To create and run ICMP ping:

1. Click on the ICMP Ping tab. A list of ICMP diagnostics appears.
2. Double-click on a row in the list to edit an existing test, or click Create to create a new test. The ICMP Ping form opens.
3. Configure the required parameters.
4. Configure the IP Address and VPRN interface in the Source panel.
5. Configure the Target IP Address and VPRN interface in the Target IP Address panel.
6. Click on the Test Parameters tab and configure the required parameters.

   The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

   The Continuous Execution option is only applicable when the Packet Interval is set to 1 second or greater.

   The Continuous Execution and Rapid options are mutually exclusive.
7. Click on the Results Configuration tab and configure the required parameters.

8. Click Apply to save the changes and confirm the action.
9. Click Execute. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

b. To create and run ICMP trace:

1. Click on the ICMP Trace tab. A list of ICMP trace diagnostics appears.
2. Double-click on a row in the list to edit an existing test, or click Create to create a new test. The ICMP Trace form opens.
3. Configure the required parameters.
4. Configure the IP Address and VPRN interface in the Source panel.
5. Configure the Target IP Address and VPRN interface in the Target IP Address panel.
6. Click on the Test Parameters tab and configure the required parameters.

   The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.
7. Click on the Results Configuration tab and configure the required parameters.
8. Click Apply to save the changes and confirm the action.
9. Choose the diagnostic from the list.
10. Click Execute. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

8

View the test results on the Results tab. The results depend on the type of test. See 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960).

Result information includes:

- Number of Probes Sent
- Time Last Response
- Number of Responses Received

Note: Results of individually run ICMP Trace tests are viewed on the Results tab. The Results tab displays only the result of the last individually run ICMP Trace test, any previous individually run test results are overwritten.
Results from ICMP Trace tests that are scheduled or part of a test suite are also stored on the Results tab. The number of scheduled test results stored corresponds to the value configured in the Probe History Size (rows) parameter. Scheduled ICMP Trace test results do not overwrite individually run test results or previously run scheduled test results.

END OF STEPS

90.21 To create a tunnel ping OAM diagnostic test from the STM

90.21.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service Transport→Tunnel Ping. The Tunnel Ping (Create) form opens.

3

Configure the required parameters.

4

Select the source of the service tunnel.

5

Save the changes and close the forms.

6

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.22 To create and run a tunnel ping OAM diagnostic test from a service tunnel

90.22.1 Steps

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Filter to list only the source and destination routers of the service tunnel and click Search. The list of service tunnels opens.

3

Double-click on a service tunnel from the list. The Tunnel (Edit) form opens.

4

Click on the Tests tab.

5

Click on the Tunnel Ping tab.

6

Click Create. The Tunnel Ping (Create) form opens. The form displays information about the circuit being tested, including the originating tunnel ID.

7

Configure the required parameters.

8

Click on the Test Parameters tab and configure the required parameters.

The Packet Interval (seconds) parameter has an effect only when multiple probes are to be sent.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

Ensure that you configure the Forwarding Class parameter to work with the services that use the tunnel.

9

Click on the Results Configuration tab and configure the required parameters.

10

Click Apply to save the changes and confirm the action.

11

Perform the tunnel ping OAM diagnostic:

1. Click Execute on the Tunnel Ping (Edit) form. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results. The diagnostic is complete.
2. Click on the Results tab. The list of tunnel ping OAM packets sent is displayed.
3. Click on the row or rows to view diagnostic information.
4. Click Properties. The OAM results form opens.

See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about the diagnostic status messages. Use the status message to interpret the diagnostic results.

5. Close the form.

12

Close the Tunnel (Edit) form when the OAM diagnostics are complete.

END OF STEPS

90.23 To create and run an MTU ping OAM diagnostic test from the STM

90.23.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Service Transport→MTU Ping. The MTU Ping (Create) form opens.

3

Configure the required parameters.

4

Select the source of the service tunnel.

5

Save the changes and close the forms.

6

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.24 To create and run an MTU ping OAM diagnostic test from a service tunnel

90.24.1 Steps

Use the MTU ping diagnostic test to find the largest valid frame size.

1

Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.

2

Filter to list only the source and destination routers of the service tunnel and click Search. The list of service tunnels appears.

3

Double-click on a service tunnel from the list. The Tunnel (Edit) form opens.

4

Click on the Tests tab.

5

Click on the MTU Ping tab.

6

Click Create. The MTU Ping (Create) form opens. The form displays information about the service tunnel being tested and the originating tunnel ID.

7

Configure the required parameters.

8

Click on the Test Parameters tab and configure the required parameters.

Note: If the step size is small but the end message size is large, the amount of time to complete the MTU ping may be many minutes. Ensure that you have an appropriate step size that reflects the range of MTU packet sizes that you want to test.

9

Click on the Results Configuration tab and configure the required parameters.

10

Click Apply to save the changes.

11

Perform the MTU ping.

1. Click execute in the MTU Ping (Edit) form. The MTU ping diagnostic starts. A deployed test is created and run. Open the deployed test from the Deployed Tests tab to view its current state. When the test is complete, the deployed test is removed, and you can view the results. The diagnostic is complete.

2. Click on the Results tab. The list of MTU ping OAM probes sent is displayed.
3. Click on the row(s) to view diagnostic information.
4. Click Properties. The OAM results form opens.

   See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about the diagnostic status messages. Use the status message to interpret the diagnostic results. For example, the status message Response Received indicates that the MTU OAM diagnostic completed successfully.
5. Close the form.

12

Close the Tunnel (Edit) form when the OAM diagnostics are complete.

END OF STEPS

90.25 To create and run a MPLS LSP ping OAM diagnostic test from the STM

90.25.1 Steps

The MPLS LSP Ping OAM diagnostic test is supported on both IPv4 and IPv6 source and destination addresses.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose MPLS→LSP Ping. The LSP Ping (Create) form opens.

3

Configure the required parameters. The available parameters change depending on the option specified for the Target Type parameter. Information about NFM-P parameters is available in the XML API Reference.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

As required, click on the NM Thresholds tab to configure threshold-crossing alarms.

Note: The NM Thresholds tab will only be displayed and configurable when the NE Schedulable parameter is not enabled. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

7

Save the changes and close the forms.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.26 To create and run a MPLS LSP trace OAM diagnostic test from the STM

90.26.1 Steps

The MPLS LSP Trace OAM diagnostic test is supported on both IPv4 and IPv6 source and destination addresses.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose MPLS→LSP Trace. The LSP Trace (Create) form opens.

3

Configure the required parameters. The available parameters change depending on the option specified for the Target Type parameter. Information about NFM-P parameters is available in the XML API Reference.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

5

Click on the Results Configuration tab and configure the required parameters.

6

As required, click on the NM Thresholds tab to configure threshold-crossing alarms.

Note: The NM Thresholds tab will only be displayed and configurable when the NE Schedulable parameter is not enabled. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

7

Save the changes and close the forms.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.27 To create and run a MPLS LDP tree trace OAM diagnostic test from the STM

90.27.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose MPLS→LDP Tree Trace. The LDP Tree Trace (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

As required, click on the NM Thresholds tab to configure threshold-crossing alarms.

Note: The NM Thresholds tab will only be displayed and configurable when the NE Schedulable parameter is not enabled. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

7

Save the changes and close the forms.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.28 To create and run a MPLS P2MP LSP ping OAM diagnostic test from the STM

90.28.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose MPLS→P2MP LSP Ping. The P2MP LSP Ping (Create) form opens.

3

Configure the required parameters.

4

Perform one of the following:

a. If you selected P2MP RSVP as your Target Type:

1. Select a P2MP dynamic LSP in the P2MP Dynamic LSP panel.
2. Select a P2MP Instance in the P2MP Instance panel.
3. Configure the Select All S2L Paths parameter.

   If you enabled the parameter, go to Step 5 .
4. Click on the S2L Paths tab.
5. Click Add and select one or more S2L paths (up to a maximum of five).

b. If you selected P2MP mLDP as your Target Type:

Note: For the test to execute successfully, ensure the following:

- The LDP protocol must be enabled on all NEs between the source and destination, including any intermediate NEs.
- LDP interfaces must be in place between the source and destination NEs.
- LDP interfaces must be up, and the Multicast Forwarding parameter must be enabled on each interface. This also applies to all intermediate NEs. See 28.52 “To configure an LDP interface” (p. 944) for information about configuring LDP interfaces.

Perform the following:

1. Configure the required parameters.
2. Select the P2MP ID you want to test in the P2MP LDP panel.
3. Click on the LDP Leaf Addresses tab.
4. Click Add and select up to five leaf NEs.

c. If you selected P2MP mLDP-SSM as your Target Type:

Note: For the test to execute successfully, ensure the following:

- The LDP protocol must be enabled on all NEs between the source and destination, including any intermediate NEs.
- LDP interfaces must be in place between the source and destination NEs.
- LDP interfaces must be up, and the Multicast Forwarding parameter must be enabled on each interface. This also applies to all intermediate NEs. See 28.52 “To configure an LDP interface” (p. 944) for information about configuring LDP interfaces.

Perform the following:

1. Select the System ID (Loopback IP Address) in the P2MP LDP SSM panel.
2. Configure the Router Instance type (either a base router or a VPRN service). If you select Service, then you must also select the required VPRN Instance.
3. Select the Source Address. The Group Address field will be automatically populated, depending on your choice of Source Address. However, you can also manually enter a multicast address in the Group Address field, as required.
4. If you selected Service as the Router Instance type, then go to Step 5 .
5. Select a Sender Address.
6. Click on the LDP Leaf Addresses tab.
7. Click Add and select up to five leaf NEs.

d. If you selected P2MP Policy as your Target Type, perform the following:

1. Select a System ID (Loopback IP Address).
2. Select the Root Address, Root Tree Id, and Root Tree Instance Id.
3. Click on the Leaf Addresses tab.
4. Click Add and select up to five leaf addresses.

5

Click on the Test Parameters tab and configure the required parameters.

The Time To Live parameter is only displayed if you selected P2MP RSVP as your Target Type.

6

Click on the Results Configuration tab and configure the required parameters.

7

As required, click on the NM Thresholds tab to configure threshold-crossing alarms. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

8

Click OK to save the changes. The Service Test Manager form opens.

9

Choose the created test from the list of OAM diagnostics and click Properties.

The P2MP LSP Ping (Edit) form opens.

10

Click Execute to start the P2MP LSP ping. A deployed test is created and run. Open the deployed test to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

Note: When this test is performed on an existing Point-to-Multipoint LSP, up to five S2L paths (for a P2MP RSVP target type) or five LDP leaf addresses (for a P2MP mLDP or P2MP mLDP-SSM target type) can be selected. If none are selected, all available paths or addresses up to the limit of five will be pinged.

11

Click on the Results tab in the P2MP LSP Ping (Edit) form.

The Results tab will show an entry for each S2L path (RSVP) or leaf address (mLDP or mLDP-SSM) tested.

12

Select the test from the list and click on Properties. The P2MP LSP Ping Packet Result form opens.

The General tab displays the test parameters you configured, along with time the test was executed, its status, and the following test information:

- Number of Packets Sent
- Packet Timeouts
- Number of Responses Received

- Number of Packets Lost
- Loss Percentage
- Last Good Packet Time

Note: The test results depend on the type of test you run. See 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960) for more information.

13

Click on the Details tab.

14

View the test results for the:

- Round Trip Time
- Outbound One Way Time
- Inbound One Way Time

15

Click on the Response Packets tab.

16

Select a response packet from the list and click on Properties. The P2MP LSP Ping Packet Result form opens.

Note: The test results depend on the type of test you run. See 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960) for more information.

17

View details on the response packet, including:

- Status
- Round-Trip Time
- Response sequence

18

Repeat Step 12 to Step 17 for all tested S2L paths (RSVP) or leaf addresses (mLDP).

END OF STEPS

90.29 To create and run a MPLS P2MP LSP trace OAM diagnostic test from the STM

90.29.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose MPLS→P2MP LSP Trace. The P2MP LSP Trace (Create) form opens.

3

Configure the required parameters.

4

Click Select in the P2MP Dynamic Lsp panel to choose the P2MP Dynamic LSP you want to test.

5

Click on Select in the P2MP Instance panel to choose the P2MP Instance you want to test.

6

Click on Select in the S2L Destination Address panel to choose the S2L Destination Address you want to test.

7

Click on the Test Parameters tab and configure the required parameters.

8

Click on the Results Configuration tab and configure the required parameters.

9

As required, click on the NM Thresholds tab to configure threshold-crossing alarms. See 89.11 “To configure threshold-crossing alarms or NM threshold-crossing alarms for an OAM diagnostic test or STM test policy” (p. 2949) for more information.

10

Click OK to save the changes. The Service Test Manager form opens.

11

Choose the created test from the list of OAM diagnostics and click Properties.

The P2MP LSP Trace (Edit) form opens.

12

Click execute to start the P2MP LSP trace. A deployed test is created and run. Open the deployed test to view its current state. When the test is complete, the deployed test is removed, and you can view the results.

13

Click on the Results tab.

14

Select the test from the list and click on Properties. The P2MP LSP Trace Result form opens.

The General tab displays the test parameters you configured, along with time the test was executed and its status.

Note: The test results depend on the type of test you run. See 89.16 “To view and compare OAM diagnostic test results on the STM” (p. 2960) for more information.

15

Click on the Hops and Probes tab.

A tree view of the Hops and Probes associated with this test is provided.

16

Right-click on a Hop from the tree view and click on Properties.

The P2MP LSP Trace Hop (Edit) form is displayed, showing the General tab.

17

View the test results for the following:

- Response Probe
- Round Trip Details
- Outbound One Way Trip Details
- Inbound One Way Trip Details

18

Right-click on a Probe from the tree view and click on Properties.

The P2MP LSP Trace Probe (Edit) form is displayed, showing the General tab.

19

View details on the probe, including:

- Response Probe

- LSP Details

20

Repeat this test for all required S2L paths. The test only evaluates one S2L path per execution.

END OF STEPS

90.30 To create and run an ATM ping OAM diagnostic test from the STM

90.30.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L1/L2→ATM Ping. The ATM Ping (Create) form opens.

3

Configure the required parameters.

The VCI parameter is only configurable when the Connection Type is set to VCC Connection.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.31 To configure an ATM OAM loopback from a device Properties form

90.31.1 Steps

The ATM OAM loopback settings are configured globally on supported devices. However, a loopback must be enabled on an IES or VPRN SAP.

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Right-click on the device where you want to configure an ATM OAM loopback and select Properties. The Network Element (Edit) form opens.

3

Click on the ATM tab.

4

Configure the required parameters.

5

Save your changes and close the form.

END OF STEPS

90.32 To create and run a BIER ping OAM diagnostic test from the STM

90.32.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→BIER Ping. The BIER Ping (Create) form opens.

3

Configure the required parameters on the General tab.

4

Select a source site in the Site panel.

5

Depending on the option selected for the BIER Target Type parameter in Step 3, perform one of the following:

a. For the BFR ID option:
   1. Configure the Sub-Domain ID parameter.
   2. Select a destination BFR ID in the BFR ID panel.

      The BFR ID must be in the configured sub-domain.

b. For the BFR ID Range option:
   1. Configure the Sub-Domain ID parameter.
   2. Configure a starting ID for the range. Click Select for the BFR ID Start parameter and choose a BFR ID from the list.
   3. Configure an ending ID for the range. Click Select for the BFR ID End parameter and choose a BFR ID from the list.

      All BFR IDs in the range must be in the configured sub-domain.

c. For the BFR Prefix option:
   1. Configure the Sub-Domain ID parameter.
   2. Click on the BFR Prefix tab.
   3. Click Create. The BIER Ping BFR Prefix (Create) form opens.
   4. Configure the parameters and select a Prefix address.

      The BFR prefix must be in the configured sub-domain.
   5. Close the BIER Ping BFR Prefix (Create) form.
   6. Configure additional BFR prefixes as required. You can configure up to 16 BIER BFR prefixes.

6

Click on the Test Parameters tab and configure the required parameters.

7

Click on the Results Configuration tab and configure the required parameters.

8

Save the changes and close the forms.

9

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.33 To create and run a BIER trace OAM diagnostic test from the STM

90.33.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→BIER Trace. The BIER Trace (Create) form opens.

3

Configure the required parameters on the General tab.

4

Select a source site in the Site panel.

5

Depending on the option selected for the BIER Target Type parameter in Step 3, perform one of the following:

a. For the BFR ID option:
   1. Configure the Sub-Domain ID parameter.
   2. Select a destination BFR ID in the BFR ID panel.

      The BFR ID must be in the configured sub-domain.

b. For the BFR Prefix option:
   1. Configure the Sub-Domain ID parameter.
   2. Select a destination prefix in the BFR Prefix panel.

      The BFR prefix must be in the configured sub-domain.

6

Click on the Test Parameters tab and configure the required parameters.

7

Click on the Results Configuration tab and configure the required parameters.

8

Save the changes and close the forms.

9

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.34 To create and run an MFIB ping OAM diagnostic test from the STM

90.34.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→MFIB Ping. The MFIB Ping (Create) form opens.

3

Configure the required parameters on the General tab.

When the Service Name parameter is configured, you can configure the Name parameter for the site.

When you set the target IP address, then choose a service ID, only service IDs from the selected site are available.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.35 To create and run an Mrinfo OAM diagnostic test from the STM

90.35.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→Mrinfo. The Mrinfo (Create) form opens.

3

Configure the required parameters.

4

Click on Select to choose a PIM interface. The Select PIM Interface form opens.

5

Select a PIM interface in the list and click OK. The Select PIM Interface form closes, and the interface information is displayed on the Mrinfo (Create) form.

6

Click on the Results Configuration tab and configure the required parameters.

7

Save the changes and close the forms.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.36 To create and run an Mtrace OAM diagnostic test from the STM

90.36.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→Mtrace. The Mtrace (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

90.37 To create and run an Mtrace2 OAM diagnostic test from the STM

90.37.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Multicast→Mtrace2. The Mtrace2 (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959). See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.38 To create and run an ICMP ping OAM diagnostic test from the STM

90.38.1 Steps

You can also perform an ICMP ping test from the Test tab of a VPRN or EIS service configuration form.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose ICMP→ICMP Ping. The ICMP Ping (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

The Continuous Execution option is only applicable when the Packet Interval is set to 1 second or greater.

The Continuous Execution and Rapid options are mutually exclusive.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.39 To create and run an ICMP trace OAM diagnostic test from the STM

90.39.1 Steps

You can also perform an ICMP trace test from the Test tab of a VPRN or EIS service configuration form.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose ICMP→ICMP Trace. The ICMP Trace (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.40 To create and run an ICMP DNS ping OAM diagnostic test from the STM

90.40.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose ICMP→DNS Ping. The DNS Ping (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.41 To create and run a PRBS test

90.41.1 Before you begin

Ensure that the administrative state of the port where the test is run and the port where the loopback is configured are set to Maintenance.

90.41.2 Steps

Set the administrative state to Maintenance

1

On the equipment tree, expand Network→1830 VWM→Shelf 1830 VWM TLU →Card Slot TLU.

2

Right-click on the port object where the PRBS test will be executed and choose Properties. The Physical Port (Edit) form opens.

3

Click on the States tab and set the Administrative State parameter to Maintenance.

4

Save your changes and close the form.

5

Repeat Step 1 to Step 4 for the port where you need to configure loopback.

Configure loopback

6

On the equipment tree, expand Network→1830 VWM→Shelf 1830 VWM TLU →Card Slot TLU.

7

Right-click on the port object where you need to configure the loopback and choose Properties. The Physical Port (Edit) form opens.

8

Click on the Port Specifics tab. In the Loop Back panel, set the appropriate LoopBack Status parameter to Enabled.

9

Save your changes and close the form.

Start the PRBS test

10

Perform one of the following:

a. From the main menu:
   1. Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.
   2. Click Create→VWM→PRBS Test. The PRBS Test (Create) form opens.

b. From the port:
   1. On the equipment tree, expand Network→1830 VWM→Shelf 1830 VWM TLU →Card Slot TLU.
   2. Right-click on the port object where you need to perform the PRBS test and choose Properties. The Physical Port (Edit) form opens.
   3. Click PRBS Test. The PRBS Test (Create) form opens.

11

Configure the required parameters and click on the PRBS tab.

12

Configure the required parameters.

Note: The parameters are populated when the test is performed from the Physical Port (Edit) form.

13

Click Apply. The PRBS Test (Create) form closes and the PRBS Test (Edit) form appears.

14

Click Execute to run the test.

Stop the PRBS test

15

Perform one of the following to stop the test:

a. Click Stop Test Execution to stop the test manually.

b. Perform the following to stop the test automatically:
   1. Click on the PRBS tab and configure the Auto Stop and Auto Stop Duration parameters.
   2. Click Apply.

View the results

16

Click on the Results tab and click search. The results are listed.

17

Choose an entry and click Properties. The PRBS Test Result form opens.

Result: The results of the test are obtained by reading the number of errors detected in the received PRBS signal in the 1830 VWM TLU or 1830 VWM ITP.

END OF STEPS

90.42 To create and run a OmniSwitch CPE SLA diagnostic test from the STM

90.42.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose L1/L2→ATM Ping. The ATM Ping (Create) form opens.

3

Configure the required parameters.

4

Click on the Test Parameters tab and configure the required parameters.

5

Click on the Results Configuration tab and configure the required parameters.

6

Save the changes and close the forms.

7

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

90.43 To configure and run OAM tests contextually

90.43.1 Reference

A number of the OAM diagnostic tests for objects and services described in this chapter can be contextually configured and executed from various entry points in the GUI. Specifically, VPLS, VPRN, Epipe, and Composite services, in addition to SAPs (access interfaces), SDP tunnels, SDP bindings, and LSPs can all be tested contextually. All of the tests are available in both Quick and Scheduled runs. Some test types also support a continuous run option. Scheduled tests use SNMP and continuous tests use accounting.

Contextual test creation and execution is available from:

- Service topology maps and flat maps (for Composite services)
- Sites and Interfaces tabs within a service properties form
- Service Tree within a service properties form
- Manage Service Tunnels form
- Applicable Manage LSP forms

This procedure is applicable to all of these access modes, since the contextual test forms that are opened are the same, regardless of the access mode you use.

Contextual tests can be created between selected service sites or between or on specific entities. The tests are accessed by selecting single or multiple entities, then right-clicking and choosing Run OAM Tests.

Refer to the following procedures for additional detailed information on configuring the tests using the STM, if required:

- CFM Loopback - 91.21 “To create and run a CFM loopback OAM diagnostic test from the STM” (p. 3123)
- CFM Link trace - 91.22 “To create and run a CFM link trace OAM diagnostic test from the STM” (p. 3125)
- CFM Two-way delay - 91.24 “To create and run a CFM two way delay OAM diagnostic test from the STM” (p. 3128)
- CFM Two-way SLM - 91.27 “To create and run a CFM two way SLM OAM diagnostic test from the STM” (p. 3132)
- VCCV Ping - 90.8 “To create and run a VCCV ping OAM diagnostic test from the STM” (p. 3009)
- VCCV Trace- 90.9 “To create and run VCCV trace OAM diagnostic test from the STM” (p. 3010)
- VPRN Ping and VPRN Trace- 90.19 “To create and run a VPRN ping or VPRN trace OAM diagnostic test from the STM” (p. 3023)
- Tunnel Ping- 90.21 “To create a tunnel ping OAM diagnostic test from the STM” (p. 3027). Either a single or multiple tunnels can be selected to run the OAM tests contextually.
- MTU Ping- 90.23 “To create and run an MTU ping OAM diagnostic test from the STM” (p. 3029). Either a single or multiple tunnels can be selected to run the OAM tests contextually.
- LSP Ping- 90.25 “To create and run a MPLS LSP ping OAM diagnostic test from the STM” (p. 3031)

- LSP Trace- 90.26 “To create and run a MPLS LSP trace OAM diagnostic test from the STM” (p. 3032)
- ICMP Ping- 90.38 “To create and run an ICMP ping OAM diagnostic test from the STM” (p. 3047)
- ICMP Trace- 90.39 “To create and run an ICMP trace OAM diagnostic test from the STM” (p. 3048)

Refer also to 4.2.3 “Working with Ethernet CFM objects” (p. 177) in Chapter 4, “Topology map management” and 4.8 “To use OAM diagnostic functions on service topology and composite service flat topology maps” (p. 181) for additional information on accessing OAM functionality from a service topology view.

Refer to the following service chapters for additional information, if required:

- VPLS - Chapter 77, “VPLS management ”
- VPRN - Chapter 79, “VPRN service management”
- VLL (for Epipe services) - Chapter 76, “VLL service management”
- Composite Services- Chapter 85, “Composite service management”
- LSPs- Chapter 31, “MPLS”
- SDPs- Chapter 33, “Service tunnels”

Additional procedural points:

- To test SAPs, the SAPs must be operationally Up.
- If you select SAPs or sites with resources such as Global MEGs and MEPs already present, you can utilize these objects and run the tests on top of them.
- To test SDP bindings, the bindings must be operationally Up. The following types are supported:
  1. Regular Spoke SDP Bindings on VPLS, MVPLS, and Epipe services
  2. Regular Spoke SDP Bindings with source interfaces on VPRN services
  3. Mesh SDP Bindings on VPLS and MVPLS with V-sites and B-sites
  4. Static Spoke SDP Bindings on VPLS
  5. Spoke FEC 128 on Epipe services. Testing these types is only available from the table or tree view. The topology view is not supported.
- Contextually-created test policies cannot be used to create regular test suites. When a contextually-created test suite is deleted, the associated test policy is also deleted.
- Contextually-created OAM tests cannot be created between VPLS B-L2 Access Interfaces.
- Contextually-created OAM tests cannot be created between Epipe sites.
- For non-CFM entities, multiple OAM Context windows can be opened to run tests on different entities at the same time.

90.43.2 Steps

To demonstrate the contextual OAM configuration and execution workflow generally, this procedure covers the testing of two SAPs or sites as examples. Testing different entities follows the same basic process. The procedure assumes that you have already accessed either a service form or service topology view.

1

Perform one of the following:

a. Select two or more SAPs and right click on one of them.

b. Select two or more sites and right click on one of them.

2

Select Run OAM Tests from the contextual menu. The OAM Contextual Test form opens.

3

Configure the OAM Test Type parameter for the type of test you want to run.

4

Configure the Test Level and Number of Test Probes, as applicable.

5

Configure and run one of the following test execution modes:

a. Quick Run. Go to Step 6 .

b. Scheduled Run. Go to Step 14 .

c. Continuous Run. Go to Step 26 .

Quick Run tests

6

Configure the Test Execution Mode for Quick Run.

Quick Run tests provide a one-time immediate test result.

7

Click Execute to run the test.

Note: When you execute a test, all required test resources are automatically created by NFM-P. After the test has been executed, the created test resources are automatically deleted. Pre-existing resources that were used for the test are not deleted.

8

To view the test results, click on the Quick Run Results tab and select an entry.

9

Click Properties. The One Time Validation Result form opens. Results for each of the selected test objects are available in the Results tab.

Note: The test results are archived for future reference.

10

Close the form.

11

Click the SAPs or Sites tab to view the objects you chose to run the test. Click View Service to open the property form of the service associated with the selected objects.

12

Click the Test Generation Log tab to view any error messages generated by the test.

13

Close the form.

Scheduled Run tests

14

Configure the Test Execution Mode for Scheduled Run.

15

Enter a Test Suite Name. This name will be used for the automatically-created test suite required for the test. The created test suite can be viewed in the STM after test execution, if required.

16

Click Select for the Test Suite Scheduler and perform one of the following in the Select OAM Context Schedules form:

a. Select an existing schedule form and go to Step 17.

b. Click Create and perform the following:

1. Configure the required parameters in the NFM-P Schedule (Create) form.

   Note:

   CFM scheduled tests do not support a schedule with a frequency less than per minute.

2. Click OK. The NFM-P Schedule (Create) form closes.
3. Click Select for the Test Suite Scheduler and select the newly-created schedule.

17

Click Execute. Tests will be executed as scheduled and the progress bar shows the current test status.

18

To view the test results, click on the Test Suite Results tab and select an entry.

19

Click Properties. The Test Suite Result form opens. General information regarding the test suite results is provided.

20

To view individual test results, select an entry in the Results tab and click Properties. The Test Result form opens. Review the test results.

21

To view the details of the test suite created for the test, click Test Suite View at the lower left-hand corner of the form. The OAM Contextual Test form closes and the Test Suite (Edit) form opens.

You can also view the SAPs or sites you selected for the test in the Contextual Tested Entity tab of the Test Suite (Edit) form.

22

Close the Test Suite (Edit) form.

23

Click the SAPs or Sites tab to view the objects you chose to run the scheduled test. Click View Service to open the property form of the service associated with the selected objects.

24

Click the Test Generation Log tab to view any error messages generated by the test.

25

Close the forms.

Continuous Run tests

26

Configure the Test Execution Mode for Continuous Run.

Continuous Run tests are accounting policy based tests.

27

Enter a Test Suite Name. This name will be used for the automatically-created test suite required for the test. The created test suite can be viewed in the STM after test execution, if required.

28

Click Execute. Tests will be run at 15-minute intervals and the progress bar shows the current test status.

29

To view the test results, click on the Test Suite Results tab and select an entry. If the continuously run test has yet not been run, its Status in the Test Suite Results tab shows as Request Sent.

30

Click Properties. The Test Suite Result form opens. General information regarding the test suite results is provided.

31

To view individual test results, select an entry in the Results tab and click Properties. The Test Result form opens. Review the test results.

32

To view the details of the test suite created for the test, click Test Suite View at the lower left-hand corner of the form. The OAM Contextual Test form closes and the Test Suite (Edit) form opens.

You can also view the SAPs or sites you selected for the test in the Contextual Tested Entity tab of the Test Suite (Edit) form.

33

Close the Test Suite (Edit) form.

34

Click the SAPs or Sites tab to view the objects you chose to run the scheduled test. Click View Service to open the property form of the service associated with the selected objects.

35

Click the Test Generation Log tab to view any error messages generated by the test.

36

Close the forms.

END OF STEPS

90.44 To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script

WARNING

Network Damage

Scripts that are not correctly created or applied can cause serious damage to the network.

Nokia recommends that system administrators clearly define user responsibilities for CLI script usage, and ensure that scripts are verified and validated before they are executed on devices in a live network.

90.44.1 Steps

Perform this procedure to create an OmniSwitch ping or traceroute OAM diagnostic test using a using the sample CLI scripts in Code Figure 90-8, “Sample OmniSwitch OAM ping script” (p. 3059) and Code Figure 90-9, “Sample OmniSwitch traceroute script” (p. 3060) .

1

Choose Tools→Scripts from the NFM-P main menu. The Scripts form opens.

2

Choose CLI Script (Scripting) from the object drop-down menu and click Create. The CLI Script (Create) form opens.

3

Configure the required parameters.

Enable the Use Latest Version parameter to associate all of the targets of the script with the latest version of the CLI script.

You must set the Content Type parameter to Velocity.

4

Perform the following steps to specify the script target types.

1. Click Add in the NE Types panel. The Select Property - CLI Script form opens.
2. Select one or more OmniSwitch NE types in the list and click OK. The Select Property - CLI Script form closes, and the NE types are listed on the CLI Script (Create) form.

5

Click Apply to apply the configuration.

6

Click on the Versions tab.

7

Create the script by clicking Create. The Script Editor *script_name* form opens.

8

Create the CLI script text by performing one of the following steps.

a. Import an existing text file with a CLI script. You can create a text file by copying and pasting the text from the sample OmniSwitch ping shown in Figure 90-8, “Sample OmniSwitch OAM ping script” (p. 3059) or the sample traceroute script shown in Figure 90-9, “Sample OmniSwitch traceroute script” (p. 3060) into a text editor.

   Perform the following steps:

   1. Choose File→Import from the Editor menu. The Import dialog box appears.
   2. Choose the file to be imported and click on Import. The script appears in the Script Editor workspace.
   3. Modify the script as required.

b. Enter or copy and paste the CLI script text from the sample OmniSwitch ping shown in Figure 90-8, “Sample OmniSwitch OAM ping script” (p. 3059) or the sample traceroute script shown in Code Figure 90-9, “Sample OmniSwitch traceroute script” (p. 3060) into the Script Editor workspace.

Note: The following ping and traceroute scripts are sample CLI scripts that can be used to run ping and traceroute tests.

Figure 90-8 Sample OmniSwitch OAM ping script

```
<velocityProperties>
  <tab><name>General</name><tooltip>The general tab</tooltip>
    <group><name>General</name><tooltip>The general group</tooltip>
      <property>
        <name>ip_address</name>
        <uiName>IP Address:</uiName>
        <tooltip>IP address of the system to ping (IPv4 xxx.xxx.xxx.
xxx)</tooltip>
        <type>String</type>
        <default>0.0.0.0</default>
        <uiOrder>1</uiOrder>
        <required>true</required>
      </property>
      <property>
        <name>count</name>
        <uiName>Count:</uiName>
        <tooltip>Number of frames to be transmitted</tooltip>
        <type>Integer</type>
        <default>6</default>
        <uiOrder>2</uiOrder>
        <required>true</required>
      </property>
```

```
      <property>
        <name>packed_size</name>
        <uiName>Packet Size:</uiName>
        <tooltip>Size of the data portion of the packet sent for this
ping, in bytes</tooltip>
        <type>Integer</type>
        <default>64</default>
        <required>true</required>
        <uiOrder>3</uiOrder>
        <min>1</min>
        <max>60000</max>
      </property>
      <property>
        <name>interval</name>
        <uiName>Interval (seconds):</uiName>
        <tooltip>Polling interval</tooltip>
        <type>Integer</type>
        <uiOrder>4</uiOrder>
        <default>1</default>
        <min>1</min>
        <max>10000</max>
      </property>
      <property>
        <name>timeout</name>
        <uiName>Timeout (seconds):</uiName>
        <tooltip>Number of seconds the program will wait for a
response before timing out</tooltip>
        <type>Integer</type>
        <uiOrder>5</uiOrder>
        <default>5</default>
        <min>1</min>
        <max>10000</max>
      </property>
    </group>
  </tab>
</velocityProperties>

ping $ip_address count $count size $packed_size interval $interval
timeout $timeout
```

Figure 90-9 Sample OmniSwitch traceroute script

```
<velocityProperties>
  <tab><name>General</name><tooltip>The general tab</tooltip>
    <group><name>General</name><tooltip>The general group</tooltip>
      <property>
        <name>ip_address</name>
        <uiName>IP Address:</uiName>
        <tooltip>IP address of the host whose route you want to trace.
```

```
(IPv4 xxx.xxx.xxx.xxx)</tooltip>
        <type>String</type>
        <default>0.0.0.0</default>
        <uiOrder>1</uiOrder>
        <required>true</required>
      </property>
      <property>
        <name>maxHopValue</name>
        <uiName>Maximum Hop:</uiName>
        <tooltip>Maximum hop count for the trace</tooltip>
        <type>Integer</type>
        <default>5</default>
        <uiOrder>2</uiOrder>
        <required>true</required>
      </property>
    </group>
  </tab>
</velocityProperties>

traceroute $ip_address max-hop $maxHopValue
```

9

Perform one of the following.

a. Save the script to the Scripts tool.

1. Choose File→Save from the Editor menu or click Save. The Comment form opens.
2. Configure the required parameters.
   - Network Element Version Information
   - Bundle ID
3. Click OK.

b. Export the script to a local or network text file.

1. Choose File→Export from the Editor menu, or click Export. A dialog box appears and prompts you to choose a file storage location in the network.
2. Scroll to the location in which you want to save the text file, and enter a filename in the appropriate field.
3. Click Export. The script version text file is saved in the specified location.

10

Choose File→Close from the Editor menu. The Script Editor form closes, and the CLI Script (Edit) form reappears with the Version tab displayed. The new version of the script appears in the list.

11

Close the forms.

See 90.45 “To configure and run an OmniSwitch OAM diagnostic ping test CLI script” (p. 3062) and 90.46 “To configure and run an OmniSwitch OAM traceroute test CLI script” (p. 3064) for information about configuring and running OmniSwitch ping and traceroute OAM scripts.

End of steps

90.45 To configure and run an OmniSwitch OAM diagnostic ping test CLI script

WARNING

Network Damage

Scripts that are not correctly created or applied can cause serious damage to the network.

Nokia recommends that system administrators clearly define user responsibilities for CLI script usage, and ensure that scripts are verified and validated before they are executed on devices in a live network.

90.45.1 Steps

The following procedure describes how to create and run an OmniSwitch OAM diagnostic ping test using the script created using 90.44 “To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script” (p. 3059) .

1

Ensure that the mediation policy for each OmniSwitch target is configured with the correct user name and password for CLI communication. See 9.17 “To configure device mediation” (p. 301) .

2

Choose Tools→Scripts from the NFM-P main menu. The Scripts form opens.

3

Choose CLI Script (script) from the object drop-down menu.

4

Configure the filter criteria. A list of scripts appears at the bottom of the Scripts form.

5

Double-click on the OmniSwitch ping script that you created using 90.44 “To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script” (p. 3059) . The CLI Script (Edit) form opens.

6

Click on the Targets tab.

7

Click Create. The Target Configuration form opens.

8

Click Create. The Select Network Elements form opens.

9

Choose one or more OmniSwitch devices and click OK. The Select Network Elements form closes, and the OmniSwitch NEs are listed on the Target Configuration form.

10

Configure the test parameters.

1. Choose a target from the Target List panel.
2. Configure the required parameters.
3. Click Apply To Selected.
4. Repeat Step 10 1 to Step 10 3 to configure additional targets, if required.
5. Click OK to close the Target Configuration form. The CLI Script (Edit) form opens with the Targets tab displayed.

11

To run the script on specified targets, choose one or more entries from the list.

12

Click execute. The results of the test appear in the form.

13

Perform one of the following to view the results of the scripts:

a. Choose a target from the list. The results of the script that was run on the specified target appears in the panel below the list.

   Perform the following:

   1. Click Save Result to save the results to the result manager. A dialog box appears. Click OK.

      Note:

      You can save the results of multiple scripts to the result manager simultaneously; choose the entries in the list and click Save Result.

   2. Export the results to a local or network text file.
   3. Click Export. A dialog box appears and prompts you to choose a file storage location on the network.
   4. Scroll to the location in which you want to save the text file, and enter a filename in the appropriate field.

5. Click Export. The results text file is saved in the specified location.

b. Choose one or more entries from the list and click View Selected. The View Selected form opens and displays the results of the script. The results for each target are separated by a comment that indicates the script version, associated target, script status, run time and date, and script parameters.

Perform the following:

1. Export the results of the script to a local or network text file. Choose File→Export from the Editor menu, or click Export. A dialog box appears and prompts you to choose a file storage location in the network. Scroll to the location in which you want to save the text file, and enter a filename in the appropriate field. Click Export. The results text file is saved in the specified location.
2. Choose File→Close from the View Selected form to close the View Selected form.

14

Click OK or Cancel to close the CLI Script (Edit) form.

15

Close the Scripts form.

END OF STEPS

90.46 To configure and run an OmniSwitch OAM traceroute test CLI script

WARNING

Network Damage

Scripts that are not correctly created or applied can cause serious damage to the network.

Nokia recommends that system administrators clearly define user responsibilities for CLI script usage, and ensure that scripts are verified and validated before they are executed on devices in a live network.

90.46.1 Steps

The following procedure describes how to configure and run an OmniSwitch traceroute using the script created using 90.44 “To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script” (p. 3059) .

1

Ensure that the mediation policy for each OmniSwitch target is configured with the correct user name and password for CLI communication. See 9.17 “To configure device mediation” (p. 301) .

2

Choose Tools→Scripts from the NFM-P main menu. The Scripts list form opens.

3

Choose CLI Script (script) from the object drop-down menu.

4

Configure the filter criteria. A list of scripts appears at the bottom of the Scripts form.

5

Double-click on the OmniSwitch traceroute script that you created using 90.44 “To create an OmniSwitch ping or traceroute OAM diagnostic test using a CLI script” (p. 3059) . The CLI Script (Edit) form opens.

6

Click on the Targets tab.

7

Click on Add. The Target Configuration form opens.

8

Click Add in the target list to add an OmniSwitch to the list. The Select Network Elements form opens with a list of OmniSwitch NEs.

9

Choose one or more OmniSwitch NEs and click OK. The Select Network Elements form closes, and the OmniSwitch NEs appear in the target list panel of the Target Configuration form.

10

Configure the test parameters.

1. Choose a target from the Target List panel.
2. Configure the required parameters.
3. Click Apply To Selected.
4. Repeat Step 10 1 to Step 10 3 to configure additional targets.
5. Click OK to close the Target Configuration form. The CLI Script (Edit) appears with the Targets tab displayed.

11

To run the script on the specified targets, choose one or more entries from the list.

12

Click Execute. The test results appear.

13

Perform one of the following to view the results of the scripts:

a. Choose a target from the list. The results of the script that was run on the specified target appears in the panel below the list.

   Perform the following:

   1. Click Save Result to save the results to the result manager. A dialog box appears. Click OK.

      Note:

      You can save the results of multiple scripts to the result manager simultaneously; choose the entries in the list and click Save Result.

   2. Export the results to a local or network text file.
   3. Click Export. A dialog box appears and prompts you to choose a file storage location on the network.
   4. Scroll to the location in which you want to save the text file, and enter a filename in the appropriate field.
   5. Click Export. The results text file is saved in the specified location.

b. Choose one or more entries from the list and click View Selected. The View Selected form opens and displays the results of the script. The results for each target are separated by a comment that indicates the script version, associated target, script status, run time and date, and script parameters.

   Perform the following:

   1. Export the results of the script to a local or network text file. Choose File→Export from the Editor menu, or click Export. A dialog box appears and prompts you to choose a file storage location in the network. Scroll to the location in which you want to save the text file, and enter a filename in the appropriate field. Click Export. The results text file is saved in the specified location.
   2. Choose File→Close from the View Selected form to close the View Selected form.

14

Click OK or Cancel to close the CLI Script (Edit) form.

15

Close the form.

END OF STEPS

90.47 To configure an advanced loopback test on an OmniSwitch from a device Properties form

90.47.1 Steps

Note: Only one test profile per port can be created, up to a maximum of 8 ports on an NE. The test profile cannot be created on a port that is part of a LAG.

1

Choose Equipment from the navigation tree. The navigation tree displays the Equipment view.

2

Expand the OmniSwitch NEs icon.

3

Right-click on the device where you want to configure an advanced loopback test and select Properties. The Network Element (Edit) form opens.

4

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

5

Expand to the port level and click on a port object. The Physical Port (Edit) form opens.

6

Click on the Advanced Loopback tab. Click Create. The AOS Advanced Loopback (Create) form opens.

7

Configure the required parameters.

Note: For OS 6250 and OS 6450 NEs, when the Traffic Type parameter value is set to Outward, the user can specify the SAP Id, provided the SAP Id already exists on the NE. If a loopback profile is created with the Traffic Type parameter value set to Outward, the SAP Id parameter will be provided with the default value of 0. If the loopback profile is created with a SAP Id, the SAP Id must exist on the NE.

8

Save your changes.

9

Choose a test entry on the Physical Port (Edit) form. Click apply. A dialog box appears.

10

Click Yes.

11

Choose a test entry on the Physical Port (Edit) form. Click Properties.

12

The AOS Advanced Loopback (Edit) form opens. Configure the Status parameter.

13

Save your changes and close the forms.

END OF STEPS

90.48 To run the F5 OAM loopback diagnostic test from a 7705 SAR-M/ME Properties form

90.48.1 Procedure support

This procedure is only supported on the 7705 SAR-M/ME equipped with a DSL module.

Note: The parameters for the F5 OAM loopback test are only visible when the DSL module for the 7705 SAR-M/ME is connected to an ISAM device with the bonding type of ATM and the device is operationally up.

90.48.2 Steps

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Navigate to the port object where you want to configure an Ethernet port. The path is Network→*NE*→Shelf→Card Slot→Daughter Card Slot→Port *n/n/n*.

3

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

4

Click on the DSL tab.

5

Choose Active for the F5 Oam Loopback parameter.

6

Click Apply to save your changes. A confirmation window opens.

7

Click Yes. The loopback test is executed.

If the test was successfully executed, the F5 Oam Loopback Status field displays Pass and the F5 Oam Loopback Time field displays the time it took for the test to execute in milliseconds.

8

View the results from the F5 OAM Loopback test by opening a Telnet session on the NE, if required. See 10.4 “To open and close an NFM-P device CLI session” (p. 329) .

END OF STEPS

90.49 To configure an 802.3ah EFM OAM diagnostic test from an NE Properties form

CAUTION

Service Disruption

Performing an 802.3ah EFM diagnostic is service-affecting.

Ensure that you consider the implications of performing this test before you proceed.

90.49.1 Steps

Note: For the 802.3ah EFM diagnostic to be successful, the local and peer ports must support the 802.3ah protocol, and be operationally and administratively up.

When a port is in loopback mode, service mirroring does not work if the port is a mirror source or a mirror destination.

Satellite ports support virtually all OAM functionality. The only exception is that no EFM-OAM link monitoring is available for such ports.

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Right-click on the device where you want to configure an 802.3ah EFM OAM diagnostic and select Properties. The Network Element (Edit) form opens.

3

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

4 Expand to the port level and click on a port object. The Physical Port (Edit) form opens.

5 Click on the Ethernet tab. The Ethernet port configuration form opens.

6 Click on the EFM-OAM tab. The EFM-OAM form opens.

7 Configure the required parameters.

8 Click Apply. A dialog box appears.

9 Click Yes.

10 Click Resync.

The EFM-OAM form is updated with the results of the 802.3ah EFM diagnostic. The values that are displayed on the form depend on the configuration of the local and peer ports.

Table 90-3, “802.3ah EFM OAM results” (p. 3070), Table 90-4, “802.3ah EFM OAM Statistics” (p. 3072) and Table 90-5, “802.3ah EFM OAM peer information” (p. 3073) describe the displayed information.

Table 90-3 802.3ah EFM OAM results

Parameter	Value	Description
Operational Status (dot3OamOperStatus)	Disabled	The Administrative State parameter on the local port is set to Disabled.
	Active Send Local	The Mode parameter is set to Active and the local port is discovering the peer port EFM OAM capabilities.
	Passive Wait	The Mode parameter is set to Passive and the local port is waiting to receive OAMPDUs from the peer device.
	Send Local and Remote	The local port discovered the peer but has not yet accepted or rejected the configuration of the peer. The local port may determine that the peer port is not compatible and decline OAM peering.
	OAM Peering Locally Rejected	The local port declined OAM peering.
	Send Local and Remote Ok	The local port accepted OAM peering.
	OAM Peering Remotely Rejected	The remote port declined OAM peering.
	Operational	The local and remote ports accepted the OAM peering.
	Non-operational Half Duplex	The Administrative State parameter is set to Enabled but the local port is in half-duplex mode.
	Link Fault	The link detected a fault and is transmitting OAMPDUs with a link fault indication.
Max. PDU Size (dot3OamMaxOamPduSize)	PDU size	The largest OAMPDU size, in bytes, that is supported by the local port: Minimum is 64 Maximum is 1518 Default is 1514
Configuration Revision (dot3OamConfigRevision)	0 to 65 535	The configuration revision of the OAM entity obtained from the latest OAMPDU sent by the OAM entity. The value is used by OAM entities to indicate that configuration changes occurred that may require the peer OAM entity to re-evaluate whether OAM peering is allowed.

Table 90-3 802.3ah EFM OAM results (continued)

Parameter	Value	Description
Functions Supported (dot3OamFunctionsSupported)	Event Support	The port can send and receive event notification OAMPDUs.
	Loopback Support	The port can initiate and respond to loopback commands.
	Unidirectional Support	The port supports the transmission of OAMPDUs on links that operate in unidirectional mode.
	Variable Support	The port can send and receive variable request and response OAMPDUs.
Loopback Status (dot3OamLoopbackStatus)	No Loopback	The port is not in a loopback condition.
	Initiating Loopback	The local device initiated a loopback, sent a loopback OAMPDU and is waiting for a response from the peer port.
	Remote Loopback	The peer port is in loopback mode.
	Terminating Loopback	The local port is in the process of terminating a loopback on the peer port.
	Local Loopback	The peer port has put the local port into loopback mode.
	Unknown	An OAMPDU that contains an unexpected message was received by the local port.

Table 90-4 802.3ah EFM OAM Statistics

Statistics	Displayed Value	Description
Frames Lost Due to OAM(dot3OamFramesLostDueToOam)	*Number of frames*	The number of frames that were dropped by the OAM multiplexer. Discontinuities of this counter can occur under some conditions.
Information Rx (dot3OamInformationRx)	*Number of OAMPDUs*	The number of information OAMPDUs received on this interface
Information Tx (dot3OamInformation Tx)	*Number of OAMPDUs*	The number of information OAMPDUs transmitted on this interface. Discontinuities of this counter can occur can occur under some conditions.
Loopback Control Rx (dot3OamLoopbackControlRx)	*Number of OAMPDUs*	The number of loopback control OAMPDUs received on this interface
Loopback Control Tx (dot3OamLoopbackConrolTx)	*Number of OAMPDUs*	The number of loopback control OAMPDUs transmitted on this interface
Unsupported Codes Rx (dot3OamUnsupportedCodesRx)	*Number of OAMPDUs*	The number of OAMPDUs received on this interface with an unsupported opcode

Table 90-4 802.3ah EFM OAM Statistics (continued)

Statistics	Displayed Value	Description
Unsupported Codes Tx (dot3OamUnsupportedCodesTx)	*Number of OAMPDUs*	The number of OAMPDUs transmitted on this interface with an unsupported opcode

Table 90-5 802.3ah EFM OAM peer information

Peer Information	Displayed Value	Description
Peer MAC Address (dot3OamPeerMacAddress)	*MAC address*	The MAC address of the peer port. The MAC address is contained in the received OAMPDU.
Peer Vendor OUI (dot3OamPeerVendorOui)	*OUI*	The OUI of the OAM peer. The OUI is part of the peer MAC address contained in the received OAMPDU. The OUI can be used to identify the vendor of the remote OAM device.
Peer Vendor Info (dot3OamPeerVendorInfo)	*Vendor information text*	The vendor information field is in the local information TLV, and can be used to determine additional information about the peer device.
Peer Mode (dot3OamPeerMode)	Active or Passive	See Mode
Peer Max PDU Size (dot3OamPeerMaxOamPduSize)	*PDU size*	The largest OAMPDU value, in bytes, that is supported by the peer port. Minimum is 64 Maximum is 1518 Default is 1514
Peer Configuration Revision (dot3OamPeerConfigRevision)	0 to 65 535	The configuration revision of the OAM device as identified in the latest OAMPDU sent by the OAM peer. The configuration revision is used by OAM devices to indicate that configuration changes have occurred which might require the peer OAM device to re-evaluate whether OAM peering is allowed.
Peer Functions Supported (dot3OamPeerFunctSupported)	Event Support	The peer port can send and receive Event Notification OAMPDUs.
	Loopback Support	The peer port can initiate and respond to loopback commands.
	Unidirectional Support	The peer port supports the transmission of OAMPDUs on links that are operating in unidirectional mode.
	Variable Support	The peer port can send and receive Variable Request and Response OAMPDUs.

11

Close the EFM-OAM form when you have completed the EFM OAM diagnostic.

END OF STEPS

90.50 To configure an 802.3ah EFM OAM diagnostic test on an OmniSwitch Properties form

90.50.1 802.3ah EFM OAM diagnostic considerations

CAUTION

Service Disruption

Performing an 802.3ah EFM diagnostic is service-affecting.

Ensure that you consider the implications of performing this test before you proceed.

Link OAM (802.3ah) is not supported on mirroring ports.

When a port is in loopback mode, service mirroring does not work if the port is a mirror source or a mirror destination.

Note: For the 802.3ah EFM OAM diagnostic to be successful, the local and peer ports must support the 802.3ah protocol, and be operationally and administratively up.

The EFM OAM diagnostic is supported at the NE level or at the Ethernet port level on the Ethernet ports in network mode.

To configure Link OAM at the NE level, go to Step 1 .To configure Link OAM and 802.3ah EFM OAM diagnostics at the port level, go to Step 6 .

90.50.2 Steps

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Expand the OmniSwitch NEs icon.

3

Right-click on the OmniSwitch device where you want to configure an 802.3ah EFM OAM diagnostic and select Properties. The Network Element (Edit) form opens.

4

Click on the Globals tab.

5

Click on the EFM-OAM tab. The EFM-OAM form opens.

6

Configure the required parameters.

7

On the Network Element (Edit) form navigation tree, expand the Shelf icon.

8

Expand to the port level and click on a port object. The Physical Port (Edit) form opens.

9

Click on the Ethernet tab. The Ethernet port configuration form opens.

10

Click on the EFM-OAM tab. The EFM-OAM form opens.

11

Configure the required parameters.

12

Save your changes.

13

Click Yes.

The EFM-OAM form is updated with the results of the 802.3ah EFM diagnostic. The values that are displayed on the form depend on the configuration of the local and peer ports.

Table 90-6, “802.3ah EFM OAM results” (p. 3075), Table 90-7, “802.3ah EFM OAM Statistics” (p. 3077) and Table 90-8, “802.3ah EFM OAM peer information” (p. 3078) describe the displayed information.

Table 90-6 802.3ah EFM OAM results

Parameter	Value	Description
Operational Status (dot3OamOperStatus)	Disabled	The Administrative State parameter on the local port is set to Disabled.
	Active Send Local	The Mode parameter is set to Active and the local port is discovering the peer port EFM OAM capabilities.
	Passive Wait	The Mode parameter is set to Passive and the local port is waiting to receive OAMPDUs from the peer device.
	Send Local and Remote	The local port discovered the peer but has not yet accepted or rejected the configuration of the peer. The local port may determine that the peer port is not compatible and decline OAM peering.
	OAM Peering Locally Rejected	The local port declined OAM peering.
	Send Local and Remote Ok	The local port accepted OAM peering.
	OAM Peering Remotely Rejected	The remote port declined OAM peering.
	Operational	The local and remote ports accepted the OAM peering.
	Non-operational Half Duplex	The Administrative State parameter is set to Enabled but the local port is in half-duplex mode.
	Link Fault	The link detected a fault and is transmitting OAMPDUs with a link fault indication.
Max. PDU Size (dot3OamMaxOamPduSize)	PDU size	The largest OAMPDU size, in bytes, that is supported by the local port: Minimum is 64 Maximum is 1518 Default is 1514
Configuration Revision (dot3OamConfigRevision)	0 to 65 535	The configuration revision of the OAM entity obtained from the latest OAMPDU sent by the OAM entity. The value is used by OAM entities to indicate that configuration changes occurred that may require the peer OAM entity to re-evaluate whether OAM peering is allowed.

Table 90-6 802.3ah EFM OAM results (continued)

Parameter	Value	Description
Functions Supported (dot3OamFunctionsSupported)	Event Support	The port can send and receive event notification OAMPDUs.
	Loopback Support	The port can initiate and respond to loopback commands.
	Unidirectional Support	The port supports the transmission of OAMPDUs on links that operate in unidirectional mode.
	Variable Support	The port can send and receive variable request and response OAMPDUs.
Loopback Status (dot3OamLoopbackStatus)	No Loopback	The port is not in a loopback condition.
	Initiating Loopback	The local device initiated a loopback, sent a loopback OAMPDU and is waiting for a response from the peer port.
	Remote Loopback	The peer port is in loopback mode.
	Terminating Loopback	The local port is in the process of terminating a loopback on the peer port.
	Local Loopback	The peer port has put the local port into loopback mode.
	Unknown	An OAMPDU that contains an unexpected message was received by the local port.

Table 90-7 802.3ah EFM OAM Statistics

Statistics	Displayed Value	Description
Frames Lost Due to OAM(dot3OamFramesLostDueToOam)	*Number of frames*	The number of frames that were dropped by the OAM multiplexer. Discontinuities of this counter can occur under some conditions.
Information Rx (dot3OamInformationRx)	*Number of OAMPDUs*	The number of information OAMPDUs received on this interface
Information Tx (dot3OamInformation Tx)	*Number of OAMPDUs*	The number of information OAMPDUs transmitted on this interface. Discontinuities of this counter can occur can occur under some conditions.
Loopback Control Rx (dot3OamLoopbackControlRx)	*Number of OAMPDUs*	The number of loopback control OAMPDUs received on this interface
Loopback Control Tx (dot3OamLoopbackConrolTx)	*Number of OAMPDUs*	The number of loopback control OAMPDUs transmitted on this interface
Unsupported Codes Rx (dot3OamUnsupportedCodesRx)	*Number of OAMPDUs*	The number of OAMPDUs received on this interface with an unsupported opcode

Table 90-7 802.3ah EFM OAM Statistics (continued)

Statistics	Displayed Value	Description
Unsupported Codes Tx (dot3OamUnsupportedCodesTx)	*Number of OAMPDUs*	The number of OAMPDUs transmitted on this interface with an unsupported opcode

Table 90-8 802.3ah EFM OAM peer information

Peer Information	Displayed Value	Description
Peer MAC Address (dot3OamPeerMacAddress)	*MAC address*	The MAC address of the peer port. The MAC address is contained in the received OAMPDU.
Peer Vendor OUI (dot3OamPeerVendorOui)	*OUI*	The OUI of the OAM peer. The OUI is part of the peer MAC address contained in the received OAMPDU. The OUI can be used to identify the vendor of the remote OAM device.
Peer Vendor Info (dot3OamPeerVendorInfo)	*Vendor information text*	The vendor information field is in the local information TLV, and can be used to determine additional information about the peer device.
Peer Mode (dot3OamPeerMode)	Active or Passive	See Mode
Peer Max PDU Size (dot3OamPeerMaxOamPduSize)	*PDU size*	The largest OAMPDU value, in bytes, that is supported by the peer port. Minimum is 64 Maximum is 1518 Default is 1514
Peer Configuration Revision (dot3OamPeerConfigRevision)	0 to 65 535	The configuration revision of the OAM device as identified in the latest OAMPDU sent by the OAM peer. The configuration revision is used by OAM devices to indicate that configuration changes have occurred which might require the peer OAM device to re-evaluate whether OAM peering is allowed.
Peer Functions Supported (dot3OamPeerFunctSupported)	Event Support	The peer port can send and receive Event Notification OAMPDUs.
	Loopback Support	The peer port can initiate and respond to loopback commands.
	Unidirectional Support	The peer port supports the transmission of OAMPDUs on links that are operating in unidirectional mode.
	Variable Support	The peer port can send and receive Variable Request and Response OAMPDUs.

14

Close the EFM-OAM form when you have completed the EFM OAM diagnostic.

END OF STEPS

90.51 To configure an ICMP Ping template

90.51.1 Before you begin

ICMP ping templates are used to populate values for ICMP ping tests when the tests are used on L3 interfaces to control the operational state. ICMP ping templates are supported for IPv4 only.

Using the NFM-P policy framework, ICMP ping templates are assigned to L3 interfaces for IES and VPRN services; see the following procedures:

- 78.50 “To assign an ICMP ping template to an IES L3 access interface” (p. 2495)
- 79.105 “To assign an ICMP ping template to a VPRN L3 access interface” (p. 2679)

Nokia recommends that you assign an NE DoS protection policy configured with ICMP-Ping-Check as the protocol to the affected interfaces; see 78.45 “To assign a DoS protection policy or DDoS protection policy to an IES L3 access interface” (p. 2491) and 79.100 “To assign an NE DoS or DDoS protection policy to a VPRN L3 access interface” (p. 2675).

90.51.2 Steps

1

Click Tools→OAM Templates→ICMP Ping Template on the NFM-P main menu. The ICMP Ping Template form opens.

2

Click Create, or select an existing template and click Properties. The ICMP Ping Template (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy.

5

Perform 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

END OF STEPS

90.52 To configure a link measurement template

90.52.1 Purpose

The link measurement template includes common test parameters that can be synchronized with the Global policy. The link measurement template can be associated with a routing instance. When that association is established the interface will execute a process to determine the operational state and detect and defect conditions that may prevent proper test execution.

90.52.2 Steps

1

Click Tools→OAM Templates→Link Measurement Template on theNFM-P main menu. The Link Measuremen Templates form opens.

2

Click Create, or select an existing template and click Properties. The Link Measurement Template, Global Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Click OK to save the template and close the form, or click Apply to save the template.

5

Perform 27.2 “To configure a routing instance or a VRF instance” (p. 826) to associate the template to a routing instance.

END OF STEPS

90.53 To configure link monitoring on an Ethernet port

90.53.1 Steps

Note: Ethernet link monitoring works between directly-connected Ethernet ports and monitors link level errors, such as Ethernet frame errors, and for a subset of MDAs, Ethernet symbol errors as well. In the event that error levels exceed configured thresholds, a signal degradation error log, or in more severe cases, a signal failure error log can be generated. In the case of signal failure error, the peer port may be notified, and the local (and peer) ports may be disabled until the error log is manually cleared.

Satellite ports support virtually all OAM functionality. The only exception is that no EFM-OAM link monitoring is available for such ports.

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Right-click the device on which you want to configure link monitoring and select Properties. The Network Element (Edit) form opens.

3

Navigate to the port on which you want to configure link monitoring. The path is Network→*NE*→Shelf→Card Slot→Daughter Card Slot→Port *n/n/n*.

4

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

5

Click on the Ethernet tab and then on the EFM-OAM tab.

6

Configure the Administrative State to Enabled.

7

Configure the Transmit Interval (x100ms) and Multiplier (Intervals) parameters to determine the transmit rate for the link monitoring PDUs.

8

Configure the Dying Gasp Notify and Critical Event Notify parameters to determine the reporting of a signal failure to the far-end device.

9

Configure the parameters in the Link Event Configuration panel.

The general parameters in this panel are used to enable the link monitoring capability and to determine the port's response to local and remote signal degradation (SD) and signal failure (SF) events. The parameters in the Errored Frame, Errored Frame Period, Errored Frame Seconds Summary, and Errored Frame Symbols sub-panels are used to enable and configure their respective link monitoring variants. These monitoring variants can operate concurrently.

10

Expand the Counters panel to view link monitoring communications counters between the local and peer Ethernet ports. These counters are not automatically updated. Click Refresh to resynchronize the MIBs related to these counters.

11

Expand the Peer Information panel to view information learned from PDUs sent by the directly-connected peer. The Peer Vendor OUI attribute is used by routers to correctly interpret the meaning of attributes sent during dying gasp or critical events when the peer experiences a signal failure threshold crossing event.

12

Expand the Link Monitoring Error Event Logs panel to view signal failure and signal degradation event logs issued from local or remote link monitoring peer ports.

Perform the following:

1. Click Refresh to obtain current event logs from the NE.
2. Select a log from the list under either the Signal Failure or Signal Degradation sub-tabs and click Properties to view the contents.
3. Click Clear to clear local, remote, or all the event logs.

   Note:

   Clearing event logs cannot be undone.

13

Save your changes and close the form.

END OF STEPS

90.54 To configure system and port level ETH-OAM Dying Gasp notification

90.54.1 Reference

Ethernet First Mile OAM 802.3ah is deployed heavily at E-NNI and UNI connections. These connection points are typically multi-vendor and have resilient connectivity. Signaling the pending removal of a link when a soft reset could interrupt service on such ports is critical in reducing unacceptably long outages for known maintenance conditions. This signaling is accomplished by setting the Dying Gasp flag in the Information OAMPDU when a soft reset message is received by EFM for those ports that will be affected by the soft reset. The Dying Gasp flag can be enabled at both the port and system levels.

90.54.2 Steps

1

Choose Equipment from the navigation tree view selector. The navigation tree displays the Equipment view.

2

Right-click the device on which you want to configure dying gasp notification and select Properties. The Network Element (Edit) form opens.

3

Click on the Globals tab and then on the OAM sub-tab.

4

Enable the System Dying Gasp Tx on Reset Enable parameter. The default state is disabled.

Both the system level and port level dying gasp functionality must be enabled for the notification capability to be active on a port.

Note: You can click on Show system EFM-OAM Information to display additional EFM-OAM details.

5

Navigate to the port on which you want to configure the dying gasp notification. The path is Network→*NE*→Shelf→Card Slot→Daughter Card Slot→Port *n/n/n*.

6

Right-click on the Port icon and choose Properties. The Physical Port (Edit) form opens.

7

Click on the Ethernet tab and then on the EFM-OAM sub-tab.

8

Configure the Administrative State to Enabled.

9

Enable the Dying Gasp on Reset Enable and the Dying Gasp Notify parameters. The Dying Gasp Notify parameter determines whether or not the local OAM entity attempts to report a dying gasp to its peer OAM entity when a dying gasp event occurs.

10

Configure the Trigger Fault parameter as required. Enabling the Trigger Fault function allows an operator to invoke an ETH-OAM fault (either Dying Gasp or Critical Event) before shutting down the port.

11

Configure the remaining parameters as required.

12

Configure the required parameters in the Link Event Configuration panel. See 90.53 “To configure link monitoring on an Ethernet port” (p. 3081) for detailed information on configuring link monitoring.

Note: The Reaction to Peer Dying Gasp parameter specifies how the local port will be impacted when it receives an Information OAMPDU from a peer with the Dying Gasp flag set.

13

Expand the Link Monitoring Error Event Logs panel to view signal failure and signal degradation event logs issued from local or remote link monitoring peer ports. The event logs on a remote peer will record active error events when receiving the trigger fault command from the source node.

Perform the following:

1. Click Refresh to obtain current event logs from the NE.
2. Select a dying gasp log from the list under the Signal Failure sub-tab and click Properties to view the contents.
3. Click Clear to clear local, remote, or all the event logs.

   Note: Clearing event logs cannot be undone.

14

Save your changes and close the form.

END OF STEPS

90.55 To run a one-time validation test on a service

90.55.1 Steps

Perform the following steps to run the one-time validation test on a service:

1

Select the required test entity using the appropriate item under the NFM-P Manage menu item, and then click Properties. The entity's property form opens.

Entities that you can test using this function include:

- VLL, VPLS, and VPRN services (Manage→Service→Services)
- Composite services (Manage→Service→Composite Services)
- Dynamic LSPs (Manage→MPLS→Dynamic LSPs)
- Service tunnels (Manage→Service Tunnels)

2

Click One Time Validation.

Note: If a regular OAM validation is currently running on the selected test entity, the One Time Validation function will not be executed.

The Choose Validator Test Policy form is displayed and lists all test policies that are applicable to the selected test entity.

3

Select the required policy from the list and click OK.

Note: A default read-only test policy (Validation-Default-Policy) is supplied for this purpose. You can copy this policy but not modify it. Modify the cloned copy as required, or you can select a test policy you have previously created. See 89.10 “To configure an STM test policy” (p. 2947) to create an STM test policy.

The One Time Validation function proceeds after you click OK.

The automated operations that occur are as follows:

1. The One Time Validation test suite is created.
2. The selected test entity is added to the test suite.
3. Any tests, MEGs, and MEPs that are appropriate to the test entity and that are required to run the validation are created.
4. The generated tests are executed.
5. The test entity’s OAM Validation Failed state indicator is updated as required (a checkmark indicates the OAM validation failed whereas a blank checkbox indicates that the validation succeeded).
6. Test results for each individual test that was created and executed to perform the validation are copied to new test result objects.
7. The test suite, tests, and any required OAM objects that were created to run the validation are deleted.

4

To view the test results, click on the tested entity’s Test tab and then on the One Time Validation Result sub-tab.

Note: These archived test results are also available through the STM. They can be viewed by selecting the One Time Validation Result (Assurance) item from the object drop-down menu. See 89.19 “To view and compare STM test suite results for a tested entity” (p. 2963) for more information.

5

Select one of the test results in the list and click Properties. The One Time Validation Result (Edit) form opens.

6

Review the results as required and then close the One Time Validation Result (Edit) form.

Note: The archived test results are not deleted after viewing. They will remain available until manually deleted by the user.

7

Close the form.

END OF STEPS

91 Ethernet CFM

91.1 Overview

91.1.1 Purpose

This chapter describes the Ethernet CFM (Connectivity Fault Management) OAM diagnostic tests, and provides information about configuring and performing Ethernet CFM diagnostic tests.

91.1.2 Contents

Ethernet CFM

91.2 Ethernet CFM test descriptions

91.2.1 Introduction

This section provides a description of NFM-P supported Ethernet CFM OAM diagnostic tests. Unless noted, all tests are accessible from the STM. See Table 90-1, “NFM-P supported OAM diagnostic tests and configurations” (p. 2981) for a list of all supported OAM diagnostic tests and their applicable procedures.

91.2.2 Global MEG check

A global MEG check detects connectivity failures between pairs of local and remote maintenance endpoints, or MEPs, in a MEG. A MEG consists of the maintenance entities that belong to the same service inside a common service OAM domain.

See 91.19 “To create and run a Global MEG OAM diagnostic test from the STM” (p. 3120) for information about how to create and run a global MEG check from the STM.

91.2.3 Continuity check

Continuity check messages are multicast messages that a MEP transmits periodically to remote MEPs in the same MEG. CCMs are used to discover a remote endpoint, check the health of a site, and detect cross-connect misconfigurations. A CFM continuity check test is automatically generated when you create an MD. When you execute the test and a connectivity fault is present, the MEP that detects the fault raises an alarm.

See 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122) for information about how to create and run a continuity check from the STM.

The loss of three consecutive CCMs or the receipt of a CCM with incorrect information, indicates a fault.

Note: If a service is modified after you associate it with an MD; for example, a new site is added, you must manually add new MEPs.

If a CFM continuity check test is running on a service when you add a new MEP to the service, you must stop the test and execute it again to make the new MEP active.

When a service is modified after a CFM continuity check is created; for example, a new B-VPLS site is added to the service, you must manually add a virtual MEP to the site.

91.2.4 CFM dual-ended loss test

A CFM dual-ended loss test functions as an optional extension of a CFM continuity check. It applies only to Y.1731 MEPs that are configured on the 7705 SAR. This type of test is used to calculate the rate of frame loss in each direction for Ethernet packets sent between two MEPs.

When a continuity check test is executed with the dual-ended loss option enabled, the option is replicated on all participating MEPs that support the test, along with the accompanying alarm threshold values. If a MEP detects that the local or remote frame loss ratio has exceeded the alarm threshold for a remote MEP, the MEP raises an alarm against the remote MEP.

See 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122) for information about how to configure the CFM dual-ended loss test option as part of creating and running a CFM continuity check from the STM.

91.2.5 CFM loopback

CFM loopback messages are sent to a unicast destination MAC address. The MEP at the destination responds to the loopback message with a loopback reply. A MEP or a MIP can reply to a loopback message if the destination MAC address matches the MAC address of the MEP or MIP. CFM loopback tests verify connectivity to a specific MEP or MIP.

See 91.21 “To create and run a CFM loopback OAM diagnostic test from the STM” (p. 3123) for information about how to create and run a CFM loopback from the STM.

You can also perform multicast loopbacks by providing a multicast address (class 1 multicast destination) that aligns with the level that the originating MEP is configured on. Only one multicast test can be run at a time per NE, and results from the previous test are deleted when a new test is started. The stored values include the responding MEP MAC address, the sequence number, and a locally-assigned rx index (allowing you to detect out-of-order responses).

91.2.6 CFM link trace

CFM link trace messages that contain a target unicast MAC address are sent to multicast destination MAC addresses. Each MIP at the same MD level replies with a link trace response. Messages are forwarded to the next hop until they reach the destination MAC address.

See 91.22 “To create and run a CFM link trace OAM diagnostic test from the STM” (p. 3125) for information about how to create and run a CFM link trace from the STM.

91.2.7 CFM Eth test

The CFM Eth test applies only to Y.1731 MEPs. This one-way test originates on a source MEP and terminates on a destination MEP. The target of a CFM Eth test is a MAC address. The test is used to perform one-way in-service diagnostic test that include verifying bandwidth throughput, frame loss, and bit errors. To perform the test, a MEP inserts frames with Eth-test information that includes specific throughput, frame size, and transmission patterns. A MIP is transparent to Eth-test frames.

See 91.23 “To create and run a CFM Eth OAM diagnostic test from the STM” (p. 3127) for information about how to create and run a CFM Eth test from the STM.

91.2.8 CFM two way delay test

The CFM two way delay test applies only to Y.1731 MEPs. In this test, the frame delay is defined as the time elapsed since the start of transmission of the first bit of the frame by the source site until the frame is received by the same site after passing through the destination site. The frame delay represents the round-trip time between the source and destination sites.

See 91.24 “To create and run a CFM two way delay OAM diagnostic test from the STM” (p. 3128) for information about how to create and run a CFM two way delay test from the STM.

91.2.9 CFM one way delay test

The CFM one way delay test applies only to Y.1731 MEPs. The test originates on one MEP and terminates on a target MEP. The results are read from the target MEP. In the test, frame delay is defined as the time elapsed since the start of transmission of the first bit of the frame by a source site until the frame is received by the destination site. The frame delay represents the one-way trip time between the source and destination sites.

See 91.25 “To create and run a CFM one-way delay OAM diagnostic test from the STM” (p. 3130) for information about how to create and run a CFM one way delay test from the STM.

91.2.10 CFM single ended loss test (7705 SAR only)

The CFM single ended loss test applies only to Y.1731 MEPs. This one-way test originates on a source MEP and terminates on a destination MEP. The target of a single-ended loss test is a destination MAC address. The test is used to calculate the rate of frame loss in each direction for Ethernet packets sent between the two MEPs.

See 91.26 “To create and run a CFM single-ended loss OAM diagnostic test from the STM” (p. 3131) for information about how to create and run a CFM single ended loss test from the STM.

91.2.11 CFM two way SLM

The CFM two way SLM test provides Synthetic Loss Measurement and is used to check packet loss for a particular MEP. See 91.27 “To create and run a CFM two way SLM OAM diagnostic test from the STM” (p. 3132) for information about how to create and run a CFM two way SLM test from the STM.

Note: On the 7705 SAR, SLM is only supported on Epipe services under the following conditions:

- single-ended (two way) loss measurement using optional TLV with a timestamp on the near-end and far-end for combined loss and delay measurement
- unicast destination addresses only
- DOWN MEP direction for Spoke-SDPs only (UP and DOWN MEP direction for SAPs)

91.2.12 CFM LM test

The CFM LM (loss measurement) test measures the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs.

See 91.28 “To create and run a CFM LM OAM diagnostic test from the STM” (p. 3134) for information about how to create and run a CFM LM test from the STM.

91.2.13 CFM Statistics

The Ethernet CFM statistics are used for viewing and analyzing the current overall processing requirements for CFM. Any packet that is counted against the CFM resource will be included in the statistics counters.

The Ethernet CFM statistics are available per NE and per MEP, with a per-OpCode breakdown. Use the **`show eth-cfm statistics`** command to display the statistics at the system level. Use the

show eth-cfm mep mep-id domain md-index association ma-index statistics command to view the per-MEP statistics.

From the NFM-P GUI, Ethernet CFM statistics can be retrieved:

- at the network level
  - Network Element (Edit) →Statistics tab
- at the MEP level
  - Tool →Service Test Manager (STM) →CFM opcodes tab

For ETH-CFM statistics at the NE level, the record types include:

- CFM Global Opcode Stats (Ethernet OAM)
- Ethernet CFM Packet Count (Ethernet OAM)
- Ethernet CFM Resource Limits (Ethernet OAM)
- OAM General Stats (Ethernet OAM)
- OAM Performance Req Types Stats (Ethernet OAM)

At the individual MEP level, CFM Opcodes → OAM_MEP_statistics can be displayed. Click Resync to see the changes in the NFM-P.

These statistics help operators to determine the busiest active MEPs on the system with a breakdown of per-OpCode processing at the system and MEP level. For eth-cfm oper up statistics collection, create an MD and a MEG and attach a service with up/down MEP. See 91.14 “To configure an automatic MEP ID assignment on an NE” (p. 3108) and 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109).

91.3 Ethernet CFM concepts and components

91.3.1 Ethernet CFM function

The NFM-P Ethernet Connectivity Fault Management, or Ethernet CFM, function is implemented based on the IEEE 802.1ag OAM standard. This standard describes protocols for detecting, isolating, and reporting connectivity faults in an Ethernet network.

You can use Ethernet CFM for the following:

- path discovery
- fault detection
- fault isolation
- fault notification

Ethernet CFM diagnostic tests are configured using the NFM-P Service Test Manager. See Chapter 89, “Service Test Manager” for more information about the STM. See Chapter 90, “OAM diagnostic tests” for more information about CFM diagnostic tests.

91.3.2 MDs

The IEEE 802.1ag OAM standard partitions a network into eight hierarchical levels called maintenance domains, or MDs. An MD is a network, or part of a network, that is provisioned with a set of maintenance entity groups, or MEGs, which are groups of service sites.

91.4 MEGs

91.4.1 MEG description

Typically, a MEG represents one service and consists of a group of maintenance end points, or MEPs. Only one MEG can be associated with a service, but one service can be associated with multiple MEGs. MDs and MEGs are distributed to NEs using NFM-P policy distribution.

Within a MEG, MEPs can be sorted into logical groupings called MEG subgroups. A MEG that contains subgroups is called a Global MEG. By default, a Global MEG has one subgroup. MEG subgroups allow you to group managed MEPs and unmanaged remote MEPs so that Ethernet CFM can be directed to a specific area of a system, a MEG subgroup, rather than the entire system.

91.5 MEPs

91.5.1 MEP description

The MEPs are configured at the edge of an MD and perform the following functions:

- periodically send CFM continuity check messages
- validate CFM PDU replies
- discard CFM PDU messages that are not in the MEP configuration
- initiate and respond to CFM test messages

MEPs can be added to services automatically or manually. When a MEP is assigned to a service manually, it associates its MEG with a single site on the service. When a MEP is assigned to a service automatically, it associates its MEG with all sites on the service. MEPs which are generated automatically can inherit test generation options from the SAP or service site. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) for more information.

You can configure an initial MEP ID for automatic MEP ID assignment on an NE.

You can also add MEPs to services during Ethernet CFM test configuration from the Service Topology map. See 4.2.3 “Working with Ethernet CFM objects” (p. 177) in Chapter 4, “Topology map management” for more information.

Each MEP is assigned an up or down direction. An up MEP is provisioned on an ingress port, and monitors the forwarding path inside a bridge NE to the egress port. A down MEP is provisioned on an egress port, and monitors the forwarding path between bridge NEs.

You can assign roles to managed MEPs and unmanaged remote MEPs during OAM test suite configuration. A MEP can be designated as a hub or spoke, and as a test source, a test target, or both. Assigning roles to MEPs in test suites can reduce the total number of automatic tests generated. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109)
.

MEPs can inherit their roles from the SAP or service site. You can propagate test generation role settings to all unmanaged remote and managed MEPs on a SAP or service site using the Propagate to MEPs button on the ETH-CFM tab of the Access Interface (Edit) and Site (Edit) forms.

91.5.2 MEP association

An up MEP can be associated with the following object types:

- Epipe and VPLS SAPs
- Epipe spoke SDP bindings
- VPLS mesh and spoke SDP bindings
- VPLS and MVPLS B-sites, as virtual MEPs
- B-L2 access interfaces
- OmniSwitch VLAN SAPs

A down MEP can be associated with the following object types:

- Epipe, Ipipe, IES, VPLS, and VPRN SAPs
- Epipe and IES spoke SDP bindings
- VPLS mesh and spoke SDP bindings
- B-L2 access interfaces
- MVPLS access interfaces
- OmniSwitch VLAN SAPs
- Network interfaces
- Ports
- LAGs

91.5.3 Virtual MEPs

A virtual MEP is an up MEP that is created on a VPLS site when a CFM continuity check test is run. Each virtual MEP transmits a CFM continuity check stream on all SAPs and SDPs of the site. A virtual MEP uses the site MAC address, if configured; otherwise, it uses the shelf MAC address. See Chapter 77, “VPLS management ” for information about assigning virtual MEPs to a VPLS.

The following rules apply to virtual MEP management:

- One virtual MEP can be configured on a VPLS or MVPLS B-site.
- All regular MEPs on SAPs and SDP bindings in the same MEG as a virtual MEP must be configured as up MEPs.
- Regular MEPs in the same MEG and on the same B-site as a virtual MEP cannot be enabled when the virtual MEP is enabled.
- Virtual MEPs can be created in a MEG only when MIP creation in the MEG is disabled.

91.5.4 Facility MEP

A facility MEP is a MEP that is down and is created on a router interface, network interface, port, or LAG. A facility MEP detects failure conditions for an Ethernet transport network using ETH-CCM or AIS and, where appropriate, propagates alarm conditions so that the services that share this common transport are aware of the failure.

A virtual tunnel facility MEP is created by configuring a LAG MEP or a port MEP with a VLAN ID. In this instance the VLAN ID must match the outer encapsulation value of the SAP that is associated with the Tunnel MEP.

The following rules apply to facility MEP management:

- Only one facility MEP can be configured on a port, a LAG, or a network interface.
- A facility MEP can be configured on a port only if the MD level is 0.
- A facility MEP must be configured in the down direction.
- A facility MEP ID must be unique within the same device.
- Port facility MEPs are supported on ports in access, hybrid, or network mode with dot1q encapsulation, and on ports in access or hybrid mode with Q in Q encapsulation.
- Port facility MEPs cannot be configured on network elements that support Ethernet tunnels.
- Tunnel facility MEPs are supported on LAG ports configured with Q in Q encapsulation in access or hybrid mode.
- Tunnel facility MEPs must have a lower MD level than any service MEPs that are created on the same LAG.
- A tunnel facility MEP VLAN ID must be unique within the same port.
- Tunnel facility MEPs cannot be configured on LAGs with a VLAN ID of 0.
- Tunnel port facility MEPs cannot be configured on LAG members.
- Fast facility MEPs, which are MEPs with a CCM interval of 10 ms or 100 ms, cannot be configured when CCM padding is enabled.
- There must be 2 facility MEPs in a MEG subgroup to execute a CCM test.

91.5.5 Unicast MEPs

The following rules apply to Unicast MEP management:

- When a CCM test is executed, the NFM-P will deploy remote MEPs (with unicast-da) to spokes for NEs that support unicast CCM.
- MEG sub-groups will ignore Unicast remote MEP lists when discovering MEG sub-groups and only use the hub remote MEPs list.
- Remote MEPs configured as unicast will update the appropriate Unicast Hub MEP’s Operational MAC Address. If an appropriate Unicast Hub MEP cannot be found, the remote MEPs will auto-create an unmanaged remote MEP with the Unicast Hub’s target pointer set.
- An unmanaged remote MEP without a MAC Address configured cannot be used as a Unicast Hub.

- Hub MEPs should be configured with remote lists that target all regular (non-unicast) MEPs and only the unicast MEPs that point to the hub. This allows a mixture of unicast and regular MEPs for CCM testing.
- You can use a multi-selection of MEPs to set a Unicast Hub MEP pointer.
- A MEP cannot point to itself as a Hub.
- A MEP cannot point to a Hub on the same site.
- If a MEG site has more than one MEP, the NFM-P sets all the MEPs to have the same Hub pointer.
- The remote MEP list of a Hub will not include spokes that do no point to it.
- When listing available Hubs, the NFM-P will also list unmanaged remote MEPs that have a MAC address from the same sub-group.
- When the NFM-P discovers a remote MEP it manages that has an appropriate MAC address, it will update the hub pointer.
- When the managed MEPs point to a hub, then regardless of what other remote MEPs or unmanaged remote MEPs are deployed, NFM-P will only deploy the hub's MEP ID with the Operational MAC Address. No other MEP IDs will be deployed.
- No local MEP is allowed when there is a remote MEP with a MAC address.
- No remote MAC address is allowed if multiple remote MEPs exist.
- If a MEP with a hub pointer exists on a MEG site, NFM-P will auto-populate the hub pointer when a new MEP is create
- When you create a remote MEP with a MAC address, the NFM-P creates an unmanaged remote MEP. This unmanaged remote MEP will be in the same sub-group as the spoke.
- Where MC-LAG redundant MEPs exist, the hub will point to the active one.

91.6 MIPs

91.6.1 MIP description

Maintenance intermediate points, or MIPs, are internal points in an MD that perform the following functions:

- validate received CFM PDUs
- validate and respond to link trace messages
- validate and respond to loopback messages

A MIP consists of two half-function objects that allow the MIP to be recognized as a MIP in one MD level and as a MEP on a higher level.

91.7 Allocating bandwidth resources for Ethernet CFM LBM/LBR SAT

91.7.1 Description

Service Activation Testing (SAT) is a critical operational task that is performed prior to service handoff to a subscriber. SAT can be performed up to speeds of 10 Gbps. To support SAT on the NFM-P, you can configure the following GUI objects to facilitate high-speed Ethernet-CFM LBMs (Loopback messages)-to-LBRs (Loopback responses) processing typically used during SAT.

- Use the Service Activation Response parameter on the MEP Properties form to specify if the MEP uses standard bandwidth processing or if additional bandwidth processing is applied for processing service activation requests/stream encapsulated in the ETH-LBM (Ethernet-Loopback message) format. When enabled, additional bandwidth resources are allocated to rapidly respond to this stream of LBM messages; the MEP will not validate any type-length-value (TLV)s, and will not increment or compute any loopback statistics. When disabled, the MEP ETH-LBMs uses standard bandwidth processing. You can configure the Service Activation Response parameter on SAPs, Spoke or Mesh bindings, and Network Interfaces MEPs associated with VLL Epipe and VPLS services.
- Use the LBM Service Activation Responder Info button on the NE Properties form (Globals tab; OAM sub-tab), to execute the “show eth-cfm lbm-svc-act-responder” CLI command. This command displays all the MEPs that have the Service Activation Response parameter enabled.
- The LBM responder is supported for both the service-based MEPs and the MEPs on the router interface.

91.8 Configuring ITU-T Y.1731 ETH-ED grace period notifications

91.8.1 ETH-ED grace period notification description

ITU-T Y.1731 performance monitoring provides standards-based Ethernet performance monitoring that encompasses the measurement of Ethernet frame delay, frame delay variation, and frame loss and throughput as outlined in the ITU-T Y-1731 specification and interpreted by the Metro Ethernet Forum (MEF).

The ETH-ED (Ethernet Expected Defect) function of the ITU-T Y.1731 standard is used by a MEP to signal to its peer MEPs, the transmission of CCM frames is expected to be interrupted, without any interruption to data frames. The consequent loss of continuity defects at the peer MEPs is also suppressed. Frames with ETH-ED information carry the MEP ID of the MEP and the expected duration of the interruption.

You can enable ETH-ED grace period notification for MEPs using the Grace Period TX Enabled parameter located on the NE Properties form (Globals tab, OAM sub-tab); see 12.5 “To modify NE properties” (p. 343).

When grace period notification is enabled, you can perform specific ITU-T Y.1731 ETH-ED protocol configurations using the following parameters located on the supported object Grace panel: Eth VSM Rx Enable, Eth VSM Tx Enable, Rx Enable, Tx Enable, Max Rx Defect Window, and Priority. See the XML API Reference for information about configurable parameters and their applicability. See Table 91-1, “Supported objects and configuration guideline for Grace ETH-ED period notifications” (p. 3097) for supported object types for grace period notifications and configuration guidelines.

Table 91-1 Supported objects and configuration guideline for Grace ETH-ED period notifications

Supported network object	Configuration guidelines
Ethernet Tunnel	See ...33.14 “To configure an Ethernet tunnel endpoint” (p. 1203) Navigate to: Service Tunnels→Ethernet Tunnel→Create→Tunnel Properties→Tunnel End points→Ethernet Tunnel Path Endpoint→MEPs→General tab (Grace panel)
Ethernet Ring Element	See ...33.16 “To configure an Ethernet Ring Element” (p. 1208) Navigate to: Ring Element→MEP tab→Create ETH-CFM Global→MEPs→General tab (Grace panel)
Routing Instance Interface	See ...27.2 “To configure a routing instance or a VRF instance” (p. 826) Navigate to: NE Routing Instance→Router Interface Properties→MEPs→General tab (Grace panel) Note -- Ensure the port is associated with the Router Interface
SAP MEPs associated with EPIPE, IPIPE, VPLS, IES, and VPRN services	See ...74.25 “To configure a MEP on a SAP” (p. 2048) Navigate to: Service type→Sites→L2 (or L3) Access Interface→SAP Properties→OAM tab→ETH-CFM tab→Create MEP→MEPs (Grace panel)
Subscriber Interface MEPs associated with IES, and VPRN services	See ...78.16 “To configure a subscriber interface on an IES” (p. 2444) or 79.47 “To configure a subscriber interface on a VPRN” (p. 2610) Navigate to: Service type→Sites→Subscriber Interfaces→Group Interfaces→Service Access Points Properties→OAM tab→ETH-CFM tab→Create MEP→MEPs (Grace panel)
Spoke SDP Binding or Mesh SDP MEP associated with EPIPE, VPLS, IES, and VPRN services	See ...77.99 “To create a MEP on a VPLS SDP binding” (p. 2398) Navigate to: Service type→Sites→Spoke SDP Bindings Properties→OAM tab→ETH-CFM tab→Create MEP→MEPs (Grace panel)
LAG MEP	See ...13.16 “To create a LAG” (p. 431) or 13.17 “To modify a LAG” (p. 435) Navigate to: NE→Logical Groups→LAGS→Properties→MEPs→Create MEP→MEPs (Grace panel)
Port/Facility MEPs	See ...91.5.4 “Facility MEP” (p. 3095) Navigate to: Port Properties→Ethernet tab→MEPs (Grace panel)

91.8.2 ETH-ED enforcement rules

The NFM-P applies the following enforcement rules for ETH-ED grace period notifications:

- When a soft-reset occurs on the node, if ETH-ED transmit and receive are both enabled, then Grace ETH-ED takes precedence over Grace ETH-VSM. CCM should be enabled when a soft-reset occurs. Only when is CCM is enabled will the grace protocols be applied.
- Transmitting Grace field will read the protocol ETH-ED for the maximum Rx defect window value.
- Any MEP transmitting the Grace ETH-ED defect protocol must have CCM enabled.
- When a soft-reset occurs, no configurations can be added or changed either on the node or using the NFM-P.
- When a soft-reset occurs on the node, if ETH-VSM transmit and receive are both enabled, then Grace ETH-VSM takes precedence over Grace ETH-ED.

91.9 Ethernet-CFM redundancy

91.9.1 Ethernet-CFM redundancy description

When you configure Ethernet-CFM redundancy, you can link MEPs to the state of the resiliency mechanism that is supported in MC-LAGs. The state of the MEP does not affect the state of resiliency mechanism.

Ethernet-CFM redundancy is configured at the NE level. When you configure Ethernet-CFM redundancy, the state of a MEP is the same as the state of the SAP, LAG, or MC-LAG with which it is associated. For example, if the redundant MC-LAG is in standby state, the MEP is also in standby state. When a MEP is in standby state, the MEP is idle, CCMs are not exchanged, and the MEP does not respond to CFM tests.

Note: The MEP, SAP, and LAG must reside within an MC-LAG that has Ethernet-CFM redundancy configured. See 43.5 “To create an MC LAG group” (p. 1365) to configure MC-LAGs.

When Ethernet-CFM redundancy is configured and the active state of an MC-LAG changes, the MC-LAG Inactive state on the MEP changes and any defects on the MEP are flagged.

The state of MEPs, SAPs, and LAGs that are configured for Ethernet-CFM redundancy appears in the MC-LAG Inactive field. The status of MC-LAGs that are configured for Ethernet-CFM redundancy appears in the Is Active field.

In addition to configuring Ethernet-CFM redundancy at the NE level, you can configure Ethernet-CFM redundancy on tunnel MEPs. See 12.5 “To modify NE properties” (p. 343) to configure Ethernet-CFM redundancy at the NE level. See 13.17 “To modify a LAG” (p. 435) to configure Ethernet-CFM redundancy on a tunnel MEP.

91.10 Sample Ethernet CFM implementation

91.10.1 MDs, MEPs, and MIPs in an IEEE 802.1ag network

Figure 91-1, “Ethernet CFM objects in example network” (p. 3100) shows an example of MDs, MEPs, and MIPs in an IEEE 802.1ag network that consists of two operator areas, for example, services in the provider network, that are joined to create a network path for customer traffic.

The number on an object identifies the associated MD, which is one of the following:

- MD 5—end-to-end customer path
- MD 3—end-to-end network provider path
- MD 2—paths within services
- MD 0—physical path

MD 5 provides access to a down MEP on each CE device, and to a MIP on each PE bridge. MD 3 provides access to an up MEP on each PE bridge, and to a MIP on each bridge between groups of operator bridges. MD 2 provides access to the up MEPs and MIPs in each service. MD 0 provides access to MEPs for checking the physical connectivity between NEs.

Figure 91-1 Ethernet CFM objects in example network

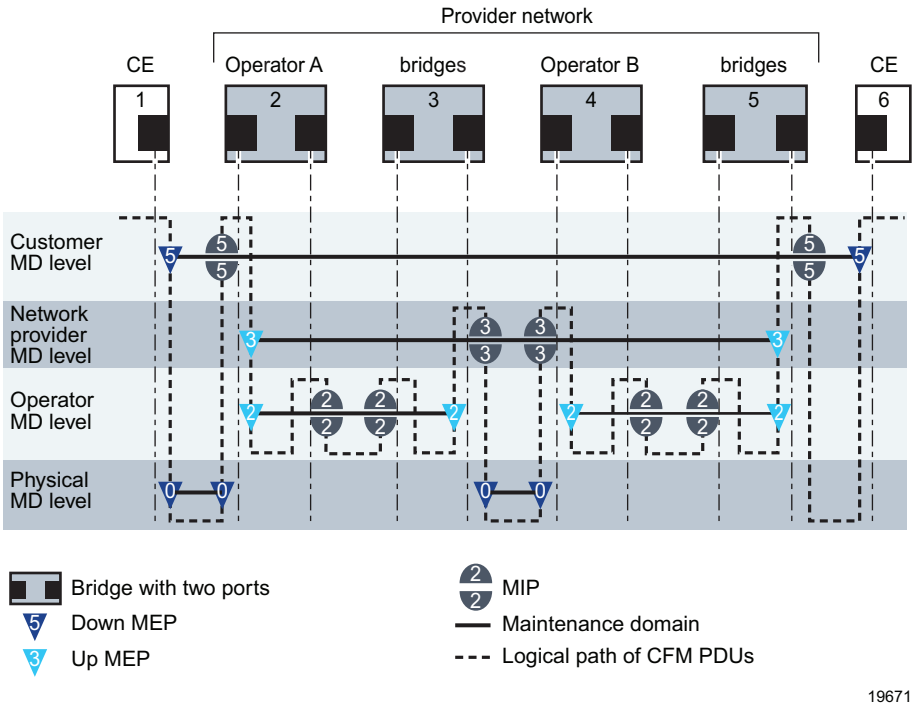

91.11 Ethernet CFM implementation for composite services

91.11.1 Composite service description

End-to-end test suite generation for composite services is available for VPLS and VLL Epipe services. In this context, a composite service is a set of services connected by one or more of the following:

- VLAN uplinks
- Spoke bindings
- CCAG connectors

Creating test suites and test policies for composite services is very similar to creating these for a regular service. Test generation options can be applied to modify the roles of managed and remote unmanaged up MEPs. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) . The Accounting Files and Continuously Executed options are also available when you create a test policy for a composite service.

Once tests have been ordered to generate, a MEG is created on each eligible endpoint SAP. Tests are then generated for this MEG that target other eligible endpoint SAPs within the composite service. Results for individual tests can then be reviewed.

Note: If there is a previously created MEG on one of the service sites (for example, a MEG created for a service before the service was included in a composite service) and the Generate Tests button is selected, tests are not generated. When a MEG with only a subset of the composite service already exists on one of the service SAPs, the generation of tests is blocked. The MEG must be deleted before the CFM tests are generated.

91.11.2 Examples of Ethernet CFM tests for composite services

The following are examples of how the Ethernet CFM test generation rules for composite services are structured to operate.The examples are based on the default test roles for MEPs. The default roles are Hub and both Test Source and Test Target.

The following figure shows a flat topology view of a simple configuration of two VPLSs (VPLS 8 and VPLS 15101) that are connected by a pair of spoke bindings between sites 192.0.2.96 and 192.0.2.65.

Figure 91-2 Composite service connected by spoke bindings

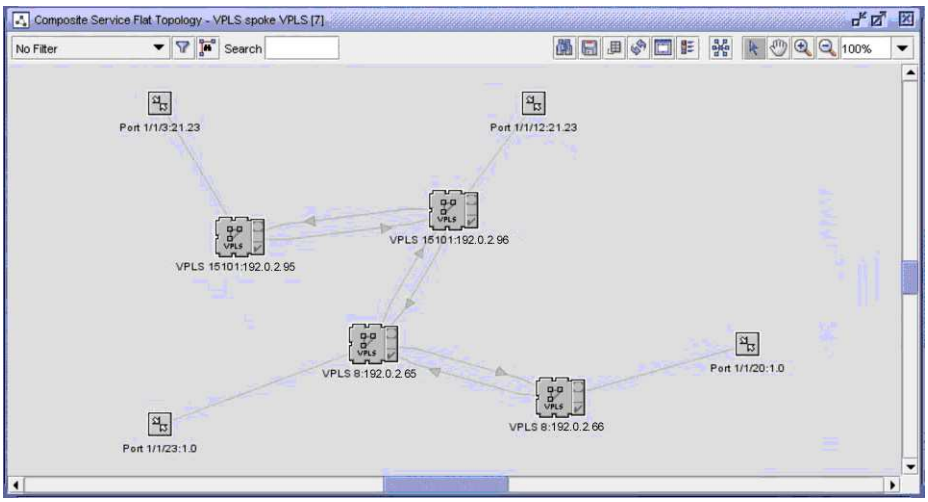

In this configuration, all four SAPs are considered as endpoints. Therefore MEPs are generated on each SAP and they all target each other. A total of twelve tests are generated. Double-clicking on an endpoint opens the interface configuration form and the MEP generated for this particular endpoint is displayed under the MEPs tab. If you then query the MEP's properties, the Tests tab in the MEP's configuration form displays the three tests that target the other MEPs in this composite service.

The following figure shows an example of a routed VPLS configuration with a VPLS (VPLS 15104) and a VPRN (VPRN 11). The services are connected by a routed VPLS interface (shown as routed:N/A:0.0).

Figure 91-3 Routed VPLS composite service

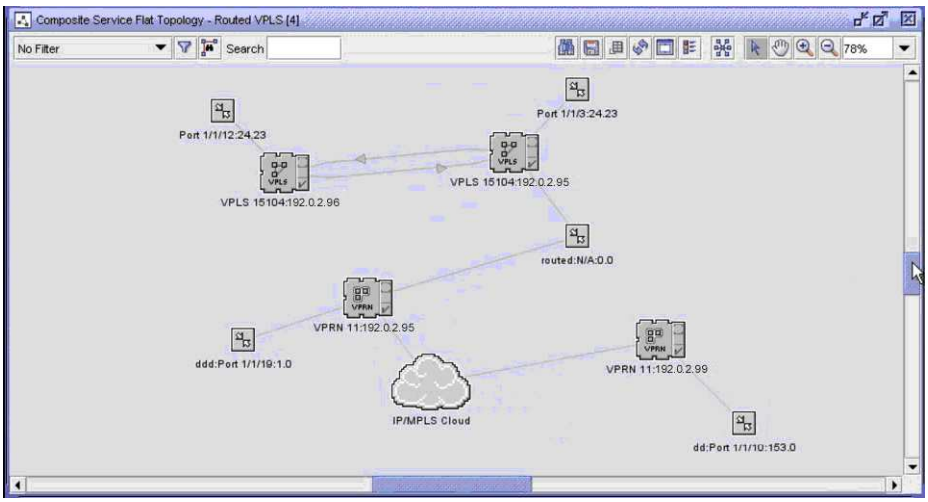

In this configuration, site 192.0.2.95 exists on both the VPLS and VPRN service. This is an overlapping site. The test generation rules would normally skip an overlapping site and no MEGs or tests would be generated for that site.

However in this case, the overlap for this site occurs in an unsupported service. The test generation rules for such a configuration will skip the unsupported VPRN service entirely and only generate tests for the supported VPLS, including site 192.0.2.95. Therefore, two MEPs and two tests (targeting each other) will be generated here, one for each of the SAPs on sites 192.0.2.96 and 192.0.2.95.

The test generation logs available in the test suite configuration form (under the Generation Error Logs tab) provide an explanation whenever the test generation rules cause items to be excluded from the tests. In this case, the log would state: “The CFM Composite Services Test Suite skipped the unsupported service”.

The following figure shows an example of a cross-connected VPLS composite service.

Figure 91-4 Cross-connected VPLS composite service

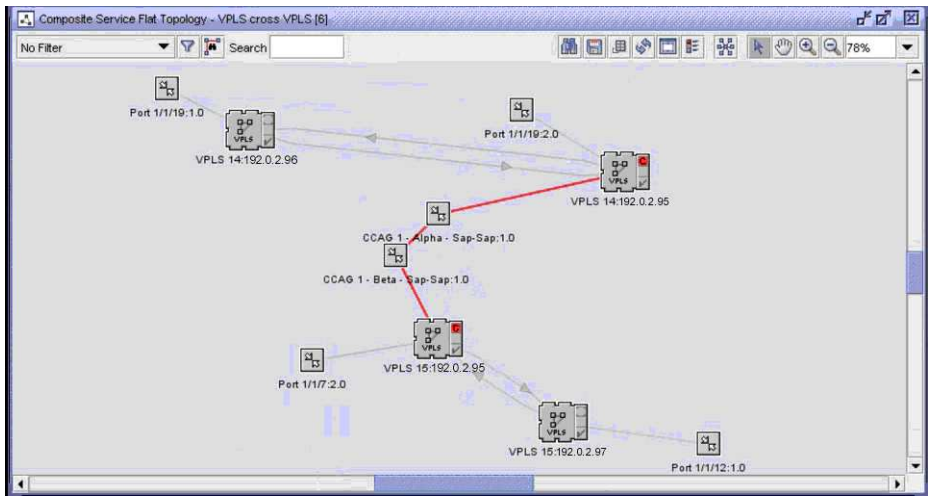

In this configuration, the two VPLSs (VPLS 14 and VPLS 15) are cross-connected by the SAPs CCAG 1-Alpha and CCAG 1-Beta. Both these SAPs originate from site 192.0.2.95, which exists in both services. It is therefore an overlapping site. The test generation rules skip overlapping sites, so no MEGs or tests are generated for any SAPs associated with that site.

Therefore, two MEPs and two tests (targeting each other) will be generated here, one for each of the SAPs on sites 192.0.2.96 and 192.0.2.97.

In addition, there will be two test generation logs available in the Generation Error Logs tab of the test suite configuration form. In this case, the logs would be identical, one stemming from each of the supported VPLSs. The logs would state: “Skipped overlapping service site”.

The following figure shows an example of an SCP-connected VPLS composite service.

Figure 91-5 SCP-connected VPLS composite service

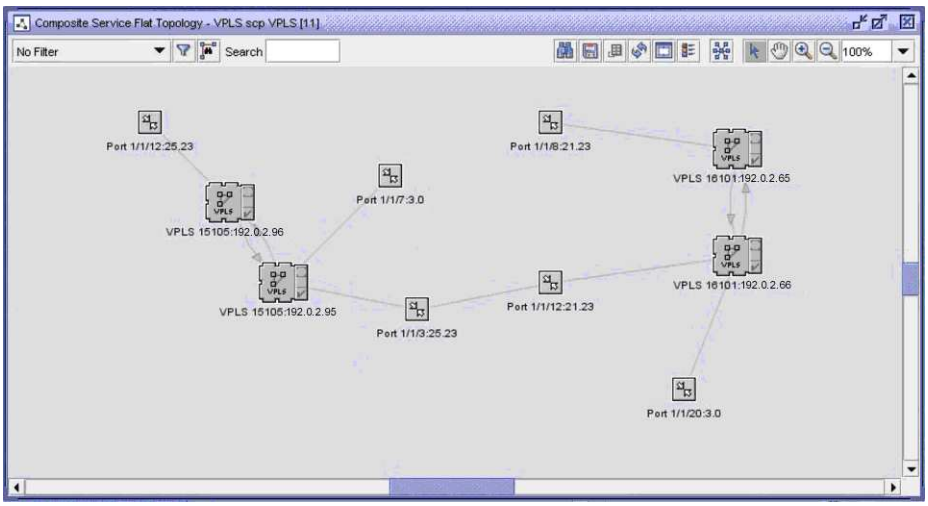

In this configuration, the two VPLSs (VPLS 15105 and VPLS 16101) are joined by an SCP connector between sites 192.0.2.95 and 192.0.2.66. The test generation rules skip any SAPs identified as part of a connector, since they are not endpoints. So no MEGs or tests are generated for either connector SAP.

MEGs and tests will however, be generated for each of the other four SAPs shown in the example. In total, four MEGs and twelve tests will be generated.

In addition, there will be two test generation logs available in the Generation Error Logs tab of the test suite configuration form. In this case again, the logs would be identical, one stemming from each of the supported VPLSs. The logs would state: “Skipped service connecting SAPs”.

91.12 Primary VLAN support for Ethernet CFM

91.12.1 Primary VLAN feature description

A VLAN assigned for service monitoring is considered the Primary VLAN. By setting a Primary VLAN identifier, you can perform additional ETH-CFM extraction checking.

By default, the ETH-CFM processing uses a SAP or SDP binding configuration to determine the proper offset for the start of the ETH-CFM e-type 0x8902 PDU. Any packet that includes ETH-CFM (e-type 0x8902) at the appropriate offset will be subjected to the extraction routines for MEPs or MIPs provided the Primary VLAN is enabled. If the ETH-CFM is not at the appropriate offset, the packets pass as user data. When a MEP is configured without a Primary VLAN, only the classic extraction is performed.

The default mapping of the ingress packet is based on the Layer 2 mapping which in turn, is mapped to the appropriate MEP. If the mapping does not match, as in the case of ambiguous or additional VLAN tags beyond the SAP or SDP binding delineation, then additional lookups for ETH-CFM are possible.

This is applicable where a single VLAN or aggregate is transporting many VLANs underneath that aggregate function. It is also applicable to dot1q networks that carry multiple tags without using the default SAP. Fully-qualified SAPs/SDP binding are also able to use the Primary VLAN function. The extraction function is the same for either fully-qualified or ambiguous SAPs/SDP binding.

A MEP or MIP must be configured on a particular SAP or SDP binding to perform any ETH-CFM extraction. The classic extraction (which is the matching of the ETH-CFM packets immediately following the SAP/SDP binding configuration) is performed first. If the Ethernet encapsulation is set to dot1q and the SAP configuration is x/y/z:*, then the ETH-CFM extraction looks for ETH-CFM to appear in the header with no VLAN tags preceding it. If there is a match on this classic extraction, then no further extraction routines are invoked. If the classic extraction fails and a second MEP is configured with Primary VLAN enabled, then the additional extraction will check to see if ETH-CFM immediately follows the specified Primary VLAN. If that is true, then the VLAN will be used as an index into the primary VLAN table and the extraction is performed based on normal criteria following the ETH-CFM rules.

The Primary VLAN function is supported for up and down MEPs, as well as ingress and egress MIPs on an Ethernet SAP for VPLS and Epipe service MEPs (including Local Switched PBB ePipe with PBB Tunnel backup). Primary VLAN support is only applicable to service MEPs and not facility MEPs (specifically, not tunnel MEPs). Facility MEPs are meant to validate the transport and not the switching capability of a network element.

Primary VLAN support for 7210 SAS NEs varies depending on chassis type, release version, and card type. See the NE documentation for more information.

When enabling this function, an operator can specify the VLAN ID for generated MEG sites, and also whether or not the MEP should be a Primary VLAN MEP or MIP (only applicable to SAP/SDP binding MEPs or MIPs on VPLS and Epipe services).

The CCM test execution routine checks if all MEPs in the MEG or MEG Sub-group match before allowing the test execution to proceed.

Points to consider when enabling the Primary VLAN function:

- A unique MA should be employed for each Primary VLAN.
- The Primary VLAN function is only applied to SAPs and spoke and mesh SDP bindings. Virtual MEPs are not supported.
- The MHF-Creation parameter must be set to the static option to force the generation of Primary VLAN MIPs.
- Primary VLAN MEPs can only be assigned at creation time.
- Primary VLAN MEPs must have an MA Service VLAN ID correctly set.
- The MA Service VLAN ID cannot be changed if Primary VLAN MEPs exist.
- Multiple Primary VLAN MEPs per SAP or SDP bindings are supported. There is no restriction on the number of MEPs that are allowed on SAPs/SDP binding for this feature. Both the classic extraction context and every Primary VLAN on the SAPs/SDP binding can support up to 16 MEPs (8 Up and 8 Down), one per MD level.
- All eight MD levels can be configured on all MEPs (0-7) that are configured on a SAP/SDP binding, within their specific context (classic or Primary VLAN). This allows an operator to have

the complete range of eight MD levels for use in the specified Primary VLAN. This avoids conflicts when the same MD Level and Layer 2 encapsulation maps to two different MEPs in the different contexts.

- Levels are specific within each of the contexts (classic and all the Primary VLANs configured on the same SAP/SDP binding). Therefore each lookup can have overlapping levels. The hierarchy is maintained within the extraction context.
- Primary VLAN MEPs and Fault Propagation are mutually exclusive.
- Primary VLAN MEPs and Fast CCM Interval are mutually exclusive.
- There is no sub-second CCM support for Primary VLAN enabled MEPs.

Procedures to configure Ethernet CFM

91.13 Ethernet CFM diagnostic test workflow

91.13.1 Stages

This workflow outlines the procedures that describe how to configure, manage and perform Ethernet CFM OAM diagnostic tests, as well as how to view the results the tests generate. The PM session tests are generally configured through the Service Test Manager. See Chapter 89, “Service Test Manager” for more information on working with the STM.

1

If the Ethernet CFM diagnostic test being configured is part of an STM test policy or STM test suite, review the STM chapter workflow to determine any prerequisite configuration requirements. See Chapter 89, “Service Test Manager” for more information on the STM.

2

Configure the initial MEP ID value on an NE for automatic MEP creation; see 91.14 “To configure an automatic MEP ID assignment on an NE” (p. 3108) .

3

Create an Ethernet CFM MD policy and subordinate objects associated with the MD such as a Global MEG, MEG, and MEP for each level at which Ethernet connectivity is to be monitored; see 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109).

4

As required, create identical MEPs on service SAPs that are part of an MC-LAG; see 91.16 “To automatically create identical MEPs on a redundant pair of service SAPs” (p. 3118).

5

As required, change the MEG sub-group association for managed MEPs or unmanaged remote MEPs; see 91.17 “To change the MEG sub-group association for managed MEPs or unmanaged remote MEPs” (p. 3119) .

6

As required, configure a default NE-level MD on an OmniSwitch; see 91.18 “To configure a default MD on an OmniSwitch” (p. 3120).

7

As required, configure one of the following CFM Ethernet diagnostic tests:

- Global MEG check; see 91.19 “To create and run a Global MEG OAM diagnostic test from the STM” (p. 3120)
- Continuity check and CFM dual-ended loss test; see 91.20 “To create and run a Continuity Check OAM diagnostic test from the STM” (p. 3122)

- CFM loopback; see 91.21 “To create and run a CFM loopback OAM diagnostic test from the STM” (p. 3123)
- CFM link trace; see91.22 “To create and run a CFM link trace OAM diagnostic test from the STM” (p. 3125)
- CFM Eth test; see 91.23 “To create and run a CFM Eth OAM diagnostic test from the STM” (p. 3127)
- CFM two way delay test; see 91.24 “To create and run a CFM two way delay OAM diagnostic test from the STM” (p. 3128)
- CFM one way delay test; see 91.25 “To create and run a CFM one-way delay OAM diagnostic test from the STM” (p. 3130)
- CFM single ended loss test; see 91.26 “To create and run a CFM single-ended loss OAM diagnostic test from the STM” (p. 3131)
- CFM two way SLM; see 91.27 “To create and run a CFM two way SLM OAM diagnostic test from the STM” (p. 3132)
- CFM LM test; see 91.28 “To create and run a CFM LM OAM diagnostic test from the STM” (p. 3134)

8

Create an STM test policy specifying at least one Ethernet CFM test definition; see 89.10 “To configure an STM test policy” (p. 2947).

9

Create an STM test suite; see 89.12 “To create an STM test suite” (p. 2951).

10

As required, execute the generated Ethernet CFM test suite; see 89.17 “To execute an STM test suite” (p. 2961).

11

As required, view and compare CFM Ethernet test results: See 89.18 “To view STM test suite results” (p. 2962) , and 89.19 “To view and compare STM test suite results for a tested entity” (p. 2963) ,

91.14 To configure an automatic MEP ID assignment on an NE

91.14.1 Steps

Perform this procedure to configure an initial MEP ID value and associated MEP parameters on an NE for automatic MEP creation.

1

On the equipment tree, right-click on a device and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

Click on the Globals tab and on the OAM tab.

3

Configure the MEP ID parameter.

4

Save your changes and close the form.

END OF STEPS

91.15 To configure an Ethernet CFM MD policy and subordinate objects

91.15.1 Related information

Many of the forms accessed in this procedure can also be accessed using the Service Topology map. See 4.2 “Working with topology maps” (p. 176) for information about working with maps.

Note: You can only perform this procedure when a service exists to associate with the MD policy or subordinate objects.

91.15.2 Steps

This policy defines the MD for Ethernet-based services such as VPLS, Epipe, VPRN, and IES, including SAPs and SDP bindings, for use in Ethernet CFM OAM diagnostic tests.

1

Choose Tools→Ethernet CFM→Maintenance Domain Policies from the NFM-P main menu. The Maintenance Domain Policies form opens.

2

Click Create or select an existing MD and click Properties. The Maintenance Domain - Global Policy (Create|Edit) form opens.

3

Configure the general parameters.

4

Click OK to save the policy and close the form, or click Apply to save the policy. See 49.6 “To release and distribute a policy” (p. 1476) to distribute the policy to NEs.

Add a Global MEG to an MD

5

To add a Global MEG to an MD:

1. Click on the Global Maintenance Entity Group tab and click Create. The Global Maintenance Entity Group (Create) form opens with the General tab displayed.
2. Configure the general parameters.
3. Configure the Name Format and Name parameters in the Maintenance Entity Group panel.

   If you select icc-based as the Name Format, then the associated Name value must be 8 to 13 characters long.
4. Configure the parameters in the Initial MEG Configuration panel.

   When you are creating a Global MEG for a 7750 SR-c4, 7750 SR-c12, or 7450 ESS-6, the Initial CCM Interval must be set to 1s.

   For Wavence nodes, the value *static* is not supported for the Initial MHF-Creation parameter. If static is selected during MEP configuration, then the default value *none* gets assigned to the node.

   You can only configure the Initial CFM Hold Down Timer (Centiseconds) parameter when the Initial CCM Interval is set to 10 ms or 100 ms.

   When the Initial CCM Interval parameter is set to 10 ms or 100 ms, you cannot configure automatic MEP creation in Step 7 .

   Note:

   When a Global MEG is created and applied, you can open the Global Maintenance Entity Group form at any time and view on the General tab two information fields that display the number of local MEGs and the number of MEPs currently associated with the Global MEG.
5. Save your changes.

Add a MEG sub-group to a Global MEG

6

To add a MEG sub-group to a Global MEG, perform one of the following:

a. Add a new MEG sub-group to a Global MEG.
   1. Select a Global MEG and click Properties. The Global Maintenance Entity Group (Edit) form opens with the General tab displayed.
   2. Click on the MEG Sub-Group tab and click Create. The MEG Sub-Group (Create) form opens.
   3. Configure the required parameters.
   4. Save your changes.

b. Automatically create a new MEG sub-group based on existing remote MEP lists.

Note: You can only create MEG sub-groups using this method when a valid remote MEP list exists and the remote MEPs on the list are CCM enabled.

Perform the following:

1. Select a Global MEG and click Properties. The Global Maintenance Entity Group (Edit) form opens with the General tab displayed.
2. Click Discover MEG Sub Groups. If this button is not visible, click More Actions and choose Discover MEG Sub Groups. When remote MEPs on valid lists are CCM enabled, new MEG sub-groups are created from the remote MEP lists. The new MEG sub-groups are listed on the Global Maintenance Entity Group (Edit) form.
3. Save your changes.

Associate a service with a Global MEG

7

To associate a service with a Global MEG, perform one of the following:

a. Automatically associate all NEs in a service with a Global MEG.

1. Select an MD and click Properties. The Maintenance Domain - Global Policy (Edit) form opens.
2. Click on the Global Maintenance Entity Group tab and click Create. The Global Maintenance Entity Group (Create) form opens with the General tab displayed.
3. Click on the Service tab and click Create. The Service (Create) form opens.
4. Select a Service.
5. Configure the Follow Service Topology Changes parameter, if required.

   Note:

   Enabling this parameter will cause relevant Ethernet OAM objects in the MEG and MEG sub-groups to automatically change in response to topology changes in the service or composite service you selected in 4 . Such changes include the addition or deletion of sites, SDP bindings, and access interfaces. Applicable CFM tests (as specified in a test suite) will also be automatically generated and/or deleted, as required. The specific Ethernet OAM objects that will respond to a service topology change are dependent on the selections you make in 7 below.
6. Configure the required parameters in the Auto-Creation of MEG(s) panel.
7. Configure the required parameters in the Auto-Creation of MEP(s) panel.

   Note:

   You cannot configure the MEP(s) Creation on Access Interfaces, MEP(s) Creation on SDP Bindings, or Virtual MEP(s) Creation on VPLS Sites parameters when the Initial CCM Interval parameter in Step 5 is set to 10 ms or 100 ms.

   For 7210 SAS-K NEs, to enable the Primary VLAN Enable parameter, you must enable the Primary VLAN parameter in the system resource profile for the device; see 12.50 “To configure the global system resource profile on a 7210 SAS or 7250 IXR” (p. 380).
8. Configure the required parameters in the Auto-Creation of MIP(s) panel.

9. Save your changes.

b. Associate a single NE on a service with a Global MEG.

1. Select an MD and click Properties. The Maintenance Domain Global Policy (Edit) form opens.
2. Click on the Global Maintenance Entity Group tab, select a Global MEG, and click Properties. The Global Maintenance Entity Group (Edit) form opens.
3. Click on the NE Maintenance Entity Group tab and click Create. The NE Maintenance Entity Group (Create) form opens with the General tab displayed.
4. Select a site.
5. Configure the required parameters.

   You can only configure the CFM Hold Down Timer (Centiseconds) parameter when the CCM interval is set to 10 ms or 100 ms.

   You cannot choose the 10 ms and 100 ms options for the CCM interval parameter when the CCM Padding Packet Size(Bytes) parameter is enabled.

   Note:

   Only the parameters that are supported on the selected site are displayed when you open the form.
6. Click on the Service tab.

   If the selected site supports only the service ID as the key identifier in an NE maintenance group, click on the Service ID tab, then click Create. The MEG Service (Create) form opens.

   If the selected site supports the service name as the key identifier in an NE maintenance group, both the Service ID and Service Name tabs are available. Click on the required tab, then click Create. The MEG Service (Create) form opens.

   On supporting NEs, you can choose either the service ID or the service name as the identifier for the associated service. To change from ID to name, or from name to ID, select the required tab, click Create, and configure the change. The previous configuration is removed, and the updated identifier is used.
7. Configure the required parameters.
8. Save your changes.

   Note:

   When an NE Maintenance Entity Group is created and applied, you can open the NE Maintenance Entity Group (Edit) form at any time and view on the General tab an information field that displays the number of MEPs currently associated with the NE Maintenance Entity Group.

Add an Ethernet path to the Global MEG

8

To add an Ethernet path to the Global MEG:

1. On the Global Maintenance Entity Group (Edit) form, click on the Ethernet Path tab and click Create. The Ethernet Path (Create) form opens.

   Note:

   You must specify an Ethernet tunnel path or Ethernet ring path on this form.
2. Select an Ethernet tunnel path or an Ethernet Ring Path.
3. Configure the required parameters.
4. Save your changes.

Add an NE MEG to the MD

9

To add an NE MEG to the MD:

1. Click on the NE Maintenance Entity Group tab and click Create. The Maintenance Entity Group (Create) form opens with the General tab displayed.
2. Select a site for the NE MEG.
3. Select a Unicast Hub MEP for the NE MEG. The Unicast Hub MEP panel is only displayed when the site you selected in 2 is a 7450 ESS, 7750 SR, or 7950 XRS NE.

   Configuring this parameter allows you to add a unicast MEP to the remote MEPs list. This enables the deployment of a hub-and-spoke topology where multiple MEPs acting as spokes can communicate with one hub MEP. In this scenario, the remote MEP has a MEP Mac Address which is the Operational MAC Address of the Unicast Hub MEP.

   See the MEP ID parameter description for additional information and constraints on configuring a unicast CCM hub.
4. Configure the required parameters. The following restrictions apply:
   - The VLAN ID parameter is configurable only when the selected NE is an OmniSwitch.
   - The CFM Hold Down Timer (Centiseconds) parameter is only configurable when the CCM interval is set to 10 ms or 100 ms.
   - The 10 ms and 100 ms options of the CCM interval parameter cannot be configured when the CCM Padding Packet Size(Bytes) parameter is enabled.
5. Click OK.
6. To associate a template with the NE MEG, select an NE MEG and click Properties. The Maintenance Entity Group (Edit) form opens. Otherwise, go to 8 .
7. Click on the Templates tab and select an associated template.
8. Click on the Service tab and click Create. The MEG Service (Create) form opens.
9. Configure the required parameters.

On supporting NEs, you can specify a user-defined ID in the Sender ID TLV of LBM and LTM PDUs. To use this option, set the ID-Permission parameter to the chassis option, and perform Procedure 12.44 “To configure the Sender-ID TLV of a CFM PDU for an NE” (p. 376).

10. Save your changes.

Add a managed MEP to the Global MEG

10

To add a managed MEP to the MEG:

1. Click on the Managed MEP tab and click Create. The MEP (Create) form opens.
2. Select a MEG.
3. Configure the required parameters. The following considerations apply:
   - The Interface Type parameter is configurable when the Type parameter is set to Regular.
   - The Interface Type Port is only available if the MD level is 0.
   - For facility MEPs, the Direction parameter must be set to Down.
4. Configure the required parameters in the Fault Alarm/Reset Time panel and the CCM panel.
5. Select a MEG subgroup to associate with the managed MEP.

   Note:

   If a MEG subgroup is not selected, the managed MEP is automatically associated with the default MEG subgroup.
6. Configure the required parameters in the Test Generation Options panel.

   Note:

   The Role, Use as Test Source, and Use as Test Target parameters are configurable only when you set the Direction parameter to Up.

   You can propagate test generation role settings to all MEPs on a SAP or service site using the Propagate to MEPs button on the ETH-CFM tab of the Access Interface (Edit) and Site (Edit) forms.
7. If the Type parameter you configured in 3 is set to Virtual, go to 12 .
8. If the Interface Type parameter you configured in 3 is set to LAG or Port, configure the Facility Fault Notify parameter.
9. If the Interface Type parameter you configured in 3 is set to LAG, configure the Facility VLAN ID.

   Note:

   If you are configuring a tunnel facility MEP, the Facility VLAN ID parameter value must match the outer encapsulation value of the SAP to which you are connecting the tunnel facility MEP.
10. In the bottommost panel, select an object of the type specified by the Interface Type parameter you configured in 3 . The Select *object_type* form opens.
11. Go to 13 .

12. Select a service site for the virtual MEP in the Service Site panel.
13. If the MD for the MEP has a Name Type of none and its Maintenance Association has a Name Format of icc-based, the Y.1731 Tests and AIS tabs are configurable. Otherwise, go to 15 .
14. Click on the Y.1731 TEST tab and AIS tab and configure the required parameters.
15. Click on the ETH-CSF tab and configure the required parameters.

    The CSF Rx Multiplier parameter is configurable when the Enable Client Signal Fault parameter is enabled.

    Note:

    ETH-CSF acts as one of the triggers for fault propagation on the receiving MEP, when the Fault Propagation parameter is enabled in 4 .

    When the CSF Rx Multiplier is 0 (meaning that CSF will never timeout) and you need to manually clear the CSF state, you must first disable and then re-enable the Enable Client Signal Fault parameter. Otherwise, CSF must receive a C-DCI flag in a PDU to clear the state automatically, whenever the multiplier is 0.

    If the CSF Rx Frame State read-only indicator displays anything other Client Defect Clear Indication, then a fault exists.

11

Save your changes.

Add a remote MEP to the MEG

12

Add a remote MEP to the MEG.

Note: When you execute a CCM test or synchronize the managed and unmanaged remote MEP lists, the NFM-P automatically distributes the local MEP and the remote MEP to the remote MEP list, except for 7705 SAR NEs, which are excluded from automatic MEP -distribution.
If you enabled the Enable Auto Remote MEP Discovery parameter in Step 7 b , the auto-discovered remote MEPs will populate this list when CCM is running on the network. You can clear the list of auto-discovered MEPs if required, by clicking on Clear Auto Discovered Remote MEPS.
When you execute a CCM test or synchronize the managed and unmanaged remote MEP lists within a MEG sub-group, the NFM-P automatically distributes the local MEP and the remote MEP within that MEG sub-group.
If the remote MEP you are adding is to act as a Unicast Hub MEP, the remote MEP's MEP Mac Address will become the Operational MAC Address of the Unicast Hub MEP. If you subsequently delete the remote MEP, the Operational MAC Address of the configured Unicast Hub MEP will also be cleared.
You can add only one entry to the remote MEP list for a 7705 SAR.

If a local MEP and the corresponding remote MEP are on the same OmniSwitch, deleting the remote MEP results also deletes the corresponding local MEP.

Perform the following:

1. Click on the Remote MEP tab. The Remote MEP (Create) form opens.
2. Configure the required parameters.

   See the MEP ID parameter description for additional information and constraints when configuring a remote MEP to act as a unicast CCM hub.
3. Save your changes.
4. You can also convert an existing remote MEP on this list that was previously automatically discovered to a static remote MEP by clicking the entry's Auto Discovered Remote Mep checkbox. The Remote MEP (Edit) form for that entry opens with the General tab displayed.
5. Disable the Auto Discovered Remote Mep checkbox and click OK. The Remote MEP (Edit) form closes.

13

Save your changes.

Add an unmanaged remote MEP to the Global MEG

14

To add an unmanaged remote MEP to the Global MEG:

1. Click on the Unmanaged Remote MEP tab and click Create. The Unmanaged Remote MEP (Create) form opens.
2. Configure the required parameters.
3. Select a MEG Sub-Group to associate the unmanaged remote MEP.

   Note:

   If a MEG sub-group is not selected, the unmanaged remote MEP is automatically associated with the default MEG subgroup.

   The outline color of a MEP icon indicates the MEG subgroup to which the MEP belongs.
4. Configure the required parameters in the Test Generation Options panel. The following considerations apply:
   - The Role, Use as Test Source, and Use as Test Target parameters are configurable only when you set the Direction parameter to Up.
   - You can only configure Test Generation Options when the unmanaged MEP has a MAC address.
   - You can propagate test generation role settings to all MEPs on a SAP or service site using the Propagate to MEPs button on the ETH-CFM tab of the Access Interface and Site (Edit) forms.
5. Save your changes.

15

Click Synchronize Remote MEPs to distribute the MEPs to the NEs, if required.

Note: If the managed MEP or unmanaged remote MEP belongs to a MEG sub-group, the remote MEP synchronization will only occur within the MEG sub-group.

16

Click Resync Remote MEP DB to re-synchronize the Remote MEPs database, if required.

Optionally, you can also view the DB state of any Remote MEP under a Managed MEP. Otherwise, go to Step 17 .

1. Click on the Managed MEP tab in the Global MEG form, select a Managed MEP, and click Properties. The MEP (Edit) form opens.
2. Click on the Remote MEP DB State tab, select a Remote MEP, and click Properties. The Remote MEP DB State (View) form opens, with the General tab displayed.

   The presented information includes:

   - MEP ID
   - Remote State
   - Grace Period RX
   - MAC Address
   - Rdi
   - Port Status
   - Interface Status
   - Last CCM

17

Click the Audit Unicast CCM item in the More Actions menu to automatically populate all the Hub pointers, if required.

Note: For discovered networks with CCM and a remote MEP list configured, Nokia highly recommends the use of this button to set up a correct HUB-spoke topology in NFM-P before executing a CCM test.
The Audit Unicast CCM action will use the remote lists from all NEs to automatically set the Hub MEP pointers to the appropriate MEPs. If an appropriate MEP is not found, an unmanaged MEP with MEP ID and MEP MAC Address matching the unicast remote MEP entry will be created. The Hub MEP pointer is then set to this unmanaged MEP.
If duplicate managed and unmanaged MEPs are detected, the unmanaged MEP is deleted and the managed MEP is used as the unicast Hub MEP. This action also ensures that the Hub MEP is in the same sub-group as the spoke pointing to it.

18

Save your changes and close the forms.

END OF STEPS

91.16 To automatically create identical MEPs on a redundant pair of service SAPs

91.16.1 LAG creation reference

See 13.16 “To create a LAG” (p. 431) for information on creating LAGs. See Chapter 43, “MC LAG groups” for more information about creating MC-LAGs.

91.16.2 Steps

Perform this procedure to create identical MEPs on service SAPs that are part of an MC-LAG. The procedure assumes that the service has been configured with MC-LAGs and redundant SAPs. The pair of redundant SAPs must have identical inner and outer encapsulations configured and they must be in the same service.

1

Select Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Topology View. The Service Topology map opens.

3

Select the OAM check box at the bottom left of the map window. MEG and MEP objects are added to the map.

4

Perform one of the following:

a. If you are creating a Global MEG:

1. Right-click on the topology map background and select Create CFM Global MEG. The Global Maintenance Entity Group (Create) form opens. See 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) for information about configuring a Global MEG.
2. When the Global MEG is created, a purple triangular MEP icon appears above each of the SAPs. Right-click on one of the MEP icons associated with the redundant pair of SAPs. The MEP (Edit) form for that SAP opens.
3. Click Properties in the Redundancy panel.

   This action automatically populates the Properties page of the redundant partner, creating identical MEPs on both members of the redundant pair. Leave the MEP (Edit) form open.
4. Right-click on the other MEP icon associated with the redundant pair of SAPs. The MEP (Edit) form for that SAP opens.
5. Confirm that the ID and Operational Mac Address parameters in the Redundant MEP panel for both MEPs are identical.

   Note:

If the redundant SAPs are in a VPRN or IES service, only the ID parameter will be synchronized between the two SAPs. MAC addresses are not supported for these services.

b. If you are working with an existing Global MEG where a MEP is missing from one of the SAPs in a redundant pair:

1. Right-click on the MEP icon that is present for a SAP in the redundant pair. The MEP (Edit) form for that SAP opens.
2. Enable the Create Redundant MEP check box in the Redundancy panel and click Apply. A new MEP icon appears above the other SAP of the redundant pair.

   This action automatically populates the Properties page of the redundant partner which was missing its MEP, thereby creating identical MEPs on both members of the redundant pair. Leave the MEP (Edit) form open.
3. Right-click on the new MEP icon. The MEP (Edit) form for that SAP opens.
4. Confirm that the ID and Operational Mac Address parameters in the Redundant MEP panel for both MEPs are identical.

   Note:

   If the redundant SAPs are in a VPRN or IES service, only the ID parameter will be synchronized between the two SAPs. MAC addresses are not supported for these services.

5

Save your changes and close the forms.

END OF STEPS

91.17 To change the MEG sub-group association for managed MEPs or unmanaged remote MEPs

91.17.1 Steps

Note: The outline color of a MEP icon indicates the MEG sub-group to which the MEP belongs.

1

Select Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Select a service and click Topology View. The Service Topology map opens.

3

Select the OAM check box at the bottom left of the map window. MEG and MEP objects are added to the map.

4

Right-click on the interface object that includes the MEP that you wish to move to another MEG sub-group and choose Ethernet CFM→Edit MEP. The Select MEP (Edit) form opens.

5

In the MEG Sub-Group panel, click Clear to remove the current MEG sub-group association.

6

Select a MEG sub-group.

7

Save your changes and close the form.

END OF STEPS

91.18 To configure a default MD on an OmniSwitch

91.18.1 Steps

1

On the equipment tree, expand an OmniSwitch NE icon, right-click on the device and select Properties. The Network Element (Edit) form opens with the General tab displayed.

2

Click on the Globals tab and on the CFM tab.

3

Configure the required parameters.

4

Save your changes and close the form.

END OF STEPS

91.19 To create and run a Global MEG OAM diagnostic test from the STM

91.19.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→Create Global MEG. The Global Maintenance Entity Group, Test (Create) form opens.

3

Configure the required general parameters.

4

Perform one of the following:

a. To choose an existing MD, select an MD in the Maintenance Domain panel.

b. To create a new MD:

1. Click Create. The Maintenance Domain, Global Policy (Create) form opens.
2. Configure the required parameters.
3. Save the changes, close the form and select the MD you created.

5

Configure the required parameters in the Maintenance Entity Group panel.

If you select icc-based as the Name Format, the associated Name value must be 8 to 13 characters.

6

Configure the required parameters in the Initial MEG Configuration panel.

When you are creating a Global MEG for a 7750 SR-c4, 7750 SR-c12, or 7450 ESS-6, the Initial CCM Interval must be set to 1s.

You can only configure the Initial CFM Hold Down Timer (Centiseconds) parameter when the Initial CCM Interval is set to 10 ms or 100 ms.

When the Initial CCM Interval parameter is set to 10 ms or 100 ms, you cannot configure automatic MEP creation.

7

As required, associate a Service to the Global MEG. Perform Step 7 of 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) .

8

As required, add an Ethernet path to the Global MEG. Perform Step 8 of 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) .

9

As required, add a managed MEP to the Global MEG. Perform Step 10 of 91.15 “To configure an Ethernet CFM MD policy and subordinate objects” (p. 3109) .

10

Save the changes and close the form.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.20 To create and run a Continuity Check OAM diagnostic test from the STM

91.20.1 Steps

Note: The NFM-P automatically creates a Continuity Check OAM diagnostic test when you create an MD and global MEG. You can execute an automatically created test using the STM, or from the Tests tab of a MEG.

If a service is modified after you associate it with an MD, you must manually add new MEPs; for example, when a new site is added.

If a continuity check test is running on a service when you add a new MEP to the service, you must stop the test and execute it again to make the new MEP active.

When a service is modified after a continuity check is created, you must manually add a virtual MEP to the site; for example, when a new B-VPLS site is added to the service.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→Create Continuity Check. The CFM Continuity Check Test (Create) form opens.

3

Configure the required parameters.

4

Select a MEG in the Maintenance Entity Group panel.

5

Configure the Duration (minutes) parameter in the Connectivity Check Timer panel.

6

If you are configuring the test for a 7705 SAR, configure the required parameters in the Initial Dual Ended Loss Test Options panel.

7

Save the changes and close the form.

8

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.21 To create and run a CFM loopback OAM diagnostic test from the STM

91.21.1 Steps

Note: The maximum number of continuous CFM loopback tests that can be executed is 300, provided that the SAA packet per second rate of 200 pps is not exceeded. The total number of continuous tests must share this maximum rate of 200 pps.

You can also run this test contextually for VPLS, Epipe, and Composite services. See 90.43 “To configure and run OAM tests contextually” (p. 3053) for more information.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Loopback. The CFM Loopback Test (Create) form opens.

3

Configure the required parameters.

4

Select a global MEG next to the Global ID parameter.

5

Select the originating MEP next to the ID parameter.

6

To select the test destination using the target MEP ID:

1. Enable the Enable Target MEP ID parameter.
2. Configure the Target MEP ID parameter or select the target MEP ID to choose the destination MEP.

   Note:

   MEP ID selection is only supported for OmniSwitch NEs. The NE Schedulable parameter must be enabled in Step 3 .

7

In the MEP Transmit Information panel, choose one of the following for the service destination test endpoint:

a. Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.

Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

Note: The MIP selection is not supported for OmniSwitch NEs.

b. Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

8

If you enabled the NE Schedulable parameter in Step 3 , go to Step 9 . Otherwise, configure the MEP Transmit LBM Information parameters.

9

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

10

Click on the Results Configuration tab and configure the required parameters.

The Trap Generation parameter is only configurable when the Accounting Files parameter is disabled.

The Test Result Storage parameter is only configurable when the Accounting Files parameter is enabled.

11

Save the changes and close the form.

12

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.22 To create and run a CFM link trace OAM diagnostic test from the STM

91.22.1 Steps

Note: The maximum number of continuous CFM link trace tests that can be executed is 300, provided that the SAA packet per second rate of 200 pps is not exceeded. The total number of continuous tests must share this maximum rate of 200 pps. However, it is recommended that the number of continuous CFM link trace tests executed be kept well below the maximum allowed limit.

You can also run this test contextually for VPLS, Epipe, and Composite services. See 90.43 “To configure and run OAM tests contextually” (p. 3053) for more information.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Link Trace. The CFM Link Trace (Create) form opens.

3

Configure the required parameters.

If you enable the NE Schedulable parameter, two additional tabs are displayed for the form: Test Parameters and Results Configuration.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

In the MEP Transmit Information panel, choose one of the following for the service destination test endpoint:

a. Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.

Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

Note: The MIP selection is not supported for OmniSwitch NEs.

b. Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

7

If you enabled the NE Schedulable parameter in Step 3 , go to Step 8 . Otherwise, configure the MEP Transmit LBM Information parameters.

8

Click on the Test Parameters tab and configure the REQUIRED parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

9

Click on the Results Configuration tab and configure the required parameters.

The Trap Generation parameter is only configurable when the Accounting Files parameter is disabled.

The Test Result Storage parameter is only configurable when the Accounting Files parameter is enabled.

10

Save the changes and close the form.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.23 To create and run a CFM Eth OAM diagnostic test from the STM

91.23.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Eth Test. The CFM Eth Test (Create) form opens.

3

Configure the required parameters.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

In the MEP Transmit Information panel, choose one of the following for the service destination test endpoint:

a. Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.

Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

Note: The MIP selection is not supported for OmniSwitch NEs.

b. Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

7

Configure the MEP Transmit LBM Information parameters.

8

Save the changes and close the forms.

9

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.24 To create and run a CFM two way delay OAM diagnostic test from the STM

91.24.1 Steps

Note: The maximum number of continuous CFM two way delay tests that can be executed is 300, provided that the SAA packet per second rate of 200 pps is not exceeded. The total number of continuous tests must share this maximum rate of 200 pps.

You can also run this test contextually for VPLS, Epipe, and Composite services. See 90.43 “To configure and run OAM tests contextually” (p. 3053) for more information.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Two Way Delay Test. The CFM Two Way Delay Test (Create) form opens.

3

Configure the required parameters.

If the NE Schedulable parameter is enabled, two additional tabs are displayed for the form, Test Parameters and Results Configuration.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

To select the test destination using the target MEP ID:

1. Enable the Enable Target MEP ID parameter.
2. Configure the Target MEP ID parameter or select the target MEP ID to choose the destination MEP.

   Note:

   MEP ID selection is only supported for OmniSwitch NEs. The NE Schedulable parameter must be enabled in Step 3 .

7

In the MEP Transmit Information panel, choose one of the following for the service destination test endpoint:

a. Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.

Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

Note: The MIP selection is not supported for OmniSwitch NEs.

b. Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

8

If you enabled the NE Schedulable parameter in Step 3 , go to Step 9 . Otherwise, configure the MEP Transmit LBM Information parameters.

9

If the NE Schedulable parameter was enabled in Step 3 , go to Step 11 . Otherwise, go to Step 10 .

10

Configure the MEP Transmit DMM Information VLAN Priority parameter. Go to Step 13 .

11

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

12

Click on the Results Configuration tab and configure the required parameters.

The Trap Generation parameter is only configurable when the Accounting Files parameter is disabled.

The Test Result Storage parameter is only configurable when the Accounting Files parameter is enabled.

13

Save the changes and close the forms.

14

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

91.25 To create and run a CFM one-way delay OAM diagnostic test from the STM

91.25.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM One Way Delay Test. The CFM One Way Delay Test (Create) form opens.

3

Configure the required parameters.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

In the MEP Transmit Information panel, choose one of the following for the service destination test endpoint:

a. Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.

Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

Note: The MIP selection is not supported for OmniSwitch NEs.

b. Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

7

Configure the Priority parameter in the MEP Transmit LBM Information panel.

8

Save the changes and close the forms.

9

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

End of steps

91.26 To create and run a CFM single-ended loss OAM diagnostic test from the STM

91.26.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Single Ended Loss Test. The CFM Single Ended Loss Test (Create) form opens.

3

Configure the required parameters.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

To select a MEP as the test destination, click Select MEP.

7

Configure the required parameters in the MEP Transmit LMM Information panel.

8

Save the changes and close the forms.

9

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.27 To create and run a CFM two way SLM OAM diagnostic test from the STM

91.27.1 Steps

Note: The maximum number of continuous CFM two way SLM tests that can be executed is 300, provided that the SAA packet per second rate of 200 pps is not exceeded. The total number of continuous tests must share this maximum rate of 200 pps.

You can also run this test contextually for VPLS, Epipe, and Composite services. See 90.43 “To configure and run OAM tests contextually” (p. 3053) for more information.

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM Two Way SLM. The CFM Two Way SLM Test (Create) form opens

3

Configure the required parameters.

If the NE Schedulable parameter is enabled, two additional tabs are displayed for the form, Test Parameters and Results Configuration.

4

Select a global MEG in the Test Object panel.

5

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

6

To select the test destination using the target MAC Address:

a. To select a MEP as the test destination, click Select MEP.

b. To select a MIP as the test destination, click Select MIP.

Note: MIP selection is not supported for OmniSwitch NEs.

c. To select an unmanaged MEP as the test destination, click Select Unmanaged MEP.

7

If the NE Schedulable parameter was not enabled in Step 3 , configure the REQUIRED parameters in the MEP Transmit SLM Information panel and go to Step 10 . Otherwise, go to Step 8 .

8

Click on the Test Parameters tab and configure the required parameters.

The Probe History parameter is only configurable when the NE Schedulable parameter is enabled.

9

Click on the Results Configuration tab and configure the required parameters.

The Trap Generation parameter is only configurable when the Accounting Files parameter is disabled.

The Test Result Storage parameter is only configurable when the Accounting Files parameter is enabled.

10

Save the changes and close the forms.

11

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

91.28 To create and run a CFM LM OAM diagnostic test from the STM

91.28.1 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager form opens.

2

Click Create and choose Ethernet CFM→CFM LM Test. The CFM LM Test (Create) form opens.

3

Configure the required parameters.

4

Disable the NE Schedulable parameter.

5

Select a global MEG in the Test Object panel.

6

Select an originating MEP in the Test Object panel.

Note: MEPs can be added to services automatically or manually. If no MEPs are listed when you perform this step, see “Ethernet CFM” (p. 3089) for information about how to configure a MEP that is required for this step.

7

To select the test destination using the target MAC Address:

a. To select a MEP as the test destination, click Select MEP.

b. To select a MIP as the test destination, click Select MIP.

Note: MIP selection is not supported for OmniSwitch NEs.

c. To select an unmanaged MEP as the test destination, click Select Unmanaged MEP.

8

Configure the required parameters in the MEP Transmit LM Information panel.

9

Save the changes and close the forms.

10

To run the OAM diagnostic test and view the results, perform 89.15 “To run one or more OAM diagnostic tests from the STM and view the test results” (p. 2959) . See 89.20 “To interpret OAM diagnostic test results on the STM” (p. 2964) for information about how to interpret the test results.

END OF STEPS

92 Performance Monitoring tests

92.1 Overview

92.1.1 Purpose

This chapter describes PM (Performance Monitoring) OAM diagnostic tests, and provides information about configuring and performing PM diagnostic tests.

92.1.2 Contents

PM testing overview

92.2 PM session tests

92.2.1 Introduction

This section describes NFM-P supported PM OAM diagnostic test functionality. Unless noted, all tests are accessible from the STM. See Table 90-1, “NFM-P supported OAM diagnostic tests and configurations” (p. 2981) for a list of all supported OAM diagnostic tests and their applicable procedures.

92.2.2 PM session test support and configuration

The NFM-P supports PM session tests for Ethernet CFM, IP TWAMP Light, and MPLS DM.

You can configure PM session related objects from several areas in the NFM-P, including:

- STM; see 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146)
- Service Topology maps; see 4.2.3 “Working with Ethernet CFM objects” (p. 177)
- Layer 3 routing instances for IP-based PM sessions; see 27.2 “To configure a routing instance or a VRF instance” (p. 826)
- VPRN service site form for IP-based PM sessions; see 79.76 “To create a TWAMP Light reflector on a VPRN site” (p. 2647)
- VPRN service form (OAM tab→TWAMP sub-tab) for IP-based PM sessions. Perform the required steps in this procedure to create a PM session and then perform 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) to create a TWAMP Light test session.
- Dynamic LSP properties form for MPLS DM PM sessions (for dynamic LSP types only); see 31.13 “To configure a Dynamic or segment routing TE LSP” (p. 1132) .

Ethernet CFM

PM session CFM testing is based on Metro Ethernet Forum Specification 35 - Service OAM Performance Monitoring Implementation Agreement, which details a standardized method to test and report network delay and loss using CFM messaging. This testing is performed in Layer 2 networks.

TWAMP IP level monitoring

The PM session testing framework can also be utilized in the IP domain to perform TWAMP (Two-Way Active Measurement Protocol) IP level monitoring, and runs over IPv4 and IPv6 addresses. The TWAMP Light and TCC tests target Layer 3 interfaces. They provide options to monitor IP SLA performance as related to KPI.

The TCA events and their associated statistics can be configured on all PM session test types.

PM session results

Results reporting occurs at standardized measurement intervals (5 minutes, 15 minutes, 1 hour, and 1 day), and comprises statistical summaries of individual test frame results. Statistical

summaries include minimum, maximum, and average values of specific test criteria. SNMP-mediated statistics can be collected to display real-time results for in-progress measurement intervals. Additionally, a histogram can be generated to characterize the distribution of test results over the same measurement intervals. Statistics can be viewed using the individual test forms, the STM, or the Statistics Manager.

PM session test suite support

While PM session tests may be configured individually, test suite support allows for repeatable, policy driven auto-generation of PM sessions and associated tests, which are ideally suited to larger scale environments. See 89.12 “To create an STM test suite” (p. 2951) for information about how to configure PM session tests as part of an STM test suite.

92.2.3 CFM DMM session test

For CFM DMM (Delay Measurement Message) tests, calculations are made to report on three criteria: frame delay, frame delay range, and interframe delay variation. DMM test frames are issued at regular intervals from a source MEP. Delay measurement information for the forward, backward, and round trip path is determined from the DMR frames received from the destination MEP.

See 92.7 “To configure a CFM DMM session OAM diagnostic test from the STM” (p. 3149) for information about configuring a CFM DMM test.

92.2.4 CFM SLM session test

The CFM SLM (Synthetic Loss Measurement) session test is an extension of the Y.1731 standard that provides a method of exchanging transmit and receive counters to determine frame loss between a MEP and the destination MAC address or remote MEP ID of another node in the network. This test is used to verify MEP-MEP connectivity in the network and can be used to approximate the frame loss of actual data traffic. CFM SLM session tests measure frame loss using synthetic frames, rather than data traffic. Frame loss is measured by calculating the difference between the number of synthetic frames that are sent and received.

See 92.8 “To configure a CFM SLM session OAM diagnostic test from the STM” (p. 3150) for information about configuring a CFM SLM session test.

92.2.5 CFM LMM session test

The CFM LMM (Loss Measurement Message, single-ended) session test is a method of exchanging transmit and receive counters between peer MEPs to determine exact loss on a point-to-point Ethernet virtual circuit.

Note: LMM tests are limited to MEPs in the Down direction only. The following validations are performed:

- The LMM test suite tests can only be created or executed for MEG subgroups with exactly two MEPs.
- The LMM test suite tests can only be executed if the source and target MEPs are not already the source or target MEPs (respectively) of a currently running LMM test.

See 92.9 “To configure a CFM LMM session OAM diagnostic test from the STM” (p. 3151) for information about configuring a PM CFM LMM test.

92.2.6 TWAMP Light session test

The TWAMP Light session test targets Layer 3 interfaces and requires a TWAMP reflector in order to execute properly. See 92.11 “To configure a TWAMP Light reflector” (p. 3153) .

You can configure a TWAMP Light session test on the following network objects:

- STM; see 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154)
- VPRN service site; see 79.76 “To create a TWAMP Light reflector on a VPRN site” (p. 2647)
- Base routing instance; see 27.2 “To configure a routing instance or a VRF instance” (p. 826)

92.2.7 TCC test

A TWAMP server is a combination of control server and session reflector, but requires an external probe called the TWAMP controller (testing head). The TWAMP Light test supports both a session sender and session reflector. TWAMP Light does not have a separate control plane, and the session reflector is configured in the NFM-P.

The TCC (TWAMP and TWAMP Light Control Client) test allows a TWAMP Light NE launch point (base router) to test TWAMP functionality against any TWAMP server that does not support TWAMP Light. The NFM-P creates the required TCP control channel with the TWAMP server, as well as the PM sessions and tests on the TWAMP Light launch point. The NFM-P also provides a report on a managed TWAMP server's capabilities (based on SNMP queries), which is accessible on the Globals→OAM sub-tab of the NE's properties form.

For managed RFC5938–compliant TWAMP servers, individual session control allows start and stop control messaging. This permits the client to define new tests and to start or stop specific tests while the TWAMP connection is active. The following operational points apply:

- When executing a test, a Supports Individual Test Control indicator on the General tab dynamically reports on server RFC5938 compliance.
- The Execute All and Stop All buttons at the bottom of the form affect all of the test sessions.
- Selective test session control is performed on the TWL Session tab, using the Execute or Stop Test Execution buttons at the right side of the form. Selected sessions must be administratively enabled to allow stopping and administratively disabled to start. A test's administrative status can be configured on its properties form, available from this tab.

See 92.13 “To configure a TCC OAM diagnostic test from the STM” (p. 3155) for information about configuring a TCC test.

92.2.8 MPLS DM session test

The MPLS DM (Delay Measurement) session test is based on RFC 6374, and provides calculations for LSPs for frame delay, frame delay range, and interframe delay variation. Measurements are one-way for dynamic RSVP and RSVP-auto LSP types, and round-trip for MPLS-TP Static LSP types. Test frame intervals are configurable from 1000 to 8000 ms, in 1000 ms increments. As a security feature, a global setting in the NE properties must be enabled for test packets to be transmitted or reflected.

See 92.10 “To configure an MPLS DM session OAM diagnostic test from the STM” (p. 3152) for information about configuring an MPLS DM test.

Workflow to configure and perform performance monitoring testing

92.3 PM diagnostic test workflow

92.3.1 Stages

This workflow outlines the procedures that describe how to configure, manage, and perform PM OAM diagnostic tests, as well as how to view the results the tests generate. PM tests are typically configured through the Service Test Manager. See Chapter 89, “Service Test Manager” for more information on working with the STM. You can also configure some PM-related objects from the Service Topology map. See 4.2.3 “Working with Ethernet CFM objects” (p. 177) in Chapter 4, “Topology map management” .

1

If the PM diagnostic test being configured is part of an STM test policy or STM test suite, review the STM chapter workflow to determine any prerequisite configuration requirements. See Chapter 89, “Service Test Manager” for more information on the STM.

2

Create and distribute a PM bin group policy that configures the statistical bins used by the various PM tests to collect their data. Perform 92.4 “To configure a PM bin group policy” (p. 3144) .

3

Create and distribute a new accounting policy, required to configure a PM session. Choose “Complete-PM” for the Type parameter. See “To configure an accounting policy” in the *NSP NFM-P Statistics Management Guide* for information on creating an accounting policy. See 49.6 “To release and distribute a policy” (p. 1476) for information on distributing a policy.

4

Configure a TWAMP Light reflector; see 92.11 “To configure a TWAMP Light reflector” (p. 3153). This is only required for proper operation of a PM TWAMP Light test session.

5

Configure a PM session to provide the session framework; see 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) .

6

To generate accounting statistics, associate the accounting policy to the PM session. If the accounting policy is not associated with the PM session, the test results will be collected as performance statistics.

To associate the accounting policy with the PM Session:

1. Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.
2. Select a PM Session and click Properties. The PM Session (Edit) form opens.
3. Click on the Components tab.
4. Right-click on Measurement Intervals and choose Create. The Measurement Interval (Create) form opens.
5. In the Accounting Policy panel, click Select and choose a policy.
6. Click OK.
7. Close the open forms.

7

As required, configure one of the following PM diagnostic tests:

- CFM DMM session test; see 92.7 “To configure a CFM DMM session OAM diagnostic test from the STM” (p. 3149) .
- CFM SLM session test; see 92.8 “To configure a CFM SLM session OAM diagnostic test from the STM” (p. 3150) .
- CFM LMM session test; see 92.9 “To configure a CFM LMM session OAM diagnostic test from the STM” (p. 3151) .
- MPLS DM session test; see 92.10 “To configure an MPLS DM session OAM diagnostic test from the STM” (p. 3152).
- TWAMP Light session test; see 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) .
- TCC test; see 92.13 “To configure a TCC OAM diagnostic test from the STM” (p. 3155) .

8

Create an STM test policy, specifying at least one PM diagnostic test definition.

See 89.10 “To configure an STM test policy” (p. 2947) for more information.

9

Create an STM test suite; see 89.12 “To create an STM test suite” (p. 2951).

10

As required, execute the generated PM test suite; see 89.17 “To execute an STM test suite” (p. 2961) .

11

As required, view and compare PM diagnostic test results:

- From a specific PM diagnostic test form; see92.14 “To collect and view PM statistics from a test form” (p. 3156) .
- From the STM; see92.15 “To view PM test results in the STM” (p. 3157) .
- From the Statistics Manager; see92.16 “To view PM test statistics in the Statistics Manager” (p. 3159) .
- For OAM PM event server performance statistics; see92.17 “To view OAM PM Event server performance statistics” (p. 3160) .
- To view STM test suite results; see 89.18 “To view STM test suite results” (p. 2962) .
- To view and compare STM test suite results; see 89.18 “To view STM test suite results” (p. 2962) .

92.3.2 Hardware configuration requirements

The PM session testing functionality requires the following hardware configurations to operate properly:

- Platform: 7750 SR 7/12, 7450 ESS 7/12, 7750 SRc 4/12 and 7950 XRS all require CPM3 or above for binning. The 7705 SAR does not support the creation of PM session tests.
- IOM support: IOM3/IMM and above
- Chassis mode: All SAPs executing this function must reside on IOM3/IMM and above. However, not every SAP in an affected service needs to be configured on IOM3/IMM or above. All active network ports on the affected NE must exist only on IOM3/IMM and above. The term given to this mode of operation is Network Chassis Mode D.

92.3.3 Recording test results

In addition to the procedures for viewing the OAM PM test results, you can also use the LogToFile method to record test results for transfer to external OSS applications. See the *NSP NFM-P XML API Developer Guide* for more information.

Procedures to configure and perform PM testing

92.4 To configure a PM bin group policy

92.4.1 General information

You must configure and distribute a PM bin group policy to define the structure of the statistical bins and delay counts used to collect statistical data by the various PM session tests.

92.4.2 Steps

1

Click Tools→PM Bin Group Policies on the NFM-P main menu. The PM Bin Group Policies form opens.

2

Click Create or select an existing bin group and click Properties. The Bin Group (Create|Edit) form opens.

3

Configure the required parameters.

4

Click on the Bins tab. A list of pre-configured default bins is displayed.

1. Click Create or select an existing bin object and click Properties. The Bin (Create|Edit) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click on the Delay-events tab to collect TCA event statistics.

1. Select a Bin Type from the list and click Properties. The TCA Config form opens.
2. Configure the required parameters.

   As an example of the bin exclusion logic, if you configure the Lowest Bin = 4, the Lowest Bin to Exclude = Bin 5, and enable Bin 0 and Bin 6 in the Excluded Bins for Average block, then the following will occur:

   - Bins 5 and 6 will be excluded from triggering a threshold crossing alarm for this Bin Group.
   - Bins 0 and 6 will be excluded when the calculation of average delay is performed for an OAM-PM measurement interval.
3. Save your changes and close the form.

6

Click OK to save the policy and close the form, or click Apply to save the policy.

7

Perform 49.6 “To release and distribute a policy” (p. 1476) to associate the policy to NEs.

END OF STEPS

92.5 To configure a streaming delay template

92.5.1 Before you begin

Streaming delay templates are associated under CFM DMM, MPLS DM, and TWAMP Light sessions to allow telemetry to collect PM streaming via gNMI.

92.5.2 Steps

1

Click Tools→OAM Templates→Streaming Delay Template on theNFM-P main menu. The Streaming Delay Template form opens.

2

Click Create, or select an existing template and click Properties. The Streaming Delay Template, Global Policy (Create|Edit) form opens.

3

Configure the required parameters.

4

Create streaming metrics.

1. Click the Streaming Delay Metrics tab and click Create. The Streaming Measurement (Create) form opens.
2. Configure the required parameters.
3. Save your changes and close the form.

5

Click OK to save the template and close the form, or click Apply to save the template.

6

Associate the template to a CFM DMM, MPLS DM, or TWAMP Light session by setting the Streaming Delay Template parameter on the associated session test configuration form.

END OF STEPS

92.6 To configure a PM session OAM diagnostic test from the STM

92.6.1 Steps

Configure the PM session framework

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create→Performance Monitoring→PM Session or select an existing PM session and click Properties. The PM Session (Create|Edit) form opens.

Note: A bin group with ID 1 must exist before you can create a PM session. See 92.4 “To configure a PM bin group policy” (p. 3144) for information about creating bin groups.

3

Configure the required parameters.

If you configure the Test Family Type as IP, then you can also select an existing TCC test in the TWAMP panel. See 92.13 “To configure a TCC OAM diagnostic test from the STM” (p. 3155) for information on configuring a TCC test.

4

Select a System ID in the Network Element panel.

If you are creating a PM session from a Routing Instance form, the Network Element System ID is automatically populated.

5

Select a Bin Group policy in the Bin Group panel.

6

Depending on the Test Family Type selected in Step 3, perform one of the following:

a. Click on the Ethernet Session tab.
   1. Configure the Priority parameter.
   2. Select a Maintenance Entity Group ID in the Source (Controller) MEP information panel.
   3. Choose an originating MEP using the ID Select button.
   4. In the Destination (Responder) MEP information panel, choose one of the following for the service destination test endpoint:

- Select a MEP, MIP, or Unmanaged MEP for the Target MAC Address parameter.
  Note: You must select a Maintenance Entity Group ID and Originating MEP prior to configuring this parameter for the Select buttons to function.

- Click the Enable Target MEP ID checkbox if you want to specify a remote Maintenance Entity Point ID (MEP ID) endpoint.

  Note: With both options, you have the choice to either enter a value manually for the Target MEP ID or Target MAC Address parameter or use the selected value.

b. Click on the IP Session tab.

1. Configure the IP Session Target Type parameter. If you select the VPRN Service option, you must also select a VPRN service for the selected NE.

   If you are creating a PM session from a VPRN service form, the Service ID is automatically populated. If the PM session is created from the VPRN site context, the Network Element System ID, IP Session Target Type, and VPRN Service ID parameters are automatically configured.

2. Select source and destination IP addresses in the Source and Destination Addresses panel.

   The source IP and destination IP addresses are required to create a working TWAMP Light or TCC test. The source IP address is selected from a list of IP addresses on the NE specified in the VPRN service or routing instance. The destination IP address is selected from all IP addresses that are not configured on the NE in the VPRN service or routing instance. The destination IP address and destination UDP port specify the location of the TWAMP Light reflector that a TWAMP Light test will use. See 92.11 “To configure a TWAMP Light reflector” (p. 3153) for information on creating TWAMP reflectors.

   If you are creating the TWAMP Light or TCC test from a service topology map, the OAM option must be enabled in the map. Selecting two L3 access interfaces automatically populates both the Source IP and Destination IP fields, based on the SAPs you select. Selecting just one SAP populates only the Source IP field.

3. Configure the required parameters in the Session Details panel.

   The Bypass Routing, Egress IF Name, and Next Hop IP parameter are used to determine what kind of forwarding is used. You can only configure one of these parameters.

   The Source UDP Port parameter should only be specified when configuring the PM session for use in a PM TCC test.

   The DSCP Name, DSCP Egress Remark, Do Not Fragment, Use Pattern, Tunnel, and Mpls Tunnel parameters are specified when configuring the PM session for use in a TWAMP Light session test.

c. Click on the MPLS Session tab.

1. Configure the LSP Type parameter.

   The panels and parameters that appear vary depending on the LSP Type selected.

2. Select an LSP in the resulting panel.

   MPLS-TP Static LSPs are bidirectional and allow forward, backward, and round-trip delay metric calculation.

   RSVP and RSVP Auto LSPs require configuration of the Return Address parameter.

3. Configure the parameters in the Details panel.

   The Force Enable MPLS-DM parameter automatically enables transmission of test frames on test NEs, and reflection of test frames on supported destination NEs. If you do not enable the Force Enable MPLS-DM parameter, you must manually enable the

MPLS-DM Admin Status parameter on the NE properties form; see 92.10 “To configure an MPLS DM session OAM diagnostic test from the STM” (p. 3152).

7

Click Apply.

Configure and enable the PM session tests

8

Click on the Components tab.

1. Right-click the Measurement Intervals item and select Create Measurement Interval. The Measurement Interval form opens.

   A PM session test will record results information to a “raw” bin by default. This is a measurement interval with infinite duration. However, one or more finite measurement intervals may also be configured for the same PM session and associated tests. Measurement intervals can be configured as 5 minutes, 15 minutes, 1 hour, or 1 day in length. If you create more than one measurement interval, the Interval Duration must be different for each one.

2. Configure the required parameters.

   The Enable events, Delay Events, and Loss Events parameters are required to collect TCA event statistics. The Loss Events parameter is not configurable for MPLS DM tests.

3. Select an accounting policy.

9

Perform one of the following:

a. If you configured the Test Family parameter in Step 3 for Ethernet, right-click the Test Sessions icon and select one of the following tests from the contextual menu:

   - Create CFM DMM Test Session. Perform 92.7 “To configure a CFM DMM session OAM diagnostic test from the STM” (p. 3149) .
   - Create CFM SLM Test Session. Perform 92.8 “To configure a CFM SLM session OAM diagnostic test from the STM” (p. 3150) .
   - Create CFM LMM Test Session. Perform 92.9 “To configure a CFM LMM session OAM diagnostic test from the STM” (p. 3151) .

   You can create one of each type of these Ethernet tests for a single PM session.

b. If you configured the Test Family parameter in Step 3 for IP, right-click the Test Sessions icon, select Create TWAMP Light Test Session, and perform 92.12 “To configure a TWAMP Light session OAM diagnostic test from the STM” (p. 3154) .

c. If you configured the Test Family parameter in Step 3 for MPLS, right-click the Test Sessions icon, select Create MPLS DM Test Session, and perform 92.10 “To configure an MPLS DM session OAM diagnostic test from the STM” (p. 3152).

10

Save your changes and close the PM Session form.

If the Administrative State of either or both of the test sessions is set to Enabled, the tests will begin to execute automatically when the changes are applied. They continue to run until manually stopped.

Additional operations

11

To stop a test from running, select the test in the STM's PM Session Test (Assurance) list and click Stop Execution. You can also set the Administrative State of a test session to Disabled.

12

To include a PM session in a test policy, perform 89.10 "To configure an STM test policy" (p. 2947) .

13

To include a PM session in a test suite, perform 89.12 "To create an STM test suite" (p. 2951) .

END OF STEPS

92.7 To configure a CFM DMM session OAM diagnostic test from the STM

92.7.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create→Performance Monitoring→CFM DMM Session Test or select an existing CFM DMM session test object and click Properties. The CFM DMM Test Session (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a PM session.

5

To view the bins configured for the test by your choice of PM sessions, perform the following:

1. Click Properties for the selected PM session name. The PM Session (Edit) form opens.
2. Click Properties for the displayed Bin Group Mgr Object ID. The Bin Group (Edit) form opens.
3. Click on the Bins tab.
4. Select the bin you want to view and click Properties. The Bin (Edit) form opens.

6

Save your changes and close the forms.

END OF STEPS

92.8 To configure a CFM SLM session OAM diagnostic test from the STM

92.8.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create→Performance Monitoring→CFM SLM Session Test or select an existing CFM SLM session test object and click Properties. The CFM SLM Test Session (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a PM session.

5

Click on the Loss-events tab to collect TCA event statistics.

1. Select a direction type object and click Properties. The TCA Configuration form opens.
2. Configure the required TCA parameters.

   Note: The Clear threshold should always be lower than the Raise threshold.

6

Save your changes and close the forms.

END OF STEPS

92.9 To configure a CFM LMM session OAM diagnostic test from the STM

92.9.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create→Performance Monitoring→CFM LMM Session Test or select an existing CFM LMM test session object and click Properties. The CFM LMM Test Session (Create|Edit) form opens.

3

Configure the required parameters.

The CFM LMM test must be disabled to make any changes to the session, including the Timing panel options.

When you configure the parameters in the Timing panel, the product of (Interval) x (Frames Per Delta-T) x (Consec Delta-T's) must not exceed 100 seconds.

The Availability parameter governs whether or not the test results will contain the availability information.

4

Select a PM session.

5

Click on the Loss-events tab to collect TCA event statistics.

1. Select a direction type object and click Properties. The TCA Configuration form opens.
2. Configure the required TCA parameters.

   Note: The Clear threshold should always be lower than the Raise threshold.

6

Save your changes and close the forms.

7

To collect LMM statistics, you must set an option to do so on the object with which the MEP is associated. Enable the Enable LMM Session Stats Collection parameter, either on the OAM tab (ETH-CFM sub-tab) or the MEPs tab in the following contexts, as required:

- IES and VPRN L3 access interfaces
- Spoke and Mesh SDP bindings
- IES and VPRN subscriber interfaces

- VLL and VPLS L2 access interfaces
- LAG member ports
- Non-LAG member ports

END OF STEPS

92.10 To configure an MPLS DM session OAM diagnostic test from the STM

92.10.1 Before you begin

Before you can execute an MPLS DM session test, the NE must be enabled to transmit or reflect the test frames. You can enable tests on NEs in the following two ways:

- Manually, by enabling the MPLS-DM Admin Status parameter on the NE properties form, Globals tab, OAM subtab, in the OAM-PM MPLS panel. See 12.5 “To modify NE properties” (p. 343).
- Automatically, by enabling the Force Enable MPLS-DM parameter on the MPLS Session tab during configuration of the PM session for the test; see 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146).

92.10.2 Steps

1

Choose Tools→Service Test Manager (STM) from the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose Create→Performance Monitoring→MPLS DM Session Test or select an existing MPLS DM test session object, and click Properties. The MPLS DM Test Session (Create|Edit) form opens.

3

Configure the required parameters.

4

Select a PM session.

5

To view the bins configured for the test by your choice of PM sessions, perform the following:

1. Click Properties for the selected PM session name. The PM Session (Edit) form opens.
2. Click Properties for the displayed Bin Group Mgr Object ID. The Bin Group (Edit) form opens.

3. Click on the Bins tab.
4. Select the bin you want to view and click Properties. The Bin (Edit) form opens.
5. Close the form after viewing the bin.

6

Save your changes and close the forms.

END OF STEPS

92.11 To configure a TWAMP Light reflector

92.11.1 Steps

A TWAMP Light test needs to be pointed at a reflector on the correct UDP port in order to execute properly and generate valid results. A TWAMP Light reflector can be configured on a base routing instance or on a VPRN site. There can only be one reflector on a particular base routing instance or a VPRN site.

1

Perform one of the following:

a. 27.2 “To configure a routing instance or a VRF instance” (p. 826) for information on creating a TWAMP Light reflector on a base routing instance.

b. 79.76 “To create a TWAMP Light reflector on a VPRN site” (p. 2647) to create a TWAMP Light reflector on a VPRN site.

c. Step 6 in 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) to create a TWAMP Light reflector while creating a PM session. When you specify a source IP address, destination IP address, and destination UDP port on a PM session, the NFM-P automatically creates a new reflector on the appropriate target type at the destination IP address and port if one does not already exist. A prefix matching the source IP address of the PM session is also created if one does not already exist. If you later delete the PM session, the NFM-P will also delete the associated prefix.

2

Refer to Step 6 in 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) to see how the reflector information is incorporated into the PM session test.

END OF STEPS

92.12 To configure a TWAMP Light session OAM diagnostic test from the STM

92.12.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click Create→Performance Monitoring→TWAMP-Light Session Test or select an existing TWAMP Light test session object and click Properties. The TWAMP Light Test Session (Create|Edit) form opens.

3

Configure the required parameters.

The Frames Per Delta-T, Consec Delta-T's, CHLI Threshold, and FLR Threshold parameters are only displayed when the TWAMP Light Statistics Type to Collect parameter is set to either Loss or Delay And Loss.

The Consecutive High Loss Interval (CHLI) Threshold parameter must be set to a value less than the Consec Delta-T's parameter value.

4

Select a PM session.

5

To view the bins configured for the test by your choice of PM sessions, perform the following:

1. Click Properties for the selected PM session name. The PM Session (Edit) form opens.
2. Click Properties for the displayed Bin Group Mgr Object ID. The Bin Group (Edit) form opens.
3. Click on the Bins tab.
4. Select the bin you want to view and click Properties. The Bin (Edit) form opens.
5. Close the form after viewing the bin.

6

Click on the Loss-events tab to collect TCA event statistics.

1. Select a direction type object and click Properties. The TCA Configuration form opens.
2. Configure the required TCA parameters.

7

Save your changes and close the forms.

END OF STEPS

92.13 To configure a TCC OAM diagnostic test from the STM

92.13.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Click the Create→Performance Monitoring→TCC Test or select an existing TCC test object and click Properties. The TCC Test (Create|Edit) form opens.

3

Configure the required parameters.

The Always configure TWAMP Server parameter defines how the NFM-P can modify the operation and configuration of a TWAMP server configured on a managed SR NE. When enabled, on execution of a TCC test, the NFM-P queries the prefix list for the TWAMP server. If no current prefix exists for the server and TWAMP test session client IP addresses, the NFM-P operationally disables the TWAMP server, add the required prefixes, and then re-enables the server.

4

Select a TWAMP server IP address in the TWAMP Server panel.

The displayed list shows system IP addresses for managed NEs. If NEs cannot be managed using in-band communication (for example, when the NFM-P has no route to the system IP address), the SR management interface must be manually entered as the TWAMP Server IP Address.

5

Click Apply.

6

Click on the PM Session tab.

7

Select a PM session or click Create to configure a PM session specifically associated to the TCC test. See 92.6 “To configure a PM session OAM diagnostic test from the STM” (p. 3146) for information on configuring PM sessions and bin groups.

When the test client and TWAMP server NEs are specified, the NFM-P automatically populates the PM session's IP session attributes. Test session IP source and destination addresses default to their respective system interface addresses. The Source UDP port is also chosen from the first one available from the range 64374-64383. If required, the addresses may be cleared and selected from a filtered list of available interfaces on the respective NEs. TWAMP Light tests may also be configured to use specific interfaces on egress, define a static next-hop IP address, or manipulate the IP TOS forwarding class, IP profile and IP TTL fields for TWAMP test packets.

8

Save your changes and close the form.

END OF STEPS

92.14 To collect and view PM statistics from a test form

92.14.1 Steps

You can generate and view PM statistics directly from the test form.

1

Click Tools→Service Test Manager (STM) on theNFM-P main menu. The Service Test Manager (STM) form opens.

2

Select a PM session test and click Properties. The test Properties form opens.

3

Click on the Statistics tab.

4

Select the required statistics record type from the object drop-down menu.

A list of all statistics record types for each PM session test type is provided in 92.15 "To view PM test results in the STM" (p. 3157) .

5

Configure the required time interval from the object menu.

6

Click Collect or Collect All to collect the associated statistics. Clicking Collect All will cause the statistics from all applicable record types to be collected. A list of statistics records is displayed.

7

Select a record entry and click Properties. The Statistics Record form opens.

8

Review the provided information on the General and Bin Stats tabs. The Bin Stats tab is only displayed for applicable record types.

9

Close the Statistics Record form.

END OF STEPS

92.15 To view PM test results in the STM

92.15.1 Steps

1

Click Tools→Service Test Manager (STM) on the NFM-P main menu. The Service Test Manager (STM) form opens.

2

Choose the PM Session (Assurance) entry from the object drop-down menu.

3

Click on the Search button. The list of PM sessions is displayed.

4

Select a PM session entry and click Properties. The PM Session (Edit) form opens.

5

Click on the Components tab.

6

Right-click on the required test under Test Sessions and choose Properties from the contextual menu. The test form opens.

You can also view the statistics for individual bin types for DMM and TWAMP Light tests from this point. Right click on the bin you want to examine under the required test, then select Properties for that bin. The bin statistics can be viewed on the Properties form Statistics tab.

7

Click on the Statistics tab. The statistics records for the selected test are displayed in the list. Statistics utilizing an accounting policy include “Accounting” in the record name (as opposed to SNMP-based statistics).

For CFM DMM test statistics, the record types include:

- CFM DMM Bin Accounting Stats (SAS PM)

- CFM DMM Bin Stats (Ethernet OAM)
- CFM DMM Session Accounting Stats (SAS PM)
- CFM DMM Session Stats (Ethernet OAM)

  These statistics are duplicated in two rows except the monitored object name is changed. For example, one row shows the data and the session name as monitored object and the other row shows the same data but the session bin type as the monitored object.
- Delay TCA Stats (Assurance)
- PM Session Base Stats (Assurance)

For accounting-based statistics, the DMM base statistics data are incorporated into the DMM Bin Stats or DMM Session Stats fields.

For CFM SLM test statistics, the record types include:

- CFM SLM Accounting Stats (SAS PM)
- CFM SLM Session Stats (Ethernet OAM)
- Loss TCA Aggregate Stats (Assurance)
- Loss TCA FWBW Stats (Assurance)
- PM Session Base Stats (Assurance)

Data from SNMP-based SLM base statistics are incorporated into the SLM accounting-based statistics.

For CFM LMM test statistics, the record types include:

- CFM LMM Accounting Stats (SAS PM)
- CFM LMM Session Stats (Ethernet OAM)
- Loss TCA FWBW Stats (Assurance)
- PM Session Base Stats (Assurance)

Data from SNMP-based LMM base statistics are incorporated into the LMM accounting-based statistics.

For TWAMP Light test statistics, the record types include:

- Delay TCA Stats (Assurance)
- Loss TCA Aggregate Stats (Assurance)
- Loss TCA FWBW Stats (Assurance)
- PM Session Base Stats (Assurance)
- TWAMP Light Bin Accounting Stats (SAS PM)
- TWAMP Light Bin Stats (Assurance)
- TWAMP Light Loss Accounting Stats (SAS PM)
- TWAMP Light Loss Session Stats (Assurance)
- TWAMP Light Session Accounting Stats (SAS PM)
- TWAMP Light Session Stats (Assurance)

For accounting-based statistics, TWAMP Light base statistics data are incorporated into the TWAMP Light Bin or TWAMP Light Session Stats fields.

8

Select a record entry and click Properties. The Statistics Record form opens.

9

Review the information on the General and related tabs, then close the forms.

END OF STEPS

92.16 To view PM test statistics in the Statistics Manager

92.16.1 Steps

You can review the OAM PM test results as statistics using the Statistics Manager. You can also use the Statistics Manager plotting function to obtain a graphical representation of the data.

1

Click Tools→Statistics→Statistics Manager in the NFM-P main menu. The Statistics Manager opens.

2

Configure the Statistics Type parameter for the required information type.

3

Select the required statistics type from the menu.

The collected bin and session statistics records for each test type are grouped under the category PM Stats (Assurance). PM Session Base statistics can also be accessed by selecting the PM Session Base Stats (Assurance) category.

4

Refer to the *NSP NFM-P Statistics Management Guide* for the required procedures on viewing statistics and also graphing statistics using the plotter function. Both historical and real-time plots are supported in OAM PM.

You can also access the plotting function for specific tests by clicking the Plotter button on the Statistics tab of a test Properties form. This capability is also available within a PM session Property form by right-clicking a specific test on the Components tab.

Additionally, a summary of the rate of PM session DB record collection is viewable via Tools -> Statistics -> Server Performance Manager (OAM PM Event) NFM-P Performance Statistics.

The tracked attributes include: time captured, periodic time, monitored object name, files received, files received periodic, results processed, and results processed periodic.

END OF STEPS

92.17 To view OAM PM Event server performance statistics

92.17.1 Steps

1

Click Tools→Statistics→Server Performance Statistics on the NFM-P main menu. The Server Performance Statistics form opens.

2

Select OAM PM Event (NFM-P Performance Statistics) from the menu.

The tracked attributes for each OAM PM event are displayed. They include: time captured, periodic time, monitored object name, files received, files received periodic, results processed, and results processed periodic.

END OF STEPS

93 Mirror services

93.1 Mirror service overview

93.1.1 Traffic packet mirroring

CAUTION

Service Disruption

Service mirroring can affect performance.

Service mirroring can affect performance across the network and in the source and destination devices, so must be planned accordingly.

The NFM-P GUI implementation of service mirroring provides mirroring of service traffic packets from any service type.

In a mirror service, packets from one or more sources are forwarded to their normal destinations and a copy of the entire packet, or a specified portion of the packet, is sent to the mirror destination. The mirrored packet can be viewed using a packet-decoding device, typically called a sniffer, that is attached to the destination port. The NFM-P does not limit the number of destination and source sites added under a mirror service. The mirrored packets are transported unidirectionally through the core network using IP or MPLS tunneling.

With pseudo-wire redundancy support, an ICB can be enabled in the mirror service spoke and remote source, which can provide bidirectional service that enables support for active and standby PE redundancy. An endpoint can be used to group the redundant objects, which may be of mirror SDP bindings or SDP and SAP. In the mirror map view, the color for the active and backup states of the redundant mirror SDP differ.

Service mirroring can be used to do the following:

- Troubleshoot problems with customer packet delivery and content.
- Help service providers meet regulations by providing itemized call records and wiretaps, as authorized by investigative authorities.
- Simplify the complex traffic-analysis networks that are often implemented as overlays to the customer-facing network.

93.1.2 Configuration methods

The NFM-P supports end-to-end mirror service configuration using the following methods:

- Tabbed configuration forms with an embedded navigation tree. The navigation tree provides a logical view of the service and acts as a configuration interface.
- Preconfigured template. A user who has the Administrator scope of command role, or the Mirror Service Management and Template Script Management roles, can create a mirror service template.

93.1.3 Operational status

The mirror service operational status is aggregated based on the status of each site.

- Aggregate status is down if all destination sites are down.
- Aggregate status is up if one of the redundant destination site is up.
- Aggregate status is down if all source sites are down.
- Aggregate status is up if one of the redundant source sites is up.
- Aggregate status is unknown if no sites are added to the service.

93.1.4 Implementation considerations

Consider the following before you implement a mirror service:

- You must be assigned the administrator or mirror service management scope of command role to create or modify a mirror service, or to view any mirror-related objects in the NFM-P.
- The default customer is automatically associated with a mirror service. You cannot associate a different customer with a mirror service.
- You can only configure one endpoint per mirror site.
- You can configure the endpoint in the destination site with SAP and SDP with ICB.
- An endpoint cannot have more than one SAP.
- You can create a mirror SDP under the endpoint in the destination site.
- You can create up to four mirror SDPs under the endpoint in the source site. One mirror SDP must have ICB. Alternatively, you can create one mirror SDP under the source site.
- When the mirror SDP binding is under an endpoint, the STP cannot be enabled on the spoke.
- Mirror sites with valid SAPs are considered as destination sites during discovery.
- You can configure a site as a destination without a SAP in the mirror service. During re-synchronization, the site type remains destination.
- Auto-SDP creation is restricted to one destination.
- You can change the redundancy setting of a mirror SDP binding by deleting and adding the mirror SDP binding on an endpoint.
- You cannot delete an MC-LAG SAP if an ICB is on the same endpoint.
- You cannot have a SAP with a non-ICB mirror spoke on the same endpoint.
- The destination of a mirror service must be an L2 SAP.
- You can mirror the ingress or egress traffic on SAPs and ports.
- You can mirror the ingress or egress traffic that is associated with one or more subscriber hosts.
- You can specify match criteria, such as IP addresses, MAC addresses, or MPLS ingress labels, to filter the mirrored traffic.
- You can configure forwarding classes and profiles for mirror service traffic on supported devices.
- Mirror service IDs are obtained from the same pool of IDs that is used by other services. When you manually assign an ID value, ensure that you do not assign an ID that belongs to another service.

- Use the packet-slicing option to copy a specific packet size from each frame. This option is useful for monitoring network usage without copying the customer data. It also limits the amount of mirrored traffic that travels through the core network.
- When the mirror destination is not on the same NE as the mirror source, a mirror service requires a service tunnel between the source and destination NEs.
- The NFM-P can automatically create a service tunnel between the source and destination sites.

  The following conditions must be met.
  - Automatic mesh SDP binding creation is enabled.
  - GRE or LDP is the transport type.
  - No other service tunnel is available between the source and destination sites.
- The mirror service source and destination encapsulation types must be the same.
- After an NE reboot or CPM activity switch, the debug configuration file for a mirror service is not by default reloaded on the NE. The NFM-P raises an alarm when this occurs. To ensure that the NE reloads the debug configuration file, you must specify the location of the file in the base NFM-P configuration. See the procedure to enable debug configuration file reloading on an NE for mirror services in the *NSP System Administrator Guide* for more information.
- When multiple mirror services reference the same packet, for example, from a SAP and from a port, the packet is mirrored only once.

  Apply these criteria in the following non-configurable order:
  - MAC or IP filters
  - MPLS ingress label
  - SAP
  - port
- For IP-only mirroring, specific considerations apply.

  Apply the following guidelines:
  - IP-only mirroring requires chassis mode C or D to be enabled.
  - By using “IP Only” as the encapsulation type, users can specify that only the IP packet is mirrored, without its original ATM/FR/POS/Ethernet DLC header.
  - When IP-only is configured on the source site, users can configure various mirroring sources, including subscribers and SAPs from Ipipe, IES, VPRN, VPLS, and MVPLS services. However, source ports and VLL services including Apipe, Epipe, and Fpipe cannot be configured.
  - For local IP-only mirroring, the source and destination MAC addresses must be configured.
  - For remote IP-only mirroring in a VPRN service, an IP mirror interface on the VPRN destination NE is used to receive the IP mirror packets and route them to the appropriate sniffer. If there are multiple sniffers connected to the VPRN using L3 interfaces, you must configure the routing policy to have IP mirror packets routed to the correct sniffer.

93.2 Sample mirror service

93.2.1 Sample mirror service configuration

The following figure shows a sample mirror service configuration.

Figure 93-1 Sample mirror service configuration

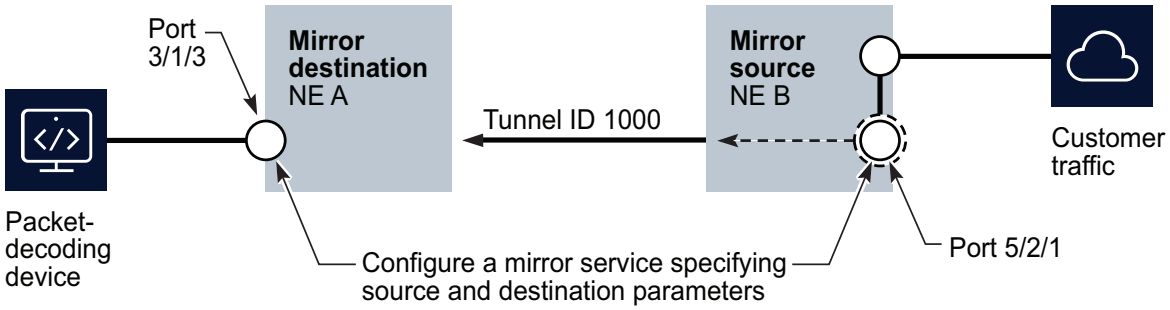

In this sample configuration, Router B is the mirror source that carries the customer packets. Router A is the mirror destination. The ingress and egress traffic on Port 5/2/1 is mirrored on the destination, Port 3/1/3. The following table lists the high-level tasks to configure the sample mirror service.

Table 93-1 Sample mirror service configuration

Task	Description
1. Connect the packet sniffer to the mirror destination.	The packet sniffer is attached to Router A, Port 3/1/3.
2. Configure the mirror destination parameters.	Port 3/1/3 on Router A is specified as the mirror destination. Specify the Tunnel ID 1000 as the transport tunnel to the mirror destination. All the parameters required to configure the type of mirroring, for example, slicing and mirror classification, are specified in the destination parameters.
3. Specify the source entity that is to be mirrored.	The egress and ingress traffic on Port 5/2/1 is to be mirrored.

93.3 Workflow to configure a mirror service

93.3.1 Prerequisites

The workflow lists the high-level steps required to configure a mirror service.

The workflow assumes the following:

- A group or a customer with the required user access privileges has been configured.
- The IP or IP/MPLS core network exists.
- Any required service tunnels are created including the static, dynamic or SR-TE LSP required to create the service tunnel; see 33.9 “To create an IP/MPLS service tunnel” (p. 1190) .
- The access ports for the service are created; see Chapter 16, “Port and channel object configuration” for more information.
- Any required pre-defined routing, QoS, scheduling, filter, accounting, and time of day suite

policies are created; see Chapter 49, “Policies overview” for more information. You do not have to create pre-defined policies if policies are created on a per-service basis.

- Any required MP-BGP for PE-to-PE routing is configured; see Chapter 28, “Routing protocol configuration” for more information about protocol configuration.
- The network services the mirror service will use are created.
- A packet-sniffing device at the L2 SAP that is the destination of the mirror service is configured.
- The SAP with the attached packet-sniffing device as the mirror destination is specified.
- The source of the packets to be mirrored is configured.

Note: Note: The functionalities described in these procedures are device-dependent. Not all functionalities described in these procedures are supported by all devices.

93.3.2 Stages

1

Create the mirror service. See 93.4 “To create a mirror service” (p. 3166) .

1. Define the general properties for the mirror service.
2. Specify the services for inclusion in the mirror service. You can specify multiple services in one operation.
3. Enable one or more routing protocols for the mirror site.
4. Create a destination site on the mirror service. See 93.5 “To create a destination site on a mirror service” (p. 3167) .
5. Create a source site on the mirror service, if required. See 93.6 “To create a source site on a mirror service” (p. 3169) .
6. Specify an endpoint on the site for redundancy, if required. See 93.9 “To create an endpoint for redundancy support on a mirror site” (p. 3173) .
7. Create an L2 access interface on the destination site, if required. See 93.10 “To create an L2 access interface on a destination site” (p. 3174) .
8. Specify one or more SAPs on the site as a mirror source, if required. See 93.11 “To specify a SAP on a mirror site as a mirror source” (p. 3177) .
9. Specify one or more ports on the site as a mirror source, if required. See 93.12 “To specify a port on a mirror site as a mirror source” (p. 3177) .
10. Specify one or more source IP filters, if required. See 93.13 “To specify a source IP filter entry as a mirror source” (p. 3178) .
11. Specify one or more source IPv6 filters, if required. See 93.14 “To specify a source IPv6 filter entry as a mirror source” (p. 3179) .
12. Specify one or more source MAC filters, if required. See 93.15 “To specify a source MAC filter entry for a mirror site” (p. 3180) .
13. Specify one or more source MPLS ingress labels, if required. See 93.18 “To specify an MPLS ingress label as a mirror source” (p. 3182) .
14. Turn up the mirror service.

2

As required, view mirror service information:

a. View the service topology map associated with a Mirror Service. See 93.20 “To view the service topology associated with a mirror service” (p. 3183) .

b. View the mirror service operational status. See 93.21 “To view mirror service operational status” (p. 3184) .

3

As required, run an OAM validation test for a mirror service. See 93.23 “To run an OAM validation test for a mirror service” (p. 3186) .

93.4 To create a mirror service

93.4.1 Before you begin

If you are configuring mirror services using dot1q or QinQ SAPs on a 7210 SAS (except for a 7210 SAS-D or 7210 SAS-K), then you must first assign a no-service mirror port. See 93.22 “To configure a mirror port on a 7210 SAS” (p. 3185) .

93.4.2 Steps

1

Choose Create→Service→Mirror from the NFM-P main menu. The Mirror Service (Create) form opens.

2

Configure the required general parameters.

The Service ID and SVC Mgr Service ID parameters are configurable when the Auto-Assign ID parameter is disabled.

3

If you enabled the Automatic SDP Binding Creation parameter, select a tunnel selection profile in the Auto SDP Binding Creation panel.

4

To configure a mirror source on a destination site, perform 93.5 “To create a destination site on a mirror service” (p. 3167) .

5

To configure a mirror source on a site other than the destination site, perform Step 3 to Step 12 of 93.6 “To create a source site on a mirror service” (p. 3169) . Otherwise, go to Step 6 .

6

Save the changes and close the forms.

END OF STEPS

93.5 To create a destination site on a mirror service

93.5.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

On the service navigation tree, right-click on Destination Sites and choose Create Destination Site to select an available site. The Mirror Site (Create) form opens.

4

Configure the required general parameters.

For mirror destination sites on network-enabled 7210 SAS NEs (except the 7210 SAS-K12 and 7210 SAS-K30 ETR), when you configure a mirror SDP as the source, you must set the Source Type parameter to Remote or Both.

5

Configure the Use Global Sampling Rate parameter to use the configured global rate. See 12.75 “To configure the global sampling rate on an NE” (p. 403) .

6

Click on the Mirroring Configuration tab and configure the required parameters.

Choose the default of 0 for the Slice Size parameter when you configure a destination mirror site, unless slicing at the destination is necessary. Packet slicing reduces the amount of mirrored traffic that traverses the network.

The Slice Size parameter is not displayed if the Encapsulation Type parameter is set to IP Only.

The Port Id Mirroring and remote source are mutually exclusive. If Enable Port ID Mirroring is set to true on the destination site, then a remote source site cannot be added to the service. If a remote source site is already on the service, then Enable Port ID Mirroring cannot be set to true. Also, if remote mirroring is used, then Enable Port ID Mirroring must be configured on all source sites.

The Enable Port ID Mirroring parameter is configurable only if the Encapsulation Type parameter is set to PPP.

For 7210 SAS NEs (except the 7210 SAS-K12 and 7210 SAS-K30 ETR), you can configure the Remote Source Capable parameter only when the Source Type parameter is set to Remote or Both in Step 4 .

7

Configure the remote source far end.

Note: The Remote Source Capable parameter configured in Step 6 must be set to True when you configure a remote source.

Perform the following:

1. Click on the Remote Source tab.
2. Click Create. The Remote Source (Create) form appears.
3. Configure the required parameters.
4. Save the changes and close the form.

8

Configure a remote source mirror SDP binding.

Note: The Remote Source Capable parameter configured in Step 6 must be set to True when you configure a remote source.
Remote source SDP bindings support only L2TPv3 or MPLS-TP tunnels. A remote source SDP binding is mutually exclusive to far end instances.

Perform the following:

1. Click on the Mirror SDP Bindings tab.
2. Click Create. The Mirror SDP Binding (Create) form opens.
3. If you are configuring an L2TPv3 tunnel for the SDP binding, go to 4 . Otherwise, configure the required parameters.
4. Specify a transport tunnel for the mirror SDP binding.

   First configure an MPLS-TP transport tunnel.
   - Select an MPLS-TP service tunnel for the spoke SDP binding in the Tunnel panel.
   - Click on the Pseudowire OAM tab and configure the Control Word parameter. If you are creating a mirror SDP binding using an MPLS-TP service tunnel for pseudowire static configuration, you must set the Control Word parameter to Preferred. See 93.8 “To configure an MPLS-TP static pseudowire on a mirror SDP binding” (p. 3172) .

   Then configure an L2TPv3 tunnel.
   - Select an L2TPv3 service tunnel for the spoke SDP binding in the Tunnel panel.
   - Configure the Ingress Cookie parameter. The Ingress Cookie parameter must match the Egress Cookie parameter that you configure on the SDP binding using L2TPv3 tunnel on the mirror source site in 93.6 “To create a source site on a mirror service” (p. 3169) .
5. Save the changes and close the form.

9

Save the changes and close the forms.

END OF STEPS

93.6 To create a source site on a mirror service

93.6.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

On the service navigation tree, right-click on Source Sites and choose Create Source Site to select an available site. The Mirror Site (Create) form opens.

4

Configure the required general parameters.

5

Click on the Mirroring Configuration tab.

6

Configure the required parameters.

The Slice Size parameter is not configurable if the Encapsulation Type parameter is set to IP Only in Step 4 .

The Enable Port ID Mirroring parameter is configurable only if you set the Encapsulation Type parameter to PPP in Step 4 .

The Port Id Mirroring and remote source are mutually exclusive. If Enable Port ID Mirroring is set to true on the destination site, then a remote source site cannot be added to the service. On the other hand, if a remote source site is already on the service, then Enable Port ID Mirroring cannot be set to true. Also, if remote mirror is used, then Enable Port ID Mirroring must be configured on all source sites.

7

Save the changes and close the form.

8

If you did not enable the Automatic SDP Binding Creation parameter in 93.4 “To create a mirror service” (p. 3166), then choose an SDP binding for the mirror source.

1. In the service navigation tree, expand the site object that you created.
2. Select Mirror SDP Binding under the source site in the service navigation tree, right-click and choose Create Mirror SDP Binding. The Mirror SDP Binding (Create) form opens.

   If you enabled the Automatic SDP Binding Creation parameter in 93.4 “To create a mirror service” (p. 3166) , go to Step 12 .

9

Specify a destination NE for the mirror SDP binding by selecting from a list of managed NEs for the Tunnel Termination Site parameter.

If the destination NE is not managed by the NFM-P, specify the system ID of the destination NE for the Tunnel Termination Site parameter.

Note: For a 7210 SAS destination site in a mirror service, you can either specify a mirror SDP binding or create a SAP, but not both.

10

If you are configuring an L2TPv3 tunnel for the SDP binding, go to Step 11 . Otherwise, configure the required parameters.

11

To specify a transport tunnel for the mirror SDP binding:

a. Let the NFM-P configure the transport tunnel automatically.
   1. Enable the Auto-Select Transport Tunnel parameter.
   2. Configure the Tunnel Auto-Selection Transport Preference parameter.

b. Configure the transport tunnel manually; select a tunnel in the Tunnel panel.

c. Configure an MPLS-TP transport tunnel.
   1. Select an MPLS-TP service tunnel for the spoke SDP binding in the Tunnel panel.
   2. Click on the Pseudowire OAM tab and configure the Control Word parameter.

      If you are creating a mirror SDP binding using an MPLS-TP service tunnel for pseudowire static configuration, then you must set the Control Word parameter to Preferred. See 93.8 “To configure an MPLS-TP static pseudowire on a mirror SDP binding” (p. 3172) .
   3. Save the changes and close the forms.

d. Configure an L2TPv3 tunnel.
   1. Select an L2TPv3 service tunnel for the spoke SDP binding in the Tunnel panel.
   2. Configure the Egress Cookie parameter. The Egress Cookie parameter must match the

Ingress Cookie parameter that you configured on the SDP binding using L2TPv3 tunnel on the mirror destination site in 93.5 “To create a destination site on a mirror service” (p. 3167) .

12

Save the changes and close the forms.

END OF STEPS

93.7 To configure a PCAP session for a mirror site

93.7.1 Steps

Mirrored packets can be captured in a PCAP file and viewed offline. A single capture provides all relevant protocol packets. Byte level details are also provided for each captured packet.

1

Create the mirror service. See 93.4 “To create a mirror service” (p. 3166) .

2

On the mirror service, create a source site. See 93.6 “To create a source site on a mirror service” (p. 3169) .

3

Select the source site and click Properties. The Mirror Service (Edit) form opens.

4

Click on the PCAP Sessions tab and click Create. The PCAP Session (Create) form opens.

5

Configure the Session Name, Capture Count, and File URL parameters.

The NFM-P can capture up to 250 packets from a mirror source.

You must provide a valid IP address for the generated PCAP file. If you want to use a new URL, you do not need to recreate the PCAP session. Delete the URL, apply any changes, and re-enter the new URL.

6

Click Apply. The new PcapSessions appears under the Mirror Service - Source Sites in the navigation tree.

7

On the MirrorPCAP session tab, click Start Pcap Session. The Session State parameter indicates the capture state of the PCAP session.

If you want to stop the PCAP session, you can use the Stop pcap capture option.

If you stopped and started a PCAP session without changing the file URL or PCAP file name, the previously captured information will be overwritten.

8

View the details of the PCAP file from the URL you specified.

Statistics for the PCAP session are available on the Statistics tab.

Scheduled statistics are not supported. PCAP session statistics are valid only when the PCAP session is in progress.

9

Save the changes and close the forms.

END OF STEPS

93.8 To configure an MPLS-TP static pseudowire on a mirror SDP binding

93.8.1 Steps

Perform this procedure to create an MPLS-TP static pseudowire on the mirror SDP binding. An MPLS-TP service tunnel must be used in the mirror SDP binding, and the Control Word parameter for pseudowire OAM must be set to Preferred in Step 11 of 93.6 “To create a source site on a mirror service” (p. 3169) .

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

On the navigation tree, expand the source site on which you want to configure the mirror SDP binding, expand Mirror SDP Bindings and click on the mirror SDP binding on which you want to configure the static MPLS-TP pseudowire. The Mirror SDP Binding (Edit) form opens.

4

Click on the Control Channel tab.

5

Configure the required parameters.

6

Click on the Static PW tab and click Create. The PW Path ID (Create) form opens.

7

Configure the Path AGI parameter.

8

Configure the parameters in the Source Attachment Individual Identifier panel.

9

Configure the parameters in the Target Attachment Individual Identifier panel.

10

Save the changes and close the forms.

END OF STEPS

93.9 To create an endpoint for redundancy support on a mirror site

93.9.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following.

a. To create an endpoint on a destination site, expand the Destination Sites object to view the configured destination sites, and expand the destination site object on which you are creating the endpoint. Right-click on Endpoints and choose Create Endpoints. The Endpoint (Create) form opens.

b. To create an endpoint on a source site, expand the Source Sites object to view the configured source sites, and expand the source site object on which you are creating the endpoint. Right-click on Endpoints and choose Create Endpoints. The Endpoint (Create) form opens.

4

Configure the required general parameters.

5

Click OK.

For the destination site, the L2 Access Interface (Test Equipment Interface) and Mirror SDP Binding objects appear under the Endpoints object.

For the source site, the Mirror SDP Binding object now appears under the Endpoints object.

6

To create an L2 access interface as the mirror destination, expand Endpoints, right-click on L2 Access Interface (Test Equipment Interface) and choose Create L2 Access Interface. The L2 Access Interface (Create) form opens.

Note: You cannot create an L2 access interface on the endpoint if you create an L2 access interface in 93.10 “To create an L2 access interface on a destination site” (p. 3174) without selecting an endpoint for redundancy.

7

Click on the Port tab and select a port for the L2 access interface.

Note: Only ports in access or hybrid mode are listed. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click Search.

8

Configure the required parameters.

9

Save the changes and close the form.

10

To create a mirror SDP binding on the endpoint, expand Endpoints, right-click on Mirror SDP Binding and choose Create Mirror SDP Binding. The Mirror SDP Binding (Create) form opens.

11

Configure the required general parameters.

12

Save the changes and close the forms.

END OF STEPS

93.10 To create an L2 access interface on a destination site

Note: You cannot create an L2 access interface on the site if you create an L2 access interface for redundancy in 93.9 “To create an endpoint for redundancy support on a mirror site” (p. 3173) .

For a 7210 SAS destination site in a mirror service, you can either create a SAP or specify a mirror SDP binding, but not both.

93.10.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

On the service navigation tree, expand Destination Sites→*destination_site*, right-click on L2 Access Interface (Test Equipment Interface), and choose Create L2 Access Interface. The L2 Access Interface (Create) form opens.

4

To configure an endpoint for redundancy, click on the Select button in the Redundancy panel to select the desired endpoint.

Note: If you select an endpoint for redundancy, the L2 access interface appears under the Endpoints object on the service navigation tree.

5

Click on the Port tab.

6

Select a port.

Note: The form lists only ports in access or hybrid mode. To choose a port that is not listed, you must set the Mode parameter for the port to Access or Hybrid. After you do this, the port is listed when you click on the Search button.
The form lists only ports that have the encapsulation type that you specified in the destination site creation, in 93.5 “To create a destination site on a mirror service” (p. 3167)

7

Configure the required parameters.

The Auto-Assign ID parameter is configurable if the port uses dot1q encapsulation. When the parameter is enabled, the NFM-P automatically configures the Outer Encapsulation Value parameter using the lowest unassigned value.

Note: You can set the Auto-Assign ID parameter to be the default parameter for dot1q encapsulation by enabling the Access Interface Encap Value (Dot1q only) parameter on the User Preferences form.

The Inner Encapsulation Value is configurable only when the port is an Ethernet or frame relay port with QinQ encapsulation.

8

To assign an egress QoS policy to the interface:

1. Click on the QoS tab.
2. Configure the Egress Mark QinQ Top Bits Only parameter.
3. Select an egress policy in the Egress Policy panel.

9

To assign a time of day suite to the interface:

1. Click on the TOD Suite tab.
2. Select a time of day suite.

10

To configure scheduling:

1. Click on the Schedulers tab.
2. Configure the Aggregate Rate Limit (kbps) parameter.

11

To configure local mirroring and if you chose IP Only as the encapsulation type in Step 4 of 93.5 “To create a destination site on a mirror service” (p. 3167) :

1. Click on the IP Mirror MAC Addresses tab.
2. Configure the required parameters.

   The Source and Destination MAC addresses on L2 Access Interface must be both null or neither null. Both null means that there are no source or destination MAC addresses configured on the interface.

   Note:

   If you want to configure remote mirroring in a VPRN service, you must create an IP Mirror Interface in that service. See Chapter 79, “VPRN service management” for more information.

12

To assign an ANCP policy to the interface:

1. Click on the ANCP Static Map tab.
2. Click Create. The ANCP Static Map (Create) form opens.
3. Configure the ANCP String parameter.
4. Select an ANCP Policy.
5. Save the changes and close the form.

13

Save the changes and close the forms.

END OF STEPS

93.11 To specify a SAP on a mirror site as a mirror source

93.11.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

In the service navigation tree, expand Site, right-click on the Source SAPs object and choose Create Source Interface. The Source Interface (Create) form opens.

4

Select an access interface to associate with the source interface.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

93.12 To specify a port on a mirror site as a mirror source

93.12.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following.

a. To create a source port on a destination site, expand Destination Sites→*destination_site*, right-click on Source Ports and choose Create Source Port. the Source Port (Create) form opens.

b. To create a source port on a source site, expand Source Sites→*source_site*, right-click on Source Ports and choose Create Source Port. The Source Port (Create) form opens.

4

Select a port.

A source port can be one of the following, depending on the NE:

- physical port
- channel
- LAG
- bundle
- CCAG

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

93.13 To specify a source IP filter entry as a mirror source

93.13.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following.

a. To specify an IP filter entry on a destination site, expand Destination Sites→*destination_site*, right-click on Source IP Filters and choose Create Source IP Filter. The Source IP Filter (Create) form opens.

b. To specify an IP filter entry on a source site, expand the Source Sites→*source_site*, right-click on Source IP Filters and choose Create Source IP Filter. The Source IP Filter (Create) form opens.

4

Select an IP filter entry.

5

Save the changes and close the forms.

END OF STEPS

93.14 To specify a source IPv6 filter entry as a mirror source

93.14.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following.

a. To specify an IPv6 filter entry on a destination site, expand Destination Sites→*destination_site*, right-click on Source IPv6 Filters and choose Create Source IPv6 Filter. The Source IPv6 Filter (Create) form opens.

b. To specify an IPv6 filter entry on a source site, expand the Source Sites→*source_site*, right-click on Source IPv6 Filters and choose Create Source IPv6 Filter. The Source IPv6 Filter (Create) form opens.

4

Select an IPv6 filter entry.

5

Save the changes and close the forms.

END OF STEPS

93.15 To specify a source MAC filter entry for a mirror site

93.15.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following.

a. To specify a MAC filter entry on a destination site, expand Destination Sites→*destination_site*, right-click on Source MAC Filters and choose Create Source MAC Filter. The Source MAC Filter (Create) form opens.

b. To specify a MAC filter entry on a source site, expand the Source Sites→*source_site*, right-click on Source MAC Filters and choose Create Source MAC Filter. The Source MAC Filter (Create) form opens.

4

Select a MAC filter entry.

5

Save the changes and close the forms.

END OF STEPS

93.16 To specify a source filter as mirror source on 7250 IXR

93.16.1 Steps

A mirror service can contain multiple aggregate filter policies. However, each ACL aggregate filter policy must contain only one filter reference. A mirror service can have multiple aggregate filter policies attached. However, the aggregate filter policy can only have one source filter type attached. See 51.3 “To configure an ACL Aggregate filter policy” (p. 1668) to configure an ACL aggregate filter policy.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

To specify a source filter on a destination site, expand Destination Sites→destination_site, right-click on Source Filters, and choose Create Source Filter. The Source Filter (Create) form opens.

4

From the Source Filter Type drop-down menu, select IP, IPv6, or MAC filter policy type.

5

Select the Aggregate Filter Name.

6

Select the IP/IPv6/MAC Filter Entry parameters.

7

Click on OK to save the changes and close the forms.

END OF STEPS

93.17 To specify a source subscriber as a mirror source

93.17.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

In the service navigation tree, expand Source Sites→*source_sites*, right-click on Source Subscribers and choose Create Source Subscriber. The Source Subscriber (Create) form opens.

4

Configure the required parameters.

5

You can further restrict the mirrored subscriber host traffic associated with the subscriber identification string by specifying SAP or SLA-profile criteria. To specify an SLA profile as a match criterion, go to Step 7 .

Note: You can configure parameters in the SAP/Subscriber Host Match Criteria panel or the SLA Profile Match Criteria panel, but you cannot configure parameters in both panels.

6

Specify a SAP on which to mirror subscriber host traffic.

Note: The form lists only dot1q- or QinQ-encapsulated SAPs that have subscriber management enabled.

Perform the following

1. Select a SAP in the SAP/Subscriber Host Match Criteria panel.
2. Configure the parameters in the SAP/Subscriber Host Match Criteria panel to specify the subscriber host match criteria for the SAP.

   Note:

   The NFM-P does not accept the parameter values unless a SAP is specified.

7

To specify an SLA profile as a match criterion, select an SLA profile in the SLA Profile Match Criteria panel.

8

Save the changes and close the forms.

END OF STEPS

93.18 To specify an MPLS ingress label as a mirror source

93.18.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

In the service navigation tree, expand Source Sites→*source_site*, right-click on Source MPLS Ingress Labels and choose Create Source Ingress Label. The Source Ingress Label (Create) form opens.

4

Configure the Ingress Label parameter.

5

Save the changes and close the forms.

END OF STEPS

93.19 To specify a source VLAN as a mirror source on the 7250 IXR

93.19.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Right-click Source VLAN and click Create Source VLAN. The Source VLAN (Create) form opens.

4

In the Mirrored VLAN, select the port or LAG from the list.

5

Assign a Port ID.

6

Click on OK to save the changes and close the forms.

END OF STEPS

93.20 To view the service topology associated with a mirror service

93.20.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Topology View. A Topology View dialog box appears.

3

Click Yes to proceed. The Service Topology - *Service Name* map opens.

See Chapter 4, “Topology map management” for more information about service topology views.

End of steps

93.21 To view mirror service operational status

93.21.1 Steps

Note: The Aggregated Service Site Operational State and State Cause indicators on the General tab of the mirror service properties form display information about service faults.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror *Service Name* (Edit) form opens.

3

View the Aggregated Service Site Operational State and State Cause indicators. When the Aggregated Operational State is Down or Partially Down, a check mark beside the appropriate State Cause indicator identifies the type of associated service fault.

4

Click on the appropriate tab to view or edit an object that is identified as faulty by a State Cause indicator.

5

Click on the Faults tab to view the alarms for the object. The alarms are grouped into various categories.

6

Close the Mirror Service (Edit) form.

7

Save the changes and close the forms.

End of steps

93.22 To configure a mirror port on a 7210 SAS

93.22.1 No-service ports on the 7210 SAS

For some implementations of mirror services on the 7210 SAS, you must configure a no-service port on the NE. You can select a physical port or a virtual port. Not all chassis types support virtual no-service ports.

If you select a physical port as a no-service port, the port cannot contain any SAPs.

For virtual no-service ports, the number of available ports and the port names vary depending on the chassis type. Virtual ports are not displayed on the navigation tree.

You cannot select the same no-service port for more than one function.

For more information about configuring no-service ports for mirror services, see the 7210 SAS OAM and Diagnostics documentation.

93.22.2 Steps

1

Open the properties form for the 7210 SAS NE that you need to configure.

2

Click on the Globals tab, then click on the Service tab.

3

Select a loopback no-service port. Perform one of the following:

a. Select a physical loopback no-service port.
   1. In the Mirror Port panel, click Clear if required, then click Select. The Select Mirror Port form opens.
   2. Choose a port and click OK.

b. Select a virtual loopback no-service port.
   1. Enable the Use Virtual Mirror Port parameter.
   2. In the Virtual Mirror Port panel, click Clear if required, then click Select. The Select Virtual Mirror Port form opens.
   3. Choose a port and click OK.

      The available virtual ports vary depending on the chassis type and card type.

4

Save your changes and close the forms.

END OF STEPS

93.23 To run an OAM validation test for a mirror service

93.23.1 Before you begin

An OAM validator test suite must be created for the tested entity. See Chapter 89, “Service Test Manager” for more information about how to create an OAM validator test suite.

93.23.2 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a service and click Properties. The Mirror Service (Edit) form opens.

3

Click Validate. If an OAM validator test suite is not associated to the service, a dialog box appears. Perform the following steps:

1. Click OK to associate the service with an existing OAM validator test suite. The Choose Validator Test Suite form appears.
2. Select an OAM validator test suite and click OK. The Choose Validator Test Suite form closes.

4

View the State Cause indicators. When the validation test fails, a check mark appears beside the OAM Validation Failure indicator.

5

Click on the Tests tab.

6

Click on the Tested Entity Result tab.

7

Select an entry and click Properties. The Tested Entity Result form opens and displays information about the validation test.

8

Close the Tested Entity Result form.

9

Close the Mirror Service (Edit) form.

10

Save the changes and close the forms.

END OF STEPS

94 Lawful Intercept

94.1 Overview

94.1.1 Purpose

This chapter describes the NFM-P Lawful Intercept (LI) function, and includes LI configuration information.

94.1.2 Contents

Lawful Intercept overview

94.2 Lawful Intercept concepts

94.2.1 Overview

Lawful Intercept (LI) is a term that describes the interception and monitoring of network subscriber traffic by authorized agencies for law-enforcement purposes. A subscriber whose traffic is intercepted using LI is called a target. The LI target traffic is replicated by a service mirror that uses subscriber information as the match criteria.

94.2.2 Security assurances and restrictions

NFM-P has no native LI functions or legal interception capability, and can only enable, configure, and disable LI functions on NEs that have native LI functions, such as the 7750 SR.

NFM-P acts only as an optional remote LI controller, and has no view of any intercepted traffic; the target address is the only information available to NFM-P.

LI security constraints

LI uses the following constraints to enforce data security.

- LI is not detectable by a target subscriber.
- LI traffic is delivered separately from other network traffic.
- LI alarms are displayed only for authorized LI users.
- LI configuration information is stored in a separate and encrypted file.
- LI uses TLS to secure the required communication channels.

94.2.3 LI requirements

LI functions are managed separately from other NFM-P functions. LI mirroring is a special type of service mirroring that requires the following:

- on each LI NE in a mirror service:
  - SSH user security
  - an LI user profile
  - an NE user account that has LI privileges
- on the NFM-P:
  - TLS on the single-user GUI clients, client delegate servers, and JMS server
  - an NFM-P user account that has LI privileges
  - an assigned Lawful Interception Management scope of command role
  - an LI mediation policy

Note: The NFM-P blocks the association of LI mediation security with any 7750 SR whose software version descriptor includes -NL. An LI configuration on such an NE cannot be completed because the NE does not support LI mediation security.

94.3 LI functional tasks by user type

94.3.1 User authorization level

LI requires a dedicated user authorization level to setup the LI infrastructure on the NFM-P and corresponding NEs to configure and manage LI functions. The following user types are required to perform LI:

- NFM-P admin
- NE admin
- NE LI user
- NFM-P LI user or LI administrator

Note: For security, all LI functions in the NFM-P GUI, for example, menus, forms, parameters, policies, filters, and alarms, are hidden from view for all NFM-P user types except an LI user or LI administrator.

94.3.2 NFM-P admin LI tasks

The NFM-P admin user is responsible for the initial setup of the LI user or LI administrator accounts on the NFM-P. The setup includes the following tasks:

- create a LI user scope of command profile and associate a non-admin user account with the LI profile
- configure an LI user profile that restricts the LI user to LI activities only and does not allow system administrator activities
- create a password for the LI user
- configure CLI and SNMP access for the NE LI user account

The NFM-P admin cannot perform the following LI functions:

- assign LI privileges to a user who is associated with the admin profile
- delete an LI user
- view, create, modify, or delete LI objects
- view or modify LI security mediation policies or LI filter policies
- view LI-related alarms

94.3.3 NE admin tasks

The NE admin is responsible for the creation of an NE LI user profile and NE LI user account using a CLI on the NE that LI will be performed on. This is usually a one-time configuration.

94.3.4 NE LI user tasks

The NE LI user configures the NE LI user security on the NE using a CLI such as changing the password of the NE LI user account so it’s unknown to NE admin and to configure the SNMP data encryption used for the NE. This is usually a one-time configuration.

94.3.5 LI users or LI administrator tasks

CAUTION

Service Disruption

An LI user password cannot be modified unless it is known. When an LI user password is not known, the LI user account is unavailable.

An LI user or LI administrator can perform the following LI related tasks on the NFM-P:

- modify the LI user password
- configure an LI mediation security policy
- configure other LI-specific policy types that are used to filter or block traffic to non-LI users for LI mirror services including:
  - LI MAC filter policy
  - LI IP filter policy
  - LI IPv6 filter policy
  - Block reservation policy
- create, configure, and view LI sources that use IP, MAC, SAP, and subscriber filters
- create the LI mirror service
- view LI-related alarms

An LI administrator is an LI user who has been designated as the lead LI user who would normally create and maintain other LI user accounts once the NFM-P System admin has created the initial LI user account. This role is dependent/optional based on your NOC security requirements.

Note: Viewing or retrieving LI user activity records requires a user account with an assigned Lawful Interception Management scope of command role. The scope is restricted to the records of users in the same LI user group.

The NFM-P does not support LI user to execute CLI commands with LI user privilege on the node by using CLI scripts or by executing XML API API calls such as executeCLI or executeMultiCLI.

The following user account creation conditions apply to LI users and LI user groups.

- To configure LI, you must assign the Lawful Interception Management scope of command role for the user type.
- The Lawful Interception Management role can be assigned to only one scope of command profile.
- A scope of command profile that has an assigned Lawful Interception Management role can include other roles except for the admin role.
- You cannot change the scope of command profile assignment for a user group when the profile includes the Lawful Interception Management role.
- An LI user group must be created as an LI user group; you cannot change a non-LI user group to an LI user group.
- You cannot change an LI user group to a non-LI user group.

- A user account can have the Lawful Interception Management or the admin role, but not both.
- A user who belongs to an LI user group cannot be changed to a non-LI user.
- An LI span of control profile restricts LI user access to specific NEs.

94.4 LI service mirroring

94.4.1 LI target traffic

The LI target traffic is replicated by a service mirror that uses subscriber information as the match criteria. The LI service mirroring has a higher priority than non-LI mirroring. The LI service mirroring must be configured by an LI user with LI privileges and is hidden from users without LI permissions.

94.4.2 LI mirror sources

Only one mirror source, which can contain one or many LI source entries, can be associated with one mirror destination service. LI takes priority over debug mirror sources.

In the configuration of an LI source, when an LI user specifies that an entry must be used as an LI entry, it is hidden from all non-LI users.

94.4.3 LI filter policies

LI filter policies are used to filter out traffic to non-LI users for LI mirror services; these filters can only be viewed and configured by LI users.

Both non-LI filter entries and special purpose LI filter entries can be referenced in an LI mirror source. LI filters are associated with non-LI filters, and entries created in the LI filters are inserted into the associated non-LI filter before the filter is applied. A LI filter block reservation can be configured to reserve a range of entries in the non-LI filter into which the LI entries are inserted. See 94.18 “To configure a Block Reservation policy” (p. 3212) for more information.

LI filter policies include:

- LI MAC filter policy; see 94.15 “To configure an LI MAC filter policy” (p. 3208)
- LI IP filter policy; see 94.16 “To configure an LI IP filter policy” (p. 3209)
- LI IPv6 filter policy; see 94.17 “To configure an LI IPv6 filter policy” (p. 3211)
- Block Reservation filter policy; see 94.18 “To configure a Block Reservation policy” (p. 3212)

If a filter entry is associated with an LI source, the filter and the filter entry cannot be modified or deleted. unless the LI filter lock is configured. The filter lock configuration controls the behaviour of filters when they are used for LI. See 94.19 “To configure the LI filter lock” (p. 3214) for more information.

94.4.4 LI for NAT

LI for NAT mirrors the traffic of configured subscribers to a mirror-destination. All traffic for the specified subscriber, including traffic associated with static port forwarding, is mirrored. See 94.30 “To configure LI on a specific NAT subscriber” (p. 3225) for more information.

Procedures to configure LI

94.5 Workflow to configure LI

94.5.1 Stages

The following is the sequence of high-level activities required to setup the LI infrastructure on both the NFM-P and corresponding NEs, and to configure and manage LI functions.

Configure the LI user on the NFM-P

1

Create an LI user on the NFM-P.

1. Plan the NFM-P LI user account creation according to the requirements for LI user access to other NFM-P functional areas. See “User account and group management” in the *NSP System Administrator Guide* for more information.
2. Create an LI scope of command profile that has an assigned Lawful Interception Management role. See 94.6 “To create an LI scope of command profile on the NFM-P” (p. 3197) .
3. Create an LI user group that is associated with the new scope of command profile. See 94.7 “To create an LI user group on the NFM-P” (p. 3198).
4. Create an LI user account and assign it to the new user group. See 94.8 “To create an LI user on the NFM-P” (p. 3199) .
5. Provide the login credentials for the LI user account to the authorized LI administrator or LI user.
6. The LI user must change their LI user account password so that it is unknown to the NFM-P admin user.

Configure TLS for LI users

2

Assign SNMPv3 access privileges to the LI user group created in Stage 1 to each NE that the LI function will be performed on. You must enable TLS for the user group. See 9.11 “To enable SNMPv3 management of a device” (p. 291) .

Note: You cannot log in to the NFM-P as an LI user unless TLS is enabled for the LI user group.

3

NFM-P GUI access for LI users requires TLS between the NFM-P and the GUI clients. If TLS is disabled on the NFM-P XML API, TLS is also disabled for GUI access.

If TLS is disabled on the NFM-P XML API, enable TLS on the interface. The NFM-P section of the security chapter in the *NSP Installation and Upgrade Guide* describes how to enable TLS for NFM-P XML API clients.

Configure the LI user on the LI NE

4

Create an NE LI user on an NE where LI is being performed.

1. Create an NE LI user profile on the NE using a CLI. See 94.9 “To create an NE LI user profile on an NE using a CLI” (p. 3200) .
2. Create an NE LI user account on the NE that is associated with the LI user profile using a CLI. See 94.10 “To create an NE LI user account on an NE using a CLI” (p. 3201) .
3. Provide the login credentials for the LI NE user account to the NE LI user.
4. Configure NE LI user security on the NE such as the LI NE user account, password, and SNMP data encryption for the NE. See 94.12 “To configure NE LI user security” (p. 3204) .
5. Specify if the NE stores the LI source configuration locally or reconfigures the LI sources after a reboot. LI source configurations are saved on an NE when the polling policy for the NE specifies LI Local Save Allowed. See 8.10 “To configure polling for a 7250 IXR, 7450 ESS, 7705 SAR, 7750 SR, 7950 XRS, VSR, or Wavence SM” (p. 260) .

Configure the LI NE mediation policies and LI filter policies

5

Perform the following steps to configure an LI security mediation policy and the appropriate LI filter policies on the NFM-P.

1. Log out of the NFM-P GUI if you are logged in as a non-LI user and log back in as an LI user.
2. Synchronize the global NE user configuration profile with the local LI NE user configuration profile which is required prior to creating the LI mediation policy. See 94.13 “To synchronize the global NFM-P NE LI user configuration profile with the local NE LI user configuration profile” (p. 3206) .
3. Create an LI mediation security policy that defines the network security model used when creating LI mirror objects. See 94.14 “To configure an LI security mediation policy” (p. 3207) .
4. Create an LI MAC filter policy if you need to define the source and destination LI MAC filter entries used to filter out traffic to non-LI users for LI mirror services. See 94.15 “To configure an LI MAC filter policy” (p. 3208) .
5. Create an LI IP or IPv6 filter policy if you need to define the source and destination LI IP or IPv6 filter entries used to filter out traffic to non-LI users for LI mirror services. See 94.16 “To configure an LI IP filter policy” (p. 3209) and 94.17 “To configure an LI IPv6 filter policy” (p. 3211) .
6. Create a Block Reservation policy if you need to define the block reservation attributes of LI filter entries such as the block start/stop-entries and the block size in the base IPv4/IPv6 ACL filter. See 94.18 “To configure a Block Reservation policy” (p. 3212) .
7. Configure the LI Filter Lock if you need to configure who can modify base IPv4 and MAC filters referenced by an LI source. See 94.19 “To configure the LI filter lock” (p. 3214) .

6

Use the NFM-P to enable LI discovery of an NE. See 94.20 “To enable NE discovery for LI” (p. 3214) .

Configure LI mirror services

7

Perform the following steps to configure LI mirror services. See the “Workflow to create a mirror service” in Chapter 93, “Mirror services” for any non-LI configuration tasks that are required to configure a mirror service.

1. Create a mirror service. See 93.4 “To create a mirror service” (p. 3166) .
2. To specify an LI MAC filter created in Stage 5 5 as the LI source, see 94.23 “To specify an LI MAC filter entry as an LI source” (p. 3218) .
3. To specify an LI IPv4 or IPv6 filter created in Stage 5 6 as the LI source, see 94.25 “To specify an LI IP filter entry as an LI source” (p. 3220) and 94.26 “To specify an LI IPv6 filter entry as an LI source” (p. 3221) .
4. To specify an LI subscriber or subscriber host as the LI source, see 94.28 “To specify an LI subscriber as an LI source” (p. 3223) .
5. To specify an LI WLAN distributed subscriber as the LI source, see 94.29 “To specify an LI WLAN distributed subscriber as an LI source” (p. 3224) .
6. To specify an LI Source Port as the LI source, see 94.24 “To specify an LI Source Port as an LI source” (p. 3219).
7. To specify an LI SAP on a mirror site as the LI source, see 94.27 “To specify an LI SAP as an LI source” (p. 3222) .
8. To configure LI on a specific NAT subscriber, see 94.30 “To configure LI on a specific NAT subscriber” (p. 3225) .
9. To view LI mirrored subscriber host service information configured with a RADIUS server, see 94.31 “To view LI mirrored subscriber hosts configured with a RADIUS server” (p. 3227) .

Monitor LI user and system activity

8

View LI user and system logs to monitor LI activity. See the section on user activity logging in the *NSP System Administrator Guide* for more information about accessing user and system logs.

94.6 To create an LI scope of command profile on the NFM-P

94.6.1 Steps

You require NFM-P admin privileges to perform this procedure.

1

Using an account with an assigned User Management scope of command role, choose Administration→Security→NFM-P User Security from the NFM-P main menu. The NFM-P User Management (Edit) form opens.

2

Click on the Scope of Command tab.

3

Click Create→Profile. The Scope of Command Profile (Create) form opens.

4

Configure the required general parameters.

5

Assign one or more scope of command roles to the profile:

1. Click on the Roles tab.
2. Click Add and select one or more roles to include in the scope of command profile, where one of the roles is the Lawful Interception Management role.

   Note:

   You cannot include the Administrator role in an LI scope of command profile.

   The NFM-P allows the creation of only one scope of command profile that contains the Lawful Interception Management role.

6

Save the changes and close the form.

END OF STEPS

94.7 To create an LI user group on the NFM-P

94.7.1 Steps

You require NFM-P admin privileges to perform this procedure.

1

Using an account with an assigned User Management scope of command role, choose Administration→Security→NFM-P User Security from the NFM-P main menu. The NFM-P User Management (Edit) form opens.

2

Click on the User Groups tab.

3

Click Create. The User Group (Create) form opens.

4

Configure the required general parameters.

5

Select the LI scope of command profile you created in 94.6 “To create an LI scope of command profile on the NFM-P” (p. 3197) in the Scope of Command panel to assign the LI scope of command profile to the user group.

6

To assign a span of control profile to the user group select a profile ID in the Span of Control panel.

7

Save the changes and close the forms.

END OF STEPS

94.8 To create an LI user on the NFM-P

94.8.1 Steps

You require NFM-P admin privileges to perform this procedure.

1

Using an account with an assigned User Management scope of command role, choose Administration→Security→NFM-P User Security from the NFM-P main menu. The NFM-P User Management (Edit) form opens.

2

Click on the Users tab.

3

Click Create. The User (Create) form opens.

4

Configure the required general parameters.

5

Select the LI user group you created in 94.7 “To create an LI user group on the NFM-P” (p. 3198) to assign the LI user to the user group.

6

Configure the parameters in the Password panel.

7

Configure the parameters in the UI Session panel.

8

Configure the parameters in the OSS Session panel.

9

Configure the parameters in the Client IP Address panel.

10

Save the changes and close the forms.

END OF STEPS

94.9 To create an NE LI user profile on an NE using a CLI

94.9.1 Steps

You require NFM-P admin privileges to perform this procedure.

1

Open an SSH session on the NE as described in Chapter 10, “Device CLI sessions” .

2

Enter the following commands in sequence at the prompt to create an NE LI user profile. The following is a sample configuration:

```
configure system security profile li-prof ↲
default-action deny-all ↲
li ↲
entry 10 ↲
match "configure li" ↲
action permit ↲
exit ↲
entry 20 ↲
match "exit" ↲
action permit ↲
exit ↲
entry 30 ↲
```

```
match "help" ↲
action permit ↲
exit ↲
entry 40 ↲
match "back" ↲
action permit ↲
exit ↲
entry 50 ↲
match "show li" ↲
action permit ↲
exit ↲
entry 60 ↲
match "info" ↲
action permit ↲
exit ↲
entry 70 ↲
match "configure system security user" ↲
action permit ↲
exit ↲
entry 80 ↲
match "admin save" ↲
action permit ↲
exit ↲
exit all ↲
admin save ↲
```

3

Close the SSH session.

4

Right-click on the NE in the NFM-P topology map and choose Resync→Resync All MIBs from the contextual menu to update the NE configuration in the NFM-P.

END OF STEPS

94.10 To create an NE LI user account on an NE using a CLI

94.10.1 Steps

You require NFM-P admin privileges to perform this procedure.

1

Open an SSH session on the NE as described in Chapter 10, “Device CLI sessions” .

2

Enter the following commands in sequence at the prompt to create an LI user:

```
configure system security snmp ↲
access group "LI_group" security-model usm security-level privacy read
"iso" write "iso" notify "iso" ↲
access group "LI_group" security-model usm security-level privacy
context "li" exact read "li-view" write "li-view" notify "iso" ↲
exit all ↲
configure system security user LI_username ↲
password LI_password ↲
snmp group LI_group ↲
console member LI_profile ↲
console no member default ↲
exit all ↲
admin save ↲
```

where

LI_group is the name of the LI user group on the NE

LI_user is the name of the LI user account on the NE

LI_password is the password for the LI user account on the NE

LI_profile is the name of an LI user profile; see 94.9 “To create an NE LI user profile on an NE using a CLI” (p. 3200) for information about creating an LI user profile

3

Close the SSH session.

4

Right-click on the NE in the NFM-P topology map and choose Resync→Resync All MIBs from the contextual menu to update the NE configuration in the NFM-P.

END OF STEPS

94.11 To create additional NE LI user accounts using the NFM-P

94.11.1 Steps

Note: You require an LI user account to perform this procedure.
Before you can perform this procedure, at least one LI user account must exist on the NE.

Before you can perform this procedure, an NFM-P LI user must enable LI discovery for the NE using 94.20 “To enable NE discovery for LI” (p. 3214) .

1

Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens.

2

Click Create. The NE User (Create) form opens.

3

Configure the required parameters.

You must specify the li, snmp, and console values for the Access parameter.

Specify an MD5 authentication key for the Password parameter. See 94.12 “To configure NE LI user security” (p. 3204) for information about generating a user authentication key.

4

To assign one or more console profiles:

1. Click on the Console Profiles tab. The list of profiles numbered one through eight appears. A user can have up to eight console profiles.
2. Select a console profile for each of the profiles 1 through 8. See the procedure to configure a user account on a managed device in the *NSP System Administrator Guide* for information about creating console profiles.

5

Configure SNMPv3.

1. Click on the SNMPv3 tab.
2. Configure the authentication and privacy parameters. Ensure the NEs support a valid combination of authentication and privacy protocols.
3. Specify an authentication key generated using 94.12 “To configure NE LI user security” (p. 3204) for the New Authentication Password and Confirm New Auth Password parameters.
4. Specify an encryption key generated using 94.12 “To configure NE LI user security” (p. 3204) for the New Privacy Password and Confirm New Privacy Password parameters.

6

Save the changes and close the forms.

END OF STEPS

94.12 To configure NE LI user security

94.12.1 Steps

Note: You require NE LI user privileges to perform this procedure.

Perform this procedure to change the password for an LI user account on an NE and to configure LI user authentication and SNMP data encryption on the NE.

1

Use the NE LI user account to open an SSH session on the NE. See Chapter 10, “Device CLI sessions” for information about opening an SSH session on an NE.

2

Enter the following command to obtain the SNMP engine ID of the NE.

`show system info` ↵

The SNMP engine ID is displayed as SNMP Engine ID.

3

Record the SNMP engine ID for use in the following steps.

4

Generate an authentication key and a privacy key.

- An authentication key is used to encrypt a user password.
- A privacy key is used to encrypt the user SNMP packets.

Note: The key authentication method determines the key length.

1. Log in to an NFM-P single-user client, client delegate server, or main server station.

   Note:

   If you log in to a RHEL main or client delegate server station, you must log in as the nsp user.

   If you log in to a single-user client station, you must log in as the user who installed the client, or as a local administrator.
2. Open a console window.
3. On a RHEL station, navigate to the *install_directory*/nms/bin directory, where *install_directory* is one of the following:
   - the NFM-P main server installation location, /opt/nsp/nfmp/server
   - the NFM-P single-user client or client delegate server installation location, typically /opt/nsp/client
4. On a Windows station, navigate to the *install_directory*\nms\bin directory, where *install_directory* is the NFM-P single-user client or client delegate server installation location, typically C:\nsp\client.
5. Enter one of the following to create an authentication key:

- on a RHEL station:
  ./nmsclient.bash password2key *method password engine_ID* ↵
- on a Windows station:
  nmsclient.bat password2key *method password engine_ID* ↵

  where

  method is the authentication method, either MD5, SHA, SHA224, SHA256, SHA384 or SHA512

  password is the authentication key password

  engine_ID is the SNMP engine ID obtained in Step 2

  Note: You must enclose a password that contains a special character in single quotation marks; for example:

  password2key *method* 'Mypa$$word'

  Only use the authentication key from the output.

6. Enter the following to create a privacy key.
   - on a RHEL station:
     ./nmsclient.bash password2key *method password engine_ID* ↵
   - on a Windows station:
     nmsclient.bat password2key *method password engine_ID* ↵

     where

     method is the authentication method, either MD5, SHA, SHA224, SHA256, SHA384 or SHA512

     password is the privacy key password

     engine_ID is the SNMP Engine ID of the SR, in hexadecimal form with 10-64 hex digits (5-32 bytes)

     Note: You must enclose a password that contains a special character in single quotation marks; for example:

     password2key *method* 'Mypa$$word'

     The list of privacy keys for each privacy method is displayed.
7. Store the generated keys for your applicable authentication and privacy methods.

5

Using the keys generated in Step 4 , enter the following commands in sequence at the CLI prompt to change the LI user password and to configure LI security for the user account

1. Enter the following sequence of commands at the prompt:

   configure system security user *username* ↵

   password *new_LI_password* ↵

   snmp ↵

   authentication *auth_method authentication_key* privacy *priv_method privacy_key* ↵

   group *SNMPv3_group* ↵

   exit all ↵

where

username is the name of the LI user account on the NE

new_LI_password is the new password for the LI user account on the NE

auth_method can be:

hmac-md5-96 hmac-sha1-96 hmac-sha2-224 hmac-sha2-256 hmac-sha2-384 hmac-sha2-512

authentication_key is the authentication key value generated in Step 4

priv_method can be:

cbc-des cfb128-aes-128 cfb128-aes-192 cfb128-aes-256

privacy_key is the privacy key value generated in Step 4

2. Enter the following to save the configuration changes:

   admin save ↵

3. Close the CLI session.

END OF STEPS

94.13 To synchronize the global NFM-P NE LI user configuration profile with the local NE LI user configuration profile

94.13.1 Steps

Note: You require an LI user account to perform this procedure. If you do not synchronize the local profile with the global profile, the NE LI user does not appear in the list of users for SNMPv3 for the LI mediation policy.

You must perform this procedure before you can create an LI mediation policy.

1

Choose Administration→Security→NE User Configuration from the NFM-P main menu. The NE User Configuration form opens with a list of NE users.

2

Choose the NE LI user you created in 94.10 “To create an NE LI user account on an NE using a CLI” (p. 3201) and click Synchronize. The Synchronize - SR Local User form opens.

3

Choose the NE on which the local policy is defined from the Available Nodes panel and click on the right-arrow to move the NE to the Selected Node panel.

4

Click Synchronize.

5

Close the Synchronize -SR Local User form.

6

Close the NE User Configuration form.

END OF STEPS

94.14 To configure an LI security mediation policy

94.14.1 Steps

Note: You require an LI user account to perform this procedure.

The LI mediation security policy defines the network security model used when creating LI mirror objects for LI users. LI users do not use the standard mediation security policy. This policy identifies the NE LI user created specifically for LI purposes, the authentication and privacy protocol types, and the required passwords.

1

Choose Administration→Mediation from the NFM-P main menu. The Mediation (Edit) form opens.

2

Click on the LI Mediation Security tab.

3

Click Create. The LI Mediation Policy (Create) form opens.

4

Select the User Name you created in 94.10 “To create an NE LI user account on an NE using a CLI” (p. 3201) in the SNMPv3 panel.

5

The SNMP passwords in the NFM-P mediation policy must match the encrypted passwords generated for the NE LI user in 94.12 “To configure NE LI user security” (p. 3204) .

Set the mediation policy passwords.

1. Click Properties. The NE User Configuration form opens.
2. Click on the SNMPv3 tab.
3. Configure the authentication and privacy parameters. Ensure the NEs support a valid combination of authentication and privacy protocols.
4. Specify the password used to generate the authentication key in 94.12 “To configure NE LI user security” (p. 3204) for the New Authentication Password and Confirm New Auth Password parameters.

5. Specify the password used to generate the encryption key in 94.12 “To configure NE LI user security” (p. 3204) for the New Privacy Password and Confirm New Privacy Password parameters.
6. Save the changes and close the form.

6

Save the changes and close the forms.

END OF STEPS

94.15 To configure an LI MAC filter policy

94.15.1 Steps

Note: You are required to be an NFM-P user with LI privileges to access the LI MAC filter policy.

You require at least one ACL MAC filter policy to complete this procedure. See 51.4 “To configure an ACL MAC filter policy” (p. 1668) for more information.

You require at least one Block Reservation policy to complete this procedure. See 94.18 “To configure a Block Reservation policy” (p. 3212) for more information.

1

Choose Policies→Filter→LI Policies→LI MAC Filter from the NFM-P main menu. The LI MAC Filter Policies form opens.

2

Click Create or select a policy and click Properties. The LI MAC Filter, Global Policy (Create|Edit) form opens.

3

Configure the required general parameters.

4

Configure LI MAC filter entries:

1. Click on the LI MAC Filter Entries tab.
2. Click Create. The LI MAC Filter Entry, LI MAC Filter, Global Policy (Create) form opens.
3. Configure the required parameters.
4. Save the changes and close the form.

5

Configure the LI MAC filter association:

1. Click on the MAC Filter Association tab.
2. Click Create. The LI MAC Filter Association, LI MAC Filter form opens.

3. Select an ACL MAC filter to associate with the LI MAC filter policy.
4. Save the changes and close the form.

6

Click Apply.

7

Click on the General tab.

8

Click Switch Mode to release the policy for distribution. The Release - LI MAC Filter Policy form opens.

9

Choose the NEs in the Available Objects to which you want to distribute the policy and click the arrow to move the NEs to the Selected Objects panel.

10

Click Distribute and close the form.

11

Close the LI MAC Filter Policies form.

END OF STEPS

94.16 To configure an LI IP filter policy

94.16.1 Steps

Note: You are required to be an NFM-P user with LI privileges to access the LI IP filter policy.

You require at least one IP filter policy to complete this procedure. See 51.5 “To configure an ACL IP filter policy” (p. 1671) for more information.

You require at least one Block Reservation policy to complete this procedure. See 94.18 “To configure a Block Reservation policy” (p. 3212) for more information.

1

Choose Policies→Filter→LI Policies→LI IP Filter from the NFM-P main menu. The LI IP Filter Policies form opens.

2

Click Create or select a policy and click Properties. The LI IP Filter (Create|Edit) form opens.

3

Configure the required general parameters.

4

Configure one or more filter entries:

1. Click on the LI IP Filter Entries tab.
2. Click Create. The Entry, LI IP Filter (Create) form opens.
3. Configure the parameters.
4. Click on the Filter Properties tab.
5. Configure the Protocol parameter.
6. Configure the parameters in the Match Criteria panel.

   The Source Port and Dest Port parameters are configurable only when the Protocol parameter is set to TCP or UDP.
7. Save the changes and close the form.

5

Associate the base IP filter with the LI IP filter:

1. Click on the IP Filter Association tab.
2. Click Create. The LI IP Filter Association (Create) form opens.
3. Select an IP filter in the IP Filter panel.
4. Save the changes and close the form.

6

Click Apply.

7

Click on the General tab.

8

Click Switch Mode to release the policy for distribution. The Release - LI IP Filter Policy form opens.

9

Choose the NEs in the Available Objects to which you want to distribute the policy and click the arrow to move the NEs to the Selected Objects panel.

10

Click Distribute and close the form.

11

Close the LI IP Filter Policies form.

END OF STEPS

94.17 To configure an LI IPv6 filter policy

94.17.1 Steps

Note: You are required to be an NFM-P user with LI privileges to access the LI IPv6 filter policy.

You require at least one ACL IPv6 filter policy to complete this procedure. See 51.6 “To configure an ACL IPv6 filter policy” (p. 1677) for more information.

You require at least one Block Reservation policy to complete this procedure. See 94.18 “To configure a Block Reservation policy” (p. 3212) for more information.

1

Choose Policies→Filter→LI Policies→LI IPv6 Filter from the NFM-P main menu. The LI IPv6 Filter Policies form opens.

2

Click Create or select a policy and click Properties. The LI IPv6 Filter (Create|Edit) form opens.

3

Configure the required general parameters.

4

Configure one or more filter entries:

1. Click on the LI IPv6 Filter Entries tab.
2. Click Create. The Entry, LI IPv6 Filter (Create) form opens.
3. Configure the required general parameters.
4. Click on the Filter Properties tab.
5. Configure the Protocol parameter.
6. Configure the parameters in the Match Criteria panel:

   The Source Port and Dest Port parameters are configurable only when the Protocol parameter is set to TCP or UDP.
7. Save the changes and close the form.

5

Associate the base IPv6 filter with the LI IPv6 filter:

1. Click on the IPv6 Filter Association tab.
2. Click Create. The LI IPv6 Filter Association (Create) form opens.

3. Select an IPv6 filter in the IPv6 Filter panel.
4. Save the changes and close the form.

6

Click Apply.

7

Click on the General tab.

8

Click Switch Mode to release the policy for distribution. The Release - LI IPv6 Filter Policy form opens.

9

Choose the NEs in the Available Objects to which you want to distribute the policy and click the arrow to move the NEs to the Selected Objects panel.

10

Click Distribute and close the form.

11

Close the LI IPv6 Filter Policies form.

END OF STEPS

94.18 To configure a Block Reservation policy

94.18.1 Steps

Note: You are required to be an NFM-P user with LI privileges to access the Block Reservation policy.

You require at least one ACL MAC filter policy. See 51.4 “To configure an ACL MAC filter policy” (p. 1668) for more information.

You require at least one IP filter policy. See 51.5 “To configure an ACL IP filter policy” (p. 1671) for more information.

1

Choose Policies→Filter→LI Policies→Block Reservation from the NFM-P main menu. The Block Reservation Policies form opens.

2

Click Create or select a policy and click Properties. The Block Reservation, Global Policy (Create|Edit) form opens.

3

Configure the required general parameters.

4

Create a MAC filter LI block reservation association:

1. Click on the MAC Filters tab.
2. Click Create. The MAC Filter LI Block Reservation Association, Block Reservation, Global Policy (Create) form opens.
3. Select an ACL MAC filter to associate with the Block Reservation policy.
4. Save the changes and close the form.

5

Create an IPv4 or IPv6 filter LI block reservation association:

1. Click on the IP Filters tab.
2. Click Create. The IP Filter LI Block Reservation Association, Block Reservation, Global Policy (Create) form opens.
3. Select an IP filter to associate with the Block Reservation policy.
4. Save the changes and close the form.

6

Click Apply.

7

Click on the General tab.

8

Click Switch Mode to release the policy for distribution. The Release - Block Reservation Policy form opens.

9

Choose the NEs in the Available Objects to which you want to distribute the policy and click the arrow to move the NEs to the Selected Objects panel.

10

Click Distribute and close the form.

11

Close the Block Reservation Policies form.

END OF STEPS

94.19 To configure the LI filter lock

94.19.1 LI Filter Lock options

This procedure allows an NFM-P LI user to configure a parameter that permits base IPv4 and MAC filters referenced by an LI Source to be modified or deleted. The LI Filter Lock parameter specifies who can modify the LI filters. The available options are:

- Locked: no user can modify the LI filters
- Unlocked For LI Users: only users with LI privileges can modify the LI filters
- Unlocked For All: all users can modify the LI filters

94.19.2 Steps

Note: You require LI privileges to view the LI Configuration Status tab and to perform this procedure.

1

Right-click on a discovered device in the Equipment navigation tree and choose Properties. The Network Element (Edit) form is displayed.

2

Click on the LI Configuration Status tab.

3

Configure the LI Filter Lock parameter.

4

Save the changes and close the form.

END OF STEPS

94.20 To enable NE discovery for LI

94.20.1 Steps

Note: You require LI privileges to perform this procedure.

Perform this procedure to enable LI discovery of an NE. Before you can enable LI discovery for an NE, the NE must be successfully discovered. See Chapter 9, “Device discovery” for information about configuring NE discovery.

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager (Edit) form opens with the Discovery Rules tab displayed.

2

Click on the Managed State tab. A list of discovered NEs is displayed.

3

Choose an NE on which 94.9 “To create an NE LI user profile on an NE using a CLI” (p. 3200) and 94.10 “To create an NE LI user account on an NE using a CLI” (p. 3201) have been performed, and click Properties. The Node Discovery Control (Edit) form opens.

4

Click on the LI Mediation Security tab.

5

Select the LI mediation policy created in 94.14 “To configure an LI security mediation policy” (p. 3207) in the Dual Read Access Mediation Policy panel.

6

Select the LI mediation policy created in 94.14 “To configure an LI security mediation policy” (p. 3207) in the Dual Write Access Mediation Policy panel.

7

Select the LI mediation policy created in 94.14 “To configure an LI security mediation policy” (p. 3207) in the Dual Trap Access Mediation Policy panel.

8

Save the changes and close the form.

9

On the Discovery Manager (Edit) form, click on the Resync Status tab.

10

Choose the NE and click Resync. The Resync Options form opens.

11

Select Choose MIB Entries and click Next. The Choose MIB Entries list form opens.

12

Select all TIMETRA-MIRROR-MIB entries in the list and click Finish.

Note: Resynchronizing only the TIMETRA-MIRROR-MIB entries resynchronizes only LI source objects and keeps the resynchronization time to the minimum required.

13

Close the forms.

END OF STEPS

94.21 To save the LI configuration of an NE

94.21.1 Steps

Note: You require LI privileges to view the LI Configuration Status tab and to perform this procedure.

This procedure allows an NFM-P LI user to save the LI configuration file (li.cfg) from the NFM-P for an NE. The li.cfg file is updated only when changes to the LI configuration are deployed to the NE.

1

Right-click on a discovered device in the Equipment navigation tree and choose Properties. The Network Element (Edit) form is displayed.

2

Click on the LI Configuration Status tab.

3

Click Save LI Configuration.

4

Close the form.

END OF STEPS

94.22 To configure Layer 3 encapsulation on a source site to allow LI-mirrored packets to be placed into a routable header

94.22.1 Related information

You can configure a Layer 3 encapsulation option to allow LI-mirrored packets to be placed into a routable header (either IP-UDP Shim or IP-GRE) and then forwarded in a routing context.

The following conditions apply if you configure the Layer 3 encapsulation option:

- Configuring Mirror SDP bindings and Mirror SAPs is mutually exclusive to the L3 Encapsulation
- You cannot configure the Session ID or Intercept ID of a source later on in this procedure if the L3 Encapsulation is not configured here.
- You cannot configure L3 Encapsulation if a mirror source object (such as: IP filter, MAC filter, and so on) is first configured for this service.

- You cannot change the routing instance of L3 Encapsulation if a gateway is first configured for this service.

You can configure L3 Encapsulation only on a source site on an NE in chassis mode D. See Chapter 15, “Shelf and card object configuration” for more information about chassis modes.

94.22.2 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

On the service navigation tree, expand Source Sites and click on the source site. The Mirror Site (Edit) form opens.

4

Click on the L3 Encapsulation tab and click Create. The Layer-3 Encapsulation null (Create) form opens

5

Configure the required parameters.

The Use Direction Bit parameter is only displayed if you configured the Header Type option for IP-UDP Shim.

6

Select a routing instance ID in the Routing Instance panel.

7

To configure a gateway:

1. Click Create in the Gateway panel. The Layer-3 Encapsulation Gateway null (Create) form opens.
2. Configure the required parameters in the IP Address panel.

   The gateway Source Address and Destination Address must either both be unconfigured or both configured. Similarly, the gateway Source Port and Destination Port must either both be unconfigured or both configured.

   The Source Port and Destination Port parameters are displayed only if you configured the Header Type option for IP-UDP Shim.
3. Save the changes and close the form.

8

Save the changes and close the forms.

END OF STEPS

94.23 To specify an LI MAC filter entry as an LI source

94.23.1 Steps

Specifying an LI MAC filter entry as the mirror source causes all the packets matching the filter to be mirrored to the mirror destination specified by the service ID of the mirror source.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following:

a. To specify an LI source LI IP filter entry on a destination site, expand the Destination Site object to view the configured destination sites, and expand the destination site object on which you are specifying the LI source LI IP filter entry.

b. To specify an LI source LI IP filter entry on a source site, expand the Source Site object to view the configured source sites, and expand the source site object on which you are specifying the LI source LI IP filter entry.

4

Right-click on LI Source LI IP Filters and choose Create LI Source LI IP Filter. The LI Source LI IP Filter (Create) form opens.

Note: A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

5

Select an LI source LI IP filter entry.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

94.24 To specify an LI Source Port as an LI source

94.24.1 Purpose

Use this procedure to configure an LI Source Port mirror on a port or LAG. The LI Source Port can mirror Egress or Ingress traffic or both directions.

Note: The following considerations apply:

- The NE must be licensed to allow LI source port mirroring.
- When an LI Source Configuration is enabled, all Mirror Service sources (such as Interfaces, Ports, IP filter, or Subscribers) can no longer be provisioned on that Mirror Service.

  The admin user must create a new unique Service Mirror for non LI source types, and not share the same Mirror Service with and LI Mirror Service(both having the same service ID number).
- All existing Mirror Service sources will be removed if LI Source Configuration is enabled, or any type of LI Mirror source is provisioned.

94.24.2 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and Properties. The Mirror Service (Edit) form opens.

3

To specify an LI source port on a source site, expand Source Sites→*source_site*, right-click on LI Source Ports and choose Create Source Port. The Source Port (Create) form opens.

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.
A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

4

Select a port.

An LI source port can be one of the following:

- physical port
- LAG

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

94.25 To specify an LI IP filter entry as an LI source

94.25.1 Steps

Specifying an LI IPv4 filter entry as the mirror source causes all the packets matching the filter to be mirrored to the mirror destination specified by the service ID of the mirror source.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following:

a. To specify an LI source LI IP filter entry on a destination site, expand the Destination Site object to view the configured destination sites, and expand the destination site object on which you are specifying the LI source LI IP filter entry.

b. To specify an LI source LI IP filter entry on a source site, expand the Source Site object to view the configured source sites, and expand the source site object on which you are specifying the LI source LI IP filter entry.

4

Right-click on the LI Source LI IP Filters object below the site in the service navigation tree and choose Create LI Source LI IP Filter. The LI Source LI IP Filter (Create) form opens.

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.

A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

5

Select an LI MAC filter entry.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

94.26 To specify an LI IPv6 filter entry as an LI source

94.26.1 Steps

Specifying an LI IPv6 filter entry as the mirror source causes all the packets matching the filter to be mirrored to the mirror destination specified by the service ID of the mirror source.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and Properties. The Mirror Service (Edit) form opens.

3

Perform one of the following:

a. To specify an LI source LI IPv6 filter entry on a destination site, expand the Destination Site object to view the configured destination sites, and expand the destination site object on which you are specifying the LI source LI IPv6 filter entry.

b. To specify an LI source LI IPv6 filter entry on a source site, expand the Source Site object to view the configured source sites, and expand the source site object on which you are specifying the LI source LI IPv6 filter entry.

4

Right-click on the LI Source LI IPv6 Filters object below the site in the service navigation tree and choose Create LI Source LI IPv6 Filter. The LI Source LI IPv6 Filter (Create) form opens.

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.

A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

5

Select an LI source LI IPv6 filter entry.

6

Configure the required parameters.

7

Save the changes and close the forms.

END OF STEPS

94.27 To specify an LI SAP as an LI source

94.27.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Right-click on the LI Source SAPs object below the site in the service navigation tree and choose Create LI Source Interface. The LI Source Interface (Create) form opens. Go to Step 4 .

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.
A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

4

Select an access interface to associate with the source interface.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

94.28 To specify an LI subscriber as an LI source

94.28.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and click Properties. The Mirror Service (Edit) form opens.

3

Expand the Source Site object to view the configured source sites, and expand the source site object on which you are specifying the subscriber as the mirror source.

4

Right-click on the LI Source Subscribers object below the site in the service navigation tree and choose Create LI Source Subscriber. The LI Source Subscriber (Create) form opens.

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.
A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

5

Configure the required parameters.

6

You can further restrict the mirrored subscriber host traffic associated with the subscriber identification string by specifying SAP or SLA-profile criteria. To specify an SLA profile as a match criterion, go to Step 8 .

Note: You can configure parameters in the SAP/Subscriber Host Match Criteria panel or the SLA Profile Match Criteria panel, but you cannot configure parameters in both panels.

7

To specify a SAP:

1. Select a SAP on which to mirror subscriber host traffic in the SAP/Subscriber Host Match Criteria panel.

   You can select only dot1q- or QinQ-encapsulated SAPs that have subscriber management enabled.
2. Configure the parameters in the SAP/Subscriber Host Match Criteria panel to specify the subscriber host match criteria for the SAP.

   Note:

The NFM-P does not accept the parameter values unless a SAP is specified.

3. Go to Step 9 .

8

To specify SLA profile match criteria, select the SLA profile in the SLA Profile Match Criteria panel.

9

Save the changes and close the forms.

END OF STEPS

94.29 To specify an LI WLAN distributed subscriber as an LI source

94.29.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Select a mirror service and Properties. The Mirror Service form opens.

3

Expand the Source Sites object to view the configured source sites, and expand the source site object on which you are specifying a WLAN distributed subscriber as the mirror source.

4

Right-click on the LI Source WLAN Distributed Subscribers object below the site in the service navigation tree and choose Create LI Source WLAN Distributed Subscriber. The LI Source WLAN Distributed Subscriber (Create) form opens.

Note: An NFM-P operator with LI privileges can view, create, and delete LI source objects.
A mirror site can be configured with debug or LI sources, but not both. LI source configuration takes priority over debug source configuration.

5

Configure the required parameters.

6

Save the changes and close the forms.

END OF STEPS

94.30 To configure LI on a specific NAT subscriber

94.30.1 Steps

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Choose a mirror service and Properties. The Mirror *Service Name* (Edit) form opens.

3

Click on the Sites tab. A list of mirror service sites are displayed.

4

Choose a site and click Properties. The Mirror Site (Edit) form opens.

5

Click on the LI NAT Sources tab.

Note: The LI NAT Sources tab appears only if the Encapsulation Type parameter for the mirror site is set to Ethernet and the L3 encapsulation Header Type parameter is set to IP-UDP Shim. See 94.22 “To configure Layer 3 encapsulation on a source site to allow LI-mirrored packets to be placed into a routable header” (p. 3216) for more information about how to set the Header Type parameter.

6

Click Configure LI Source Configuration.

7

Click on the Ethernet Header tab.

8

Configure the required parameters.

9

Click on the L2 Aware Subscribers tab.

10

Click Create. The LI Source NAT L2 Aware Subscriber (Create) form opens.

11

Configure the required parameters.

12

Click OK. The LI Source NAT L2 Aware Subscriber (Create) form becomes a tab.

13

Click OK. The LI Source NAT L2 Aware Subscriber (Create) tab closes and the Mirror Service (Edit) form reappears with the Sites tab displayed.

14

Choose a site and click Properties. The Mirror Site (Edit) form opens.

15

Click on the LI NAT Sources tab. The Ethernet Header tab is displayed.

16

Click on the Dual Stack Lite LSN Subscribers tab.

17

Click Create. The LI Source Dual Stack Lite LSN Subscriber (Create) form opens.

18

Configure the required parameters.

19

Click OK. The LI Source Dual Stack Lite LSN Subscriber (Create) form becomes a tab.

20

Click OK. The LI Source Dual Stack Lite LSN Subscriber (Create) tab closes and the Mirror Service (Edit) form reappears with the Sites tab displayed.

21

Choose a site and click Properties. The Mirror Site (Edit) form opens.

22

Click on the LI NAT Sources tab. The Ethernet Header tab is displayed.

23

Click on the Classic LSN Subscribers tab.

24

Click Create. The LI Source Classic LSN Subscriber (Create) form opens.

25

Configure the required parameters.

26

Click OK. The LI Source Classic LSN Subscriber (Create) form becomes a tab.

27

Click OK. The LI Source Classic LSN Subscriber (Create) tab closes and the Mirror Service (Edit) form reappears with the Sites tab displayed.

28

Click on the NAT 64 LSN Subscribers tab.

29

Click Create. The LI Source NAT 64 LSN Subscriber (Create) form opens.

30

Configure the required parameters.

31

Click OK. The LI Source NAT 64 LSN Subscriber (Create) form becomes a tab.

32

Click OK. The LI Source NAT 64 LSN Subscriber (Create) tab closes and the Mirror Service (Edit) form reappears with the Sites tab displayed.

33

Save the changes and close the forms.

END OF STEPS

94.31 To view LI mirrored subscriber hosts configured with a RADIUS server

94.31.1 Steps

Note: You must have NFM-P LI user privileges to perform this procedure.

Before you can perform this procedure, at least one LI user account must exist on the NE.

Before you can perform this procedure, an NFM-P LI user must enable LI discovery for the NE using 94.20 “To enable NE discovery for LI” (p. 3214) .

Only PPPoE subscriber hosts can be mirrored with RADIUS.

1

Choose Manage→Service→Mirror Services from the NFM-P main menu. The Manage Mirror Services form opens.

2

Choose a mirror service and click Properties. The Mirror *Service Name* (Edit) form opens.

3

Click on the Sites tab. A list of mirror service sites is displayed.

4

Choose a site and click Properties. The Mirror Site (Edit) form opens.

5

If an LI source configuration object has not been created, the LI Source Configuration tab is dimmed. Click Create LI Source Configuration at the bottom of the form. A dialog box appears.

Note: You must create an LI source configuration object for the NE to mirror subscriber hosts for LI.
The Create LI Source Configuration button is a toggle that also lets you delete an LI source configuration object. If you delete an LI source configuration object, all associated LI source objects are deleted.
An LI source configuration object is automatically created when an LI source object is created.

6

Click OK. The LI Source Configuration tab is enabled, and an LI source configuration object is created.

7

Click on the LI Source Subscribers Via RADIUS tab.

8

Choose a subscriber and click Properties. The LI Source Subscriber Host form opens.

9

View the information on the form.

10

Close the LI Source Subscriber Host form.

END OF STEPS

95 RCA audit

95.1 Overview

95.1.1 Purpose

This chapter describes the NFM-P root cause analysis functionality, and provides information about configuring RCA.

95.1.2 Contents

RCA audit overview

95.2 RCA audit concepts

95.2.1 RCA audit description

You can perform on-demand or scheduled root cause analysis and verification of the configuration of routing instances, services, and other network objects to identify possible configuration problems. Except for physical links, the NFM-P provides a solution, which, at your request, can automatically be implemented to make all of the required configuration changes.

Note: The adjustments are made only to the NFM-P database and are not deployed to the network.

Only VPLS, VPRN, and VLL services RCA audits and golden configuration RCA audits can be scheduled.

95.2.2 RCA audit objects

You can perform RCA audits of the following objects:

- physical links
- VPLS, VPRN, and VLL services and service site memberships
- OSPF interfaces, areas, and area sites (NFM-P/CPAM integration only)
- IS-IS interfaces and sites (NFM-P/CPAM integration only)

95.3 NFM-P service audit

95.3.1 Service audit description

An NFM-P service is defined as a collection of service sites with the same customer ID, service type, and service ID. The NFM-P discovers services that are configured on a 7210 SAS, 7450 ESS, 7750 SR, or 7950 XRS, using the service ID on the NE. Configuration errors may occur in networks where the services were created and deployed on the NEs using CLI before the NEs were managed by the NFM-P.

For example, when two VPLSs with the same service ID and the same mesh VC ID on an NE are discovered by the NFM-P, they are discovered as a single service in the NFM-P.

Another example of a configuration error is when a switching Epipe has multiple Epipe sites and multiple VC IDs for different segments. If different service IDs are used when the sites are created, the NFM-P assumes that the sites are connected and creates multiple VLL services within a composite service.

RCA audit policies allow you to modify the component membership of your NFM-P services to detect possible configuration problems. In addition, you can use the RCA audit to correct most configuration problems that are discovered in the audit.

Note: An NFM-P user who is assigned the Administrator or RCA scope of command role can create, modify, and execute all RCA audit policies. An NFM-P user who is assigned the Administrator or Service scope of command role can execute service audit adjustments.

95.3.2 Service membership rules

The NFM-P uses the rules outlined in the following figure to identify whether two service sites are connected. When two sites are identified as connected, they should belong to same service.

Figure 95-1 Rules to check whether two sites are connected

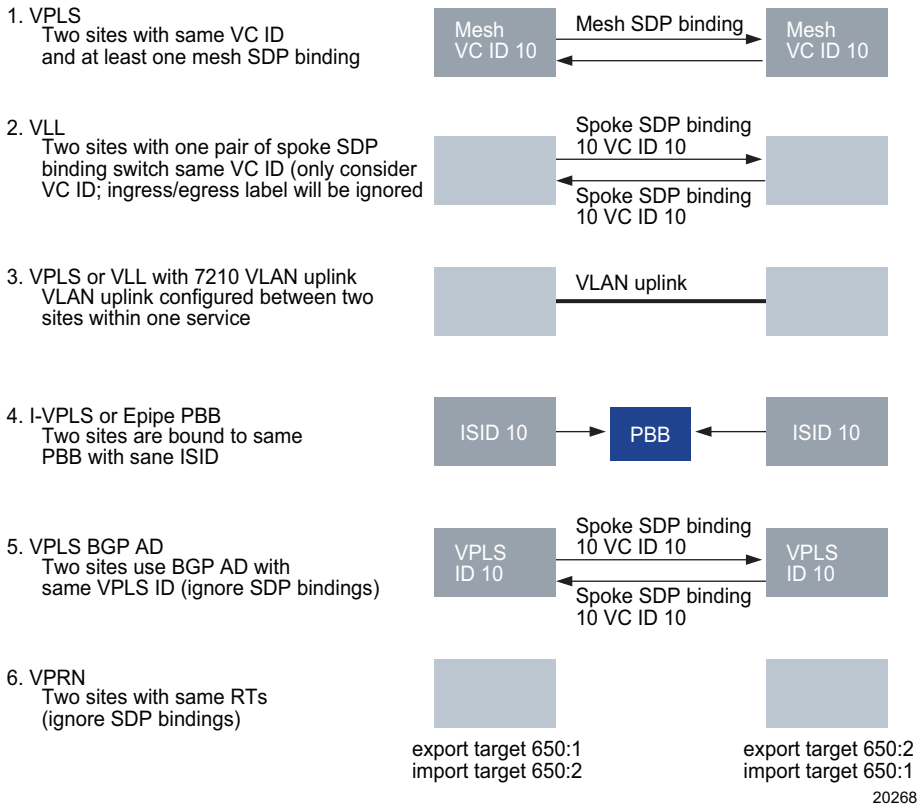

If service sites have different customer IDs, the NFM-P discovers the sites as belonging to different services. After the audit, the NFM-P generates a problem for each service. You can specify to which service the NFM-P should move the site and whether the empty service should be removed after the sites are moved.

For VPLS or VLL services, if the same service ID is used for two groups of sites that are not connected, the NFM-P detects a duplicate service ID. Two problems are generated for one of the group of sites, and the solution is to separate the services. If there are other groups of isolated sites, additional problems are generated for each group.

Consider the following when you perform an audit of a VPLS.

- Although different service IDs may be used, I-sites that are bound to same backbone VPLS and have the same ISID are considered to be in the same VPN.
- For VPLS sites that use BGP auto-discovery, the VPLS ID is used to determine service membership.
- H-VPLS is discovered as composite service in the NFM-P if different service IDs are used.

Consider the following when you perform an audit of a VLL service.

- Service sites with different customer IDs are discovered as two services, and can be reconfigured as one service.
- More than two groups of sites that are connected can be detected and can be separated into different services.
- Redundant VLLs are not affected by the audit and are considered correctly configured.
- If two Epipe sites are connected to the same PBB with same B-VPLS service ID, and the source and destination MAC addresses match, the NFM-P determines that the two sites are connected. If they are in different NFM-P services, a problem is generated during the audit.

Consider the following when you perform an audit of a VPRN service.

- VPRN service membership is based on the RTs defined in the VRFs.
- Sites that are in different services but have common import and export route targets. Import target of Site 1 is equal to the Export target of Site 2.

95.3.3 RCA audit for services with multi-vendor NE sites

The NFM-P attempts to place MV sites into the correct service during discovery. For MV sites which are not placed into the correct service, the MV Complete Service Topology function moves the MV sites and matches them with the service. When you click the MV Complete Service Topology button on the Manage Services form, the NFM-P performs a search for services containing MV sites. If a service does not contain any MV NEs, the NFM-P skips that service. For each service, the NFM-P performs an RCA audit and the following actions occur:

- Epipe service

  The NFM-P determines if an Epipe service has at least one MV site, and then checks all other Epipe services for a site with the same VC ID and with SDP bindings pointing to each other. If a matching site is found, the NFM-P moves the site from the first service to the second (matching) service. As a result, the two services are collated. If a service has two or more MV sites which do not meet the same criteria, one site is moved to either the correct matching service, or a new service. This action results in two services potentially assigned to different customers being merged into one service. No changes are made on the NE.

- VPLS service

  A VPLS service can have one or more sites. The RCA audit only supports mesh SDPs. The MV Complete Service Topology function does not support VPLS sites with spoke SDPs.

  When the NFM-P finds a VPLS service with at least one MV site, it checks all VPLS services for a site with the same mesh VC ID and the mesh SDPs pointing to each other. All of the matching sites that are in different services are moved to a single service. Any mismatched sites are moved to either the correct matching service, or a new service. This action results in two services potentially assigned to different customers being merged into one service. No changes are made on the NE.

- VPRN service

  Note: Do not re-use route targets on the NEs. Route targets must be unique because the RCA audit uses route targets to match VPRN sites that are in different NFM-P services or that have different service IDs. Reusing route targets causes the sites to be moved to the same service and the sites are physically separate.

  When the NFM-P finds a VPRN service with a MV site, it checks for all of the VPRN sites (MV or SR) that use the same RT value. The NFM-P then moves the MV sites from the first service to the second service that has other MV or SR sites. This step is repeated for all of the MV sites that have the same RT value. This action results in two services previously assigned to different customers being merged into one service.

95.4 Physical link audits

95.4.1 Physical link description

Physical links represent the actual physical configuration of network connections between ports. You can view and manage physical links from the equipment window, physical topology map, and the Manage Equipment list form of the NFM-P, on each router using the CLI. Because several key parameters on each end of the physical link depend on each other, configuration errors are possible. An RCA audit can identify configuration errors in physical links. The NFM-P does not provide a solution for configuration problems that the RCA audit identifies for physical links.

95.4.2 Configuration errors in physical links

By default, the RCA audit detects the following configuration errors in physical links:

- **physical port parameters**
  - mismatched MTU values
  - mismatched speeds
- **Ethernet port parameters**
  - Auto-negotiate parameter misconfiguration
  - Duplex parameter misconfiguration

You can configure the RCA audit policy for the physical link to include additional physical link properties in the audit.

95.5 Viewing and analyzing RCA audit results

95.5.1 RCA audit results

After you manually perform an RCA audit, or the RCA audit executes as per the schedule details, you can determine if problems were detected, when the last audit was performed, and view the correction plan for a problem. The RCA Audit Problem(s) indicator on the General tab of a network object properties form identifies whether configuration problems were detected in previous audits. The Last Audit Time indicator displays a timestamp of the last audit that was performed. The RCA Audit button is located at the bottom of the service properties form, or under the More Actions button.

After you associate an audit policy with the object, you can perform an RCA audit and view the results. If no problems are detected, the RCA Result tab does not appear.

The properties form of a problem displays the following information:

- problem severity
- probable cause
- description
- solution

The Caused By Objects tab lists the network objects that caused the problem. For service audits, the sites that should be moved out of a service, and the service they should move to, if there is only one destination service, are listed. If only one group of sites is listed, a new service is created and the sites are moved to the created service.

95.5.2 Correcting detected configuration problems

NOTICE

Service-disruption hazard

Modifying the nms-server.xml file can have serious consequences that can include service disruption.

To view the result in the server or client log, you need to enable the logging option in the nms-server.xml or nms-client.xml file. Contact your Nokia technical support representative before you attempt to modify the nms-server.xml file.

You can use an RCA audit to correct a detected configuration problem from the RCA Result tab on the network object properties form. The NFM-P lists the operations to fix the problem, as shown in the following figure.

Figure 95-2 Correction plan for a problem

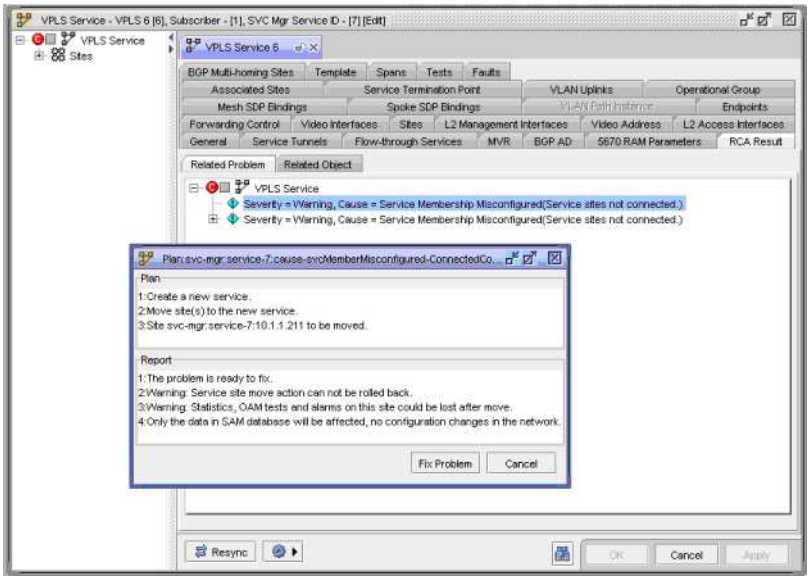

When you accept the proposed solution, a summary of the correction operation that the NFM-P implemented appears, as shown in Figure 95-3, “Correction report” (p. 3235). To view the result in the server or client log, you must enable the logging option in the nms-server.xml or nms-client.xml file.

Figure 95-3 Correction report

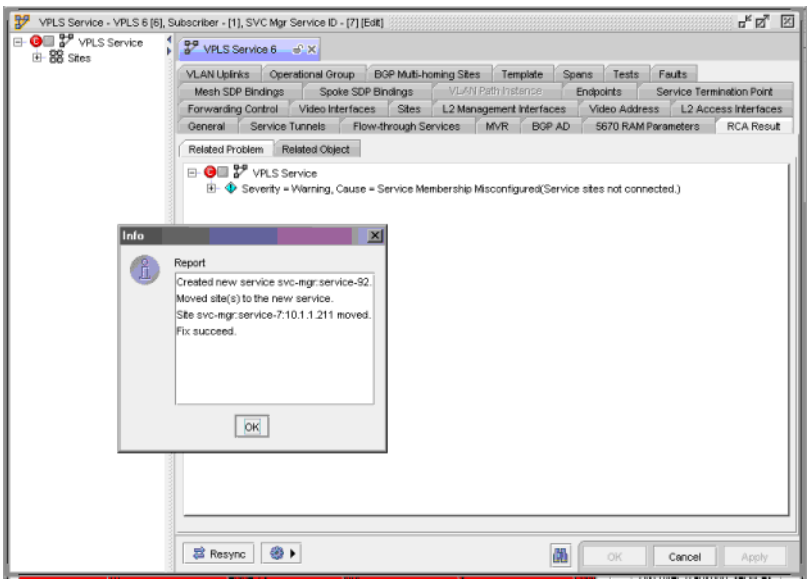

The following table describes the probable causes of problems.

Table 95-1 Probable causes

Probable cause	Description
Admin Down	The object is administratively down.
Aggregated	Aggregated cause
B-VPLS Backbone Attributes misconfigured	One or more B-VPLS backbone parameters are not configured correctly.
Control B-VPLS SPB ECT-FID mapping configurations are not consistent across all sites in this B-VPLS	The ECT algorithm to FID range mapping in a control SPB B-VPLS is not consistent on all of the sites in the same service. Each B-site must have the same number of ECT-to-FID mappings. In addition, the following ECT-to-FID range mapping parameters must be the same on all of the sites in the SPB-enabled B-VPLS: • FID Range Start • FID Range End • FID Range Algorithm See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) for more information.
Misconfiguration	There is a misconfiguration error.
Missed service members for Multi-Segment PW	Missed service members for Multi-Segment PW
No members configured for this Multi-homing site	The BGP VPLS multi-homing site does not have other sites as members that share the same multi-homing ID in order to comprise a group.
No valid Route Targets configured for VPLS BGP Multi-homing	No valid Route Targets currently exist for a VPLS site that has BGP multi-homing sites configured under it.
One or more SPB SAP(s) or SPB Spoke SDP Binding(s) exist on User B-VPLS that do not exist on Control B-VPLS. These extra links are not fate shared with the Control B-VPLS	One or more SPB SAPs or spoke SDP bindings on the user B-VPLS are not on the control B-VPLS. You must remove the SAPs or spoke SDP bindings, or create them on the control SPB B-VPLS site.
Route Targets are misconfigured for VPLS BGP Multi-homing	One of the following misconfigurations exists: • No matching RT for peered multi-homing sites • There are multiple RTs configured for a VPLS site that has BGP multi-homing sites configured under it. Multiple RTs will make the BGP multi-homing sites appear under multiple topologies or multiple services. • Multi-homing sites have the same multi-homing ID but different RTs (different RTs mean the sites are in different topologies or services)
Service Membership Misconfigured	The service membership is not configured correctly.
One or more SPB SAP(s) or SPB Spoke SDP Binding(s) configured on Control B-VPLS are missing from the User B-VPLS	One or more SPB SAPs or spoke SDP bindings are configured on the control B-VPLS that do not exist on this user B-VPLS. You must remove the SAPs or spoke SDP bindings, or create them on the user SPB B-VPLS site.

Table 95-1 Probable causes (continued)

Probable cause	Description
SPB is not enabled on one or more B-Sites in this B-VPLS	The SPB Mode parameter is not set to Control or User on one or more of the sites in the SPB-enabled B-VPLS service. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for more information.
SPB Sites within this B-VPLS service do not have their SPB Modes (Control/User) consistently configured	One or more B-sites in the SPB-enabled B-VPLS service do not have the same value configured for their SPB Mode parameter as the other sites within the service. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for more information.
The B-VPLS SPB FIDs are not identical across all sites in this B-VPLS service	The Forwarding Identifier (FID) parameter is not the same for all the sites in the SPB B-VPLS service. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) and 77.62 “To enable SPB on a user B-VPLS site” (p. 2326) for more information.
The Control B-VPLS SPB unicast forwarding tree topology configurations are not consistent across all sites in this B-VPLS	The Unicast Forwarding Tree Topology parameter is not configured consistently on all of the sites in the same control SPB B-VPLS service. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) for more information.
Underlying Resource Operational Down	The underlying resource of the object is operationally down.
Underlying Resource Admin Down	The underlying resource of the object is administratively down.
Underlying Resource Missing	An underlying resource is missing.
Underlying Resource Problem	Problem with an underlying resource
Unknown	The probable cause could not be determined.
Within this Control B-VPLS, not all SAPs or Spoke-SDPs are SPB enabled	SPB is not enabled on one or more SAPs or spoke SDP bindings in the control B-VPLS service. You must add the SPB interface to the SAP or spoke SDP binding. See 77.61 “To enable SPB on a control B-VPLS site” (p. 2323) for more information.

Procedures to configure and schedule an RCA audit

95.6 Workflow to configure and schedule an RCA audit

95.6.1 Stages

1

Create or configure an RCA audit policy. See 95.7 “To configure an RCA audit policy” (p. 3238) .

2

Run the RCA audit policy for a specific service or object.

a. Perform an RCA audit of a service or multiple services. See 95.8 “To perform an RCA audit of a service or multiple services” (p. 3240) .

b. Perform an RCA audit of a physical link. See 95.9 “To perform an RCA audit of a physical link” (p. 3242) .

3

As required, schedule the RCA audit policy to perform the RCA audit at a designated time; see 95.10 “To schedule an RCA audit” (p. 3243) .

4

Identify the problems and view the suggested solutions. Solutions are not provided for physical links. See “Viewing and analyzing RCA audit results” for more information.

5

Implement the changes, as required.

6

Delete an RCA audit policy, as required. See 95.11 “To delete an RCA audit policy” (p. 3245).

95.7 To configure an RCA audit policy

95.7.1 Steps

1

Choose Policies→RCA Audits from the NFM-P main menu. The RCA Audits form opens.

2

Select Audit Policy (RCA) from the drop down list.

3

Click Create or choose a policy and click Properties. The Audit Policy (Create|Edit) form opens with the General tab displayed.

4

Configure the general parameters.

5

Select a policy type and click OK.

6

Click Apply.

7

Click on the Entry tab. Depending on the option specified in Step 5 , a list of RCA audit policy entries is displayed.

The RCA audit policy entry for a VLL RCA audit is: RCA Audit for VLL Service Membership.

The RCA audit policy entry for a VPRN RCA audit is: RCA Audit for VPRN Service Membership.

The following are the RCA audit policy entries for a VPLS RCA audit:

- RCA Audit for B-VPLS PBB
- RCA Audit for VPLS Service Membership
- RCA Audit for BGP Multi-homing
- RCA Audit for B-VPLS SPB

The following are the RCA audit policy entries for a physical link RCA audit:

- RCA Audit For Physical Ports of Physical Links
- RCA Audit For Ethernet Port Specifics of Physical Links

8

Choose an entry from the list and click Properties. The Audit Policy Entry - RCA Audit Policy - RCA Audit For *Network_Object* (Edit) form opens.

9

Configure the Enabled and Remove Empty Service parameters.

10

To configure an RCA audit policy for a physical link:

1. Click on the Attributes tab. A list of default attributes for the physical link entry is displayed.
2. Perform one of the following:
   - To configure default attributes, go to 3 .
   - To add an attribute, go to 5 .
3. Activate or deactivate the RCA audit for each attribute.

4. Choose one of the problem severity options for each attribute by clicking on the entry in the Severity column of the attribute.
5. Click Add to add an attribute. The Adding new Attribute(s) - RCA Audit For *entry_type* of Physical Links form opens with a list of attributes associated with the physical link entry.
6. Choose one or more attributes in the list and click OK.

11

Save your changes and close the form.

END OF STEPS

95.8 To perform an RCA audit of a service or multiple services

95.8.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Choose a service or services that you want to audit and click RCA Audit. The Select RCA Policy to run the audit form opens with a list of configured policies.

Note: In addition to the standard audit checks, for BGP VPLS multi-homing sites, the audit also verifies:
Perform the following:

- if no valid Route Targets are configured under a VPLS site that has BGP multi-homing sites configured under it.
- if there are no matching RTs for peered multi-homing sites
- if there are multiple RTs configured under a VPLS site that has BGP multi-homing sites configured under it. Multiple RTs will make the BGP multi-homing sites appear under multiple topologies (or multiple services, since the RT defines the service)
- if multi-homing sites have the same multi-homing ID but different RTs (different RTs mean the sites are in different topologies or services)
- if a multi-homing site does not have other sites as members (sharing the same multi-homing ID to comprise a group)

In addition to the standard audit checks, for SPB- enabled B-VPLS sites, the audit also verifies that:

- all user B-VPLS SAPs or spoke SDP bindings are fate-shared with a control B-VPLS
- control and user B-VPLS instances in the same service on different sites have the same FID
- FID-to-ECT mapping is consistent on all of the sites for the active SPB B-VPLS services

- control B-VPLS SPB level Unicast Forwarding Tree Topology parameter is consistent on all of the sites in the same service

3

Choose an RCA audit policy and click OK. The Select RCA Policy form closes and a dialog box appears with the audit information.

4

Click OK.

5

If you selected multiple services to audit, perform the following. Otherwise, go to Step 6 .

1. Scroll to the RCA Audit Problem column and set the filter to = true.
2. Click Search. A list of services with problems appears.
3. Select a service and click Properties. The *Service* (Edit) form appears with the General tab displayed.

6

Click on the RCA Result tab. The RCA audit information is displayed on the Related Problem tab.

Note: The RCA Result tab appears only if the RCA audit detects one or more problems.

7

If required, expand the object in the problems tree to view the problems, the associated severity, and the cause.

8

Double-click on a problem icon to view information about the problem. The Problem (Edit) form opens with the General tab displayed.

9

Configure the Disable Fix Window parameter, if applicable.

10

Click on the Related Problem tab to view related problems.

Note: The NFM-P does not provide a solution for physical link configuration errors.

11

Click on the Caused By Objects tab to view a list of objects that are causing the problem.

12

Choose an entry from the list and click Properties. The properties form for the object opens.

13

Click on the tabs to view information about the configuration.

14

Configure the parameters, as required.

15

Save the configuration and close the forms.

16

To fix a problem:

1. Click on the Related Problem tab.
2. Right-click on a problem icon and choose Fix Problem. The Plan form appears.

   Note:

   The Fix Problem menu option is disabled if you set the Disable Fix Window parameter to Enabled in Step 9 .
3. Check the recommended plan to fix the problem in the Plan panel.
4. Check the report about fixing the problem in the Report panel.
5. Click Fix Problem.

17

Close the forms.

END OF STEPS

95.9 To perform an RCA audit of a physical link

95.9.1 Steps

1

Choose Manage→Equipment→Equipment from the NFM-P main menu. The Manage Equipment form opens.

2

Select Physical Link (Network) or Discovered Physical Link (Network) from the Select Object Type menu.

3

Choose an entry from the list and click Properties. The Physical Link (Edit) form opens with the General tab displayed.

4

Check the following indicators:

- RCA Audit Problem(s)
- Last Audit Time

5

Perform Step 6 to Step 15 in 95.8 “To perform an RCA audit of a service or multiple services” (p. 3240) .

END OF STEPS

95.10 To schedule an RCA audit

95.10.1 Steps

Note: Only VPLS, VPRN, and VLL services RCA audits and golden configuration RCA audits can be scheduled.

1

Choose Policies→RCA Audits from the NFM-P main menu. The RCA Audits form opens.

2

Choose a service audit policy and click Properties. The Audit Policy (Edit) form opens with the General tab displayed.

3

Click Schedule and choose Create RCA Audit Schedule Task. The Service Audit Scheduled Task (Create) form opens.

4

Configure the required parameters.

5

Click Select in the Schedule panel. The Select Schedule - Service Audit Scheduled Task form opens.

6

Perform one of the following:

a. Create a schedule to associate with the scheduled task.

b. Associate an existing schedule to the scheduled task. Go to Step 10 .

7

Click Create. The NFM-P Schedule (Create) form opens.

8

Configure the required parameters.

The Current Client End Time parameter is configurable when the Ongoing parameter is disabled and the Frequency parameter value is set to something other than Once.

When an NFM-P Schedule is not Ongoing and is assigned to a task, the NFM-P raises an alarm when the Current Client End Time expires.

Note: The Run Every parameter is not configurable when the Ongoing parameter is enabled.

9

Save your changes. The NFM-P Schedule (Create) form closes and the Select Schedule - Service Audit Scheduled Task form reappears.

10

Choose an entry from the list and click OK. The Select Schedule - Service Audit Scheduled Task form closes and the Service Audit Scheduled Task (Create) form reappears.

11

Ensure that the Administrative State parameter is set to Enabled and click OK. The Service Audit Scheduled Task (Create) form closes. The RCA Audit Policy form reappears with additional tabs.

12

To apply a filter to the schedule:

1. Click on the Schedule Filter tab.
2. Select a filter in the Audit Filter panel.

   Note:

   If no filters appear, click Audit Filter to create a new filter. See Chapter 1, “NFM-P GUI” for more information about how to create a filter.
3. Save your changes. The Select Audit Filter - Audit Policy form closes and the Audit Policy (Edit) form reappears with additional tabs.

13

Click on the Scheduled Objects to be Audited tab to view the services that are included in the audit.

14

Click on the Scheduled Result tab to view the results of the scheduled audit.

15

Click on the Problems tab to view the problems associated with the audit.

16

Close the forms.

END OF STEPS

95.11 To delete an RCA audit policy

95.11.1 Object deletion sequence

The following dependent objects must be deleted, in sequence, before you can delete an RCA audit policy:

- an associated scheduled task, if applicable
- any problems identified during the audit, if applicable

95.11.2 Steps

1

To delete the associated scheduled task:

1. Choose Tools→Schedules→Scheduled Task from the NFM-P main menu. The Scheduled Task form opens.
2. Choose NFM-P Scheduled Task (Schedule) from the Select Object Type menu.
3. Select the scheduled task, choose Task Action→Shut Down, and click Delete.
4. Confirm your changes and close the forms.

2

To delete the problems identified during the audit:

1. Choose Policies→RCA Audits from the NFM-P main menu. The RCA Audits form opens.
2. Choose Audit Policy (RCA) from the drop-down menu.
3. Select the RCA audit policy to be deleted, and click Properties. The Audit Policy (Edit) form opens.
4. Click on the Problems tab, select all of the problems, and click Delete.
5. Confirm your changes and close the forms.

3

To delete the RCA audit policy:

1. Choose Policies→RCA Audits from the NFM-P main menu. The RCA Audits form opens.
2. Choose Audit Policy (RCA) from the drop-down menu.
3. Select the RCA audit policy to be deleted, and click Delete.
4. Confirm your changes and close the forms.

END OF STEPS

96 Service throughput configuration

96.1 Overview

96.1.1 Purpose

This chapter describes the service throughput configuration, and provides information about configuring service throughput.

96.1.2 Contents

Service throughput configuration overview

96.2 Service throughput concepts

96.2.1 Throughput configuration description

The NFM-P allows you to create a copy of the SAP configuration of an end-user service that emulates the bandwidth, throughput, and QoS requirements of the service. You can perform any required OAM diagnostics on the service while sending the traffic through the NEs, to compare the customer SLA with current measurements and report the results.

96.2.2 Throughput configuration support

Throughput configurations are supported as shown in the following table.

NEs	Service types	Notes
• 7210 SAS • 7705 SAR • 7750 SR	• VLL Epipe, Apipe, and Cpipe • VPLS • composite services that contain VLL service (Epipe, Apipe, or Cpipe), VPLS, or both	—
7250 IXR	• VLL Epipe, Apipe, and Cpipe • VPLS • composite services that contain VLL service (Epipe, Apipe, or Cpipe), VPLS, or both	When the 7250 IXR is used as a source, the supported test methods are SAP-to-SAP Function and Mac Swap Loopback SAP/SDP Function. When the 7250 IXR is used as a destination, only the SAP-to-SAP Function is supported, with the Swap SAP option.

You can create a SAP based on the configuration of an existing SAP to prepare the service throughput configuration, where the port under the new SAP connects to the test entry or a loopback port as the destination.

Consider the following when you prepare a throughput test environment:

- The SAP to be replaced must reside under a VLL Epipe, Apipe or Cpipe site, including PBB Epipes, or a VPLS site including I-VPLS and B-VPLS.
- The following SAPs are not supported for the throughput test:
  - CCAGs
  - LAGs
  - SAPs that reside on an Ethernet tunnel element
  - SAPs that reside on an Ethernet ring element
  - SAPs used by a mirror service
- E-OAM objects, such as MEPs, are not copied to the test SAP.

You can view existing service throughput configurations on the Throughput Configuration tab of a service. Services and sites with a service throughput configuration have the Modified for Throughput Test parameter enabled on the General tab of the service or site.

Procedures to prepare for and restore from a service throughput configuration

96.3 Workflow to prepare for and restore from a service throughput configuration

96.3.1 Stages

The following workflow describes the high-level tasks required to prepare a throughput configuration. The workflow assumes that an NFM-P end user is challenging the offered SLA.

1

Open the service on which you need to create the test environment.

2

Set the Current OLC State parameter to Maintenance. See the *NSP System Administrator Guide* for more information.

3

Prepare the configuration. See 96.4 “To prepare an Epipe, Apipe, Cpipe, VPLS, or composite service throughput configuration” (p. 3249) .

4

Perform OAM diagnostics, as required. See Chapter 89, “Service Test Manager” for more information about the NFM-P service test manager.

5

Restore the service to the original configuration. See 96.6 “To restore a service after a throughput configuration” (p. 3255) .

6

Set the Current OLC State parameter to In Service. See the *NSP System Administrator Guide* for more information.

96.4 To prepare an Epipe, Apipe, Cpipe, VPLS, or composite service throughput configuration

96.4.1 Before you begin

The SAP that is replaced for the throughput configuration must exist under a supported VLL or VPLS site.

A throughput configuration cannot be performed on a SAP that:

- resides on CCAG or LAG

- resides on Ethernet tunnel element
- resides on Ethernet ring element
- is used by a mirror service

Nokia recommends that you set the Current OLC State parameter of the service to Maintenance before you prepare a throughput configuration. See the *NSP System Administrator Guide* for more information.

96.4.2 Steps

1

Perform one of the following:

a. To prepare a service throughput configuration on an Epipe, Apipe, Cpipe, or VPLS:
   1. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
   2. Choose a VPLS or VLL Epipe|Apipe|Cpipe and click Properties. The VPLS Service (Edit) or Epipe|Apipe|Cpipe (Edit) form opens.

b. To prepare a service throughput configuration on a composite service:
   1. Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Composite Services form opens.
   2. Choose a composite service and click Properties. The Composite Service (Edit) form opens.

2

Click on the Throughput Configuration tab and click Create. The Service Throughput Configuration form opens.

Note: Alternatively, you can choose a service throughput-capable service from the Manage Services or Manage Composite Services form, click on the Throughput Configuration button, and choose Create from the contextual menu.

3

Configure the required parameters.

4

In the Traffic Test Entry panel, select a SAP that is configured on the service that is to be replaced for testing.

5

Select a physical port or channel in the Traffic Test Entry panel that is connected to the test entry on which to create the SAP. You can choose an access or hybrid port on the service site on which the SAP that you selected in Step 4 resides that has the same encapsulation type as the selected SAP.

By default, the displayed port is the underlying port of the selected SAP. If you choose a different port, the NFM-P creates a SAP on the selected port by copying the configuration from the original SAP during the test preparation. E-OAM objects are not copied to the SAP.

6

Select a traffic destination site in the Traffic Test Destination panel.

7

To replace an existing SAP with a new SAP on a different port:

1. Set the Option parameter to Swap SAP.
2. Select a SAP in the Traffic Test Destination panel.
3. Select the port to which you need to switch in the Traffic Test Destination panel. The port must have the same encapsulation type as the selected SAP.

8

To loopback a port used by an existing SAP:

1. Set the Option parameter to Loopback Port Used By SAP.
2. Select the destination SAP in the Traffic Test Destination panel and configure the parameters on the port in the Loopback Parameters panel. The available parameters change depending on the NE on which the destination SAP is configured.

9

To replace an existing SAP with a new SAP on a different port and loopback the new port:

1. Set the Option parameter to Swap SAP and Loopback New Port.
2. Select the destination SAP and loopback port to which you need to swap in the Traffic Test Destination panel and configure parameters on the port. The available parameters change depending on the NE on which the destination SAP is configured.

10

To loopback an existing SAP:

1. Set the Option parameter to Loopback SAP.
2. Select the destination SAP in the Loopback SAP panel and configure the parameters in the Loopback Parameters panel.

11

To loopback an existing SDP binding:

1. Set the Option parameter to Loopback SDP Binding.
2. Select an SDP binding in the Loopback SDP Binding panel and configure the parameters in the Loopback Parameters panel.

12

Save the changes and close the form.

13

Click Apply on the VPLS Service (Edit) or Epipe|Apipe|Cpipe Service (Edit) form to:

- create a SAP on the selected port if the selected port differs from the port on the SAP
- configure the SAP with the same configuration as the original SAP, including all child objects that are copied over to the SAP
- delete the original SAP
- propagate port loopback parameter changes to the test port
- change the test SAP to a loopback SAP with the configured loopback parameters
- change the test SDP binding to a loopback SDP binding with the configured loopback parameters
- update the throughput information on the site
- update the throughput information on the service

Note: A site can only be used in one throughput configuration at a time. If a throughput configuration exists that uses a site as either a source or a destination, that site cannot be used in another throughput configuration.

14

Perform OAM diagnostics.

15

Save the changes and close the forms.

END OF STEPS

96.5 To configure an Epipe, Apipe, or Cpipe as a test service for MPLS-TP service tunnels

96.5.1 Preliminary considerations

Perform this procedure to create a site on a VLL Epipe, Apipe, or Cpipe service that is used exclusively for carrying test traffic for MPLS-TP service tunnels.

The admin lock test can be configured on a spoke SDP binding that is bound to a VLL Epipe, Apipe, Cpipe, and spoke termination on VPLS.

Consider the following:

- an MPLS-TP service tunnel must be used in the SDP binding, and the Control Word parameter on the Pseudowire OAM tab of the SDP binding must be set to Preferred
- the Control Channel Status parameter on the Control Channel tab of the spoke SDP binding must be set to Disabled

- a PW path must be created; see 76.15 “To configure an MPLS-TP static pseudowire on a VLL spoke SDP binding” (p. 2136)
- an ICB SDP binding must not be created on an endpoint
- no other service throughput configuration can exist on the site
- for Apipe and Cpipe services, the VC Type of the service and the test service must match

96.5.2 Steps

Create a test service site

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Create a VLL by clicking Create→VLL→Epipe|Apipe|Cpipe, or choose a VLL Epipe, Apipe, or Cpipe and click Properties. The Epipe|Apipe|Cpipe Service (Create|Edit) form opens.

3

If you are creating a service, configure the required general parameters.

4

On the navigation tree, right-click on the Sites icon and choose Create *Epipe*|*Apipe*|*Cpipe* Site. The Select Network Elements form opens.

5

Choose a site. The site you choose is also a service site on the host service of the throughput configuration you are configuring in Step 12 .

6

Click OK. The VLL Site (Create) form opens.

7

Configure the required general parameters.

8

Set the VLL Site Type to Terminating.

9

Enable the Test Service parameter .

Note: You must only create one site for a test service. This site defines a SAP that is the ingress and egress for test traffic.

The Test Service parameter is configurable only during site creation.

10

Save the changes and close the form.

Configure service throughput on the host service

11

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

12

Choose the VLL Epipe, Apipe, or Cpipe on which you want to configure the service throughput and click Properties. The Epipe|Apipe|Cpipe Service (Edit) form opens.

13

On the Epipe|Apipe|Cpipe Service (Edit) form of the host service, click on the Throughput Configuration tab and click Create. The Service Throughput Configuration form opens.

Note: Alternatively, you can choose a service throughput-capable service from the Manage Services form, click on the Throughput Configuration button, and choose Create from the contextual menu.

14

Configure the ID and Description parameters.

15

Set the Test Method parameter to MPLS-TP Lock Instruct Function.

Configure the Admin Lock test

16

A spoke SDP binding is administratively locked by locking the host service. Test traffic is injected into the spoke SDP binding using a test SAP, defined within the test service you created in Step 9 . When the admin lock is configured on a spoke SDP binding and a test service ID is specified, all of the traffic is forwarded to and from a SAP defined in the separate test service, which must be compatible with the spoke SDP binding. Traffic to and from the original SAP is dropped.

In the Traffic Test Entry panel, select a site on the host service. The Site ID is the site on which you are configuring the admin lock and the same site that you configured on the test service in Step 9 . The Admin Lock panel appears.

17

In the Admin Lock panel, select the spoke SDP binding to lock on the host site you selected in Step 16 .

18

In the Test Service panel in the Admin Lock panel, select the test service site you created in Step 5 to Step 10 .

Configure the Loopback test

19

If a loopback is configured on a spoke SDP binding, all traffic on the ingress direction of the spoke SDP binding and associated with the ingress VC-label is forwarded to the egress direction of the spoke SDP binding. The terminating site with the loopback configured to send test traffic back to the SDP binding you selected in Step 17 is automatically selected.

In the Traffic Test Destination panel, you can clear the selection or manually select the terminating, or destination, site.

20

In the Loopback panel, the spoke SDP binding on which you are configuring the loopback test is automatically selected. You can clear the selection or manually select an SDP binding.

21

Save the changes and close the form.

Note: A site can only be used in one throughput configuration at a time. If a throughput configuration exists that uses a site as either a source or a destination, that site cannot be used in another throughput configuration.

END OF STEPS

96.6 To restore a service after a throughput configuration

96.6.1 Service restoration considerations

The NFM-P performs the following when you restore the service:

- deletes the newly created SAPs
- re-creates the original SAPs
- on the loopback port: resets the loopback type to none if it was changed by the throughput configuration, and changes the loopback parameters to their default values
- updates the site throughput status information
- updates the service throughput status information

96.6.2 Steps

1

Perform one of the following:

a. Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

b. Choose Manage→Service→Composite Services from the NFM-P main menu. The Manage Services form opens.

2

Choose the service for which you need to restore the configuration and click Properties. The *Service_type* Service (Edit) form opens.

3

Click on the Throughput Configuration tab. A list of throughput configurations on the service appears.

4

Select the throughput configuration you need to remove, and click Delete.

Note: Alternatively, you can select a service or composite service that is modified for throughput configuration from the Manage Services form or Manage Composite Services, click on the Throughput Configuration button, and choose Delete (all) from the contextual menu. This removes all the service throughput configurations on the service.

5

Close the forms.

END OF STEPS

Part VIII: Appendices

Overview

Purpose

This part provides information about NFM-P user and system parameters that are not otherwise documented in the XML API Reference.

Contents

A Parameters

A.1 Overview

A.1.1 Purpose

This appendix lists and describes the NFM-P user and system parameters. Nodal parameters are not included in this appendix but can be found in the XML API Reference.

The parameters in this appendix are found on NFM-P forms and child forms that may be launched from the main menu or contextual menus related to the configurable object.

A.1.2 Contents

A.2 Alarm Settings parameters

A.2.1 Display Buffer Time (minutes)

The Display Buffer Time (minutes) parameter specifies the buffer time, in minutes, that will be applied to the time period for which current and historical alarms are displayed. The range is 1 to 59. The default is 5.

A.2.2 Reason for change

The Reason for change parameter specifies the justification for editing an existing note to an alarm. The options are:

- unspecified (default)
- unknown
- comment
- clarification
- update
- confirmation
- correction
- dispute

A.2.3 Search for Alarms

The Search for Alarms parameter specifies the search for historical alarms received during a specified time period. The Oldest Alarm Searched and the Oldest Alarm in Database fields are updated.

The options are:

- Past 1 hour
- Past 12 hours
- Past 1 Day
- Past 7 days
- Past 30 days
- All Available

A.2.4 Show Additional Text Button on Properties Forms

The Show Additional Text Button on Properties Forms parameter specifies whether the Additional Text Find icon is displayed on the object property forms, based on a client session.

The options are:

- Enabled
- Disabled (default)

A.3 Discovery Manager parameters

A.3.1 Ignore Timestamps

The Ignore Timestamps parameter specifies whether entries with unchanged last change timestamps are processed.

The options are:

- Enabled
- Disabled (default)

A.4 Generic NE Manager parameters

A.4.1 Use Default Additional Text

The Use Default Additional Text parameter specifies whether the alarm contains only the default additional text. When the parameter is enabled, the Additional Text parameter is set to the default value and cannot be configured.

The options are:

- Enabled (default)
- Disabled

A.5 NE CPM Filter parameters

A.5.1 MAX

The MAX parameter specifies whether the PIR (kbps) parameter or CIR (kbps) parameter is set to infinity or can be configured.

The options are:

- disabled (default)
- enabled

When the MAX parameter is set to enabled, you cannot configure the CIR or PIR parameters.

A.5.2 Service Id

The Service Id parameter specifies the service ID that is applied as a MAC filter match criterion in an NE CPM or management access filter.

A.5.3 Source IP Prefix

The Source IP Prefix parameter specifies a list of IP prefixes to use as sources for the filter policy. When the parameter is enabled, click on the Select button and choose a prefix list.

A.6 NE Maintenance parameters

A.6.1 Command to Apply After Backup

The Command to Apply After Backup parameter specifies whether to certify the configuration after a successful backup. The parameter is configurable only on an OmniSwitch.

The options are:

- Certify
- No command applied (default)

A.7 NFM-P User Security parameters

A.7.1 Enable

The Enable parameter specifies whether to enable or disable the ability to configure the Maximum Sessions Allowed parameter.

The options are:

- enabled
- disabled (default)

A.7.2 Password Change Required

The First Time Login Password Change parameter specifies whether a user is prompted to enter a new password before they can login for the first time.

The options are:

- disabled (default)
- enabled

A.7.3 Threshold Reporting State

The Threshold Reporting State parameter specifies whether to generate threshold alarms when the number of log entries exceeds Max Log Record parameter value.

The options are:

- Up (default)
- Down

A.7.4 User inactive greater than or equal to

The User inactive greater than or equal to parameter specifies the NFM-P user inactivity period, in days, for a custom user inactivity search. The range is 0 to 365. There is no default.

A.8 System Preferences parameters

A.8.1 Changing system preferences

CAUTION

Service Disruption

Changing the parameter value of a System Preference may affect the behavior of an existing NFM-P service.

Do not change the parameter value from the default without contacting Nokia technical support.

Note: The System Preferences parameters are configurable only by a user with administration privileges.

A.8.2 Create CSV files using UTF-8 with BOM (byte order mark) encoding

The Create CSV files using UTF-8 with BOM (byte order mark) encoding parameter specifies whether a BOM is added to CSV files. This is used to view CSV files in multi-language environments. By default, when the parameter is disabled, no BOM is added to CSV files. When the parameter is enabled, UTF-8 BOM will be added to all files created.

The options are:

- disabled (default)
- enabled

A.8.3 Maximum Number of Golden Config Audit Problems

The Maximum Number of Golden Config Audit Problems parameter specifies the maximum number of problems that a golden configuration audit can identify before the NFM-P automatically aborts the audit. The parameter is designed to preserve NFM-P system resources when comparing two extremely divergent NE configurations that may have been selected accidentally. The default is 500.

A.8.4 Propagate Name and Description of Site to Service

The Propagate Name and Description of Site to Service parameter specifies the default behavior when a site is added to a service with no other sites configured or a service with one site is found during node discovery. When the parameter is enabled, the service inherits the name and description parameters of the site configured on the service. Services with multiple sites or services with names and descriptions that are manually configured are not affected. The parameter is mutually exclusive with the A.8.6 “Whenever a site is added to a service in SAM, propagate the Service Name to Site Name” (p. 3266) parameter. The options are:

- disabled (default)
- enabled

A.8.5 Tabs

The Tabs parameter specifies the default behavior for tab customization on configuration forms.

Table A-1 Tabs parameter

Option	Description
Allow NFM-P to pre-hide some tabs on forms. Customization Allowed. (default)	Some tabs are automatically hidden by default. The user can customize the hidden tabs with the tab selector.
Always show all tabs on forms. No Customization Allowed.	All tabs are displayed by default. The tab selector is not available on configuration forms.
Initially show all tabs on forms. Customization Allowed.	All tabs are displayed by default. The user can customize the hidden tabs with the tab selector.

A.8.6 Whenever a site is added to a service in SAM, propagate the Service Name to Site Name

(propagateServiceNameToSites)

The “Whenever a site is added to a service in NFM-P, propagate the Service Name to Site Name” parameter specifies the default behavior of whether to populate the site name property on an NE using the service name when the NE is added as a service site. When the parameter is disabled, the site name property on the NE is unpopulated after the NE is added as a service site. The options are:

- disabled (default)
- enabled

A.9 Manage Workspaces parameters

A.9.1 Finished adding New Menu Items

When enabled, the Finished adding New Menu Items parameter specifies that the New Menu Items option no longer displays on the Menus tab of the Workspace (Edit) form, and that the Finished adding New Menu Items parameter also no longer displays. These functions remain inaccessible until a new upgrade is performed.

The options are:

- Enabled
- Disabled (default)

A.9.2 Overwrite Existing Workspace(s)

When enabled, the Overwrite Existing Workspace(s) parameter specifies that an imported workspace overwrites an existing workspace with the same name and scope.

The options are:

- Enabled
- Disabled (default)

A.9.3 Overwrite User Name

When enabled, the Overwrite User Name parameter specifies that the username of the currently logged-in user is assigned to the imported workspace. When the parameter is disabled, the username associated with the saved workspace is retained.

The options are:

- Enabled
- Disabled (default)

A.9.4 Scope

The Scope parameter on the Workspace form indicates the ownership of the custom workspace.

The options are:

- Private
- Public

A private workspace is visible only to the user that created the workspace, and to an administrator.

A public workspace is visible to all users.

A.9.5 Show Current OLC State

When enabled, the Show Current OLC State parameter specifies that navigation tree objects below the level of the NE display their current OLC state in their labels.

The options are:

- Enabled
- Disabled (default)

See the *NSP System Administrator Guide* for more information.

A.9.6 Tab Preferences

The Tab Preferences parameter specifies how tab preferences are derived for configuration forms in the workspace.

The options are:

- Local (default)
- Custom

Select an option by choosing one of the three buttons on the form.

Table A-2 Tab Preferences configuration buttons

Button	Description
Set to Local	Tab preferences are Local, and are derived from local tab selector settings for the current user. The workspace has no effect on tab display settings. Buttons on the tab selector are enabled and tab preferences can be changed manually.
Set to Current	Tab preferences are Custom, and are derived from the workspace. The workspace saves and uses the tab display settings in effect for the GUI client when Set to Current is selected. Buttons on the tab selector are disabled.
Set from Workspace...	Tab preferences are Custom, and are derived from the workspace. The workspace saves and uses tab display settings copied from another selected workspace. Buttons on the tab selector are disabled.

A.10 Task Manager parameters

A.10.1 autoRefreshInterval

The autoRefreshInterval parameter specifies how often, in s, the Task Manager searches for new tasks when the Task Manager is open. The range is 0, or 5 to 600. The default is 20. A value of 0 means the parameter is disabled. The parameter does not take effect until the client is restarted.

A.10.2 failedTasksPurgeInterval

The failedTasksPurgeInterval parameter specifies how often, in min, to remove all of the tasks that are not in the In Progress state. The range is 0, or 5 to 10 080. The default is 1440. A value of 0 means the parameter is disabled.

A.10.3 maxNumRetainedTasks

The maxNumRetainedTasks parameter specifies the maximum number of monitored tasks that the Task Manager displays. The count includes the top-level tasks and all sub-tasks. When the value is reached, the system automatically deletes successful tasks, starting with the earliest. The deleted tasks appear in XML format in the TaskTracker.log file, which is located in the log/taskmgmt directory on the NFM-P server. The range is 200 to 50 000 tasks. The default is 10 000.

A.10.4 numTasksToPurgeWhenFull

The numTasksToPurgeWhenFull parameter specifies the number of successful tasks to remove when the limit specified by the A.10.3 “maxNumRetainedTasks” (p. 3268) parameter is reached. The range is 20 to 500 tasks. The default is 100.

A.10.5 successfulTasksPurgeInterval

The successfulTasksPurgeInterval parameter specifies, in min, the interval at which tasks that are in the Succeeded state are removed. For a specific interval, the tasks from the previous interval that have a Succeeded state are removed. The range is 0, or 2 to 2880. The default is 10. A value of 0 means the parameter is disabled.

A.11 User Preferences parameters

A.11.1 Access Interface Encap Value (Dot1q only)

The Access Interface Encap Value (Dot1q only) parameter specifies whether the Auto-Assign ID parameter is the default parameter for dot1q encapsulation.

The options are:

- Enabled
- Disabled (default)

A.11.2 Apply User Span of Control

The Apply User Span of Control parameters specifies whether the GUI automatically filters list forms, trees, and maps to display only the objects in the user Edit Access spans.

The options are:

- Enabled
- Disabled (default)

A.11.3 Browser Path

The browser path parameter specifies the file path to the browser application that the NFM-P launches as a web portal. The browser path parameter can be configured manually or by clicking the Browse button and navigating to the file path.

A.11.4 Command Helper Key

The Command Helper Key parameter allows you to designate the key that activates the command helper function in the Script Editor.

A.11.5 Debug STM Mode

The Debug STM Mode parameter provides access to additional forms on the STM that can be used to configure OAM diagnostic test limits and view additional test configuration information on managed devices. The default is disabled.

When the parameter is enabled, you can select the following additional menus on the Service Test Manager (STM) form:

- NE Test Agent (Assurance): Allows you to specify the maximum number of pings and traces allowed to be performed on individual managed devices during the execution of a test, or limit the number of tests that can be performed on all managed devices. Also provides information about the node ID, concurrent ping and trace information, LTT concurrent requests, deployed test counts, and alarm status.
- Deployed Test (Assurance): Provides information about the node ID, name and description, management ownership, deployment and execution state, schedule status and start time, and alarm status.
- NE Schedulable Test (Assurance): Provides the node ID, name and description, management

ownership, test deployed, mode, state, runs and failures, accounting policy ID, accounting suppression, card TCA profile and alarm status.

A.11.6 Default Client Time Zone

The Default Client Time Zone parameter specifies the time zone that is applied to the NFM-P client by default. Choose a time zone from the drop-down menu. The default is the NFM-P server installation time zone. The Current Client Time Zone will automatically reflect any change made to the Default Client Time Zone.

A.11.7 Default Polling Interval (seconds)

The Default Polling Interval parameter specifies the default interval, in seconds, for the polling interval for real-time statistics in the Statistics Plotter form. The range is 10 to 3600. The default is 10.

A.11.8 Enable Command Helper

The Enable Command Helper parameter specifies whether the Command Helper function in the Script Editor is enabled.The Command Helper is a one-key function which completes commonly occurring script commands. The command helper key can be designated by the user.

A.11.9 Enable Confirmation for Bulk Change Actions

The Enable Confirmation for Bulk Change Actions parameter specifies whether a confirmation message is displayed before the NFM-P carries out a bulk change operation.

The options are:

- false
- true (default)

A.11.10 Enable properties forms with navigation trees when ‘View Alarmed Object’ button is applied

The Enable properties forms with navigation trees when ‘View Alarmed Object’ button is applied parameter, allows you to specify whether alarm information is displayed when you click on the View Alarmed Object button. The disabled option allows you to view the Properties form for an alarmed object.

The options are:

- Enabled (default)
- Disabled

A.11.11 GUI Builder in the Editor

The GUI Builder in the Editor parameter specifies whether the Velocity GUI Builder in the Script Editor is enabled.

The options are:

- Enabled (default)
- Disabled

A.11.12 Maximum Data Retention Time (seconds)

The Maximum Data Retention Time (Seconds) parameter specifies the number of seconds to keep statistics data in the Statistics Plotter form. The range is 3600 to 86400. The default is 43200.

A.11.13 Save To File Default Extension

The Save To File Default Extension parameter specifies the file type that is selected as default when saving list form information to file.

The options are:

- CSV (default)
- HTML

A.11.14 Show Alarm Flags

The Show Alarm Flags parameter specifies whether the monitoring flag panel toolbar at the top of the Dynamic alarm list Alarm Window is displayed.

The options are:

- Enabled
- Disabled (default)

A.11.15 Show Correlated Alarms

The Show Correlated Alarms parameter specifies whether correlated alarms are displayed in the alarm window. The parameter is automatically enabled when you enable the A.11.16 “Show Correlated Alarms by Default” (p. 3272) parameter. You can configure the parameter on a per-session basis.

The options are:

- Enabled (default)
- Disabled

A.11.16 Show Correlated Alarms by Default

The Show Correlated Alarms by Default parameter specifies whether the A.11.15 “Show Correlated Alarms” (p. 3271) parameter is enabled by default.

The options are:

- Enabled (default)
- Disabled

A.11.17 Show Toolbar

The Show Toolbar parameter specifies whether the toolbar at the top of the window is displayed.

The options are:

- Enabled (default)
- Disabled

A.11.18 Specify # of Items Per Page

The Specify # of Items Per Page parameter specifies the number of items returned per page by a search. The range is 1 to 9999. The default is 1000.

A.11.19 Suppress Containing Window Warning

The Suppress Containing Window Warning parameter specifies whether containing window warnings are suppressed.

The options are:

- Enabled
- Disabled (default)

When a child object configuration form is launched from a parent object, and the child object configuration is changed, a warning message opens. The warning message indicates that changes to the child form are not committed until they are applied in the parent object. You must acknowledge the message. You can suppress this warning message by enabling the parameter on the User Preferences form. You continue to receive a warning message that changes must be applied for parent objects before you can close the parent object configuration form.

A.11.20 Suppress Template Execution Warning

The Suppress Template Execution Warning specifies whether warning messages are suppressed when you execute a template.

The options are:

- Enabled
- Disabled (default)

A.11.21 Suppress Template Generation Message

The Suppress Template Generation Message parameter specifies whether template generation windows are suppressed.

The options are:

- Enabled
- Disabled (default)

A.11.22 Suppress Service and Composite Service Map Load Warning

The Suppress Service and Composite Service Map Load Warning parameter specifies whether the load time warnings for service and composite service maps are suppressed.

The options are:

- Enabled
- Disabled (default)

A.11.23 Turn on Audible Alarms

The Turn on Audible Alarms parameter specifies whether there is an alarm bell when an incoming alarm is registered.

The options are:

- Enabled (default)
- Disabled

A.12 Common Create menu parameters

A.12.1 Aggregated Service Site Operational State

The Aggregated Service Site Operational State parameter cannot be configured. The value is derived from the operational states of the sites that are part of the service.

The values are:

- Up—All sites are operationally up
- Partially Down—At least one site is operationally down
- Down—All sites are operationally down
- Unknown—The service has no provisioned sites

When the Aggregated Service Site Operational State is Partially Down or Down, a check mark appears beside the appropriate State Cause indicator to identify the type of fault to the NFM-P admin or mirror service management user. You can view alarms on the Faults page.

A.12.2 Enable

The Enable parameter specifies whether the Lease Populate parameter is configurable.

The options are:

- Enabled
- Disabled (default)

A.12.3 Enable WLAN GW

The Enable WLAN GW parameter specifies whether or not WLAN GW functionality is enabled on a routing instance. This parameter must be enabled in order to configure an APN Network Identifier or Mobile GW Address Map.

The options are:

- disabled (default)
- enabled

A.12.4 IPv6 Allow Unmatching Prefixes

The IPv6 Allow Unmatching Prefixes parameter specifies whether the interface allows unnumbered IPv6 prefixes that do not match the list of numbered prefixes configured on the NE.

The options are:

- enabled
- disabled (default)

A.12.5 Loopback Type

The Loopback Type parameter specifies the loopback configuration of the destination port in the service throughput test.

The options are:

- Internal
- Line
- None (default)

The options that are displayed depend on the NE on which you are configuring the port.

A.12.6 Loopback Time (seconds)

The Loopback Time (seconds) parameter specifies how long the selected destination port in the service throughput test remains as loopback. The parameter is configurable if the port is capable for timed loopback. The range is 30 to 86 400. The default is 0 (disabled).

A.12.7 No Egress Aggregate Rate Limit

The No Egress Aggregate Rate Limit parameter specifies that the transmission rate of all egress queues for the access interface is unlimited.

The options are:

- Enabled (default)
- Disabled

A.12.8 SHCV Enabled

The SHCV Enabled parameter specifies whether SHCV is enabled on the interface.

The options are:

- Enabled
- Disabled (default)

A.12.9 Template Description

The Template Description parameter specifies a description for the template type. The range is an interface name of 0 to 80 characters. The default is an empty string.

A.13 Physical Link parameters

A.13.1 Bandwidth (Mbps)

The Bandwidth parameter specifies the maximum bandwidth allowed for each specified CoS. The default value is derived from the speed of the physical port. Because there are eight classes of service, the default value for each CoS is the speed of the port divided by eight. The XML values are shown in Table A-3, “Bandwidth (Mbps) parameters” (p. 3274) .

Note: This parameter is only available if service CAC is configured. See the *NSP System Administrator Guide* for more information.

Table A-3 Bandwidth (Mbps) parameters

Bandwidth (Mbps) parameter name	XML string	Maximum bandwidth allowed for:
Bandwidth (Mbps) first field in column	cos0BW	CoS 0
Bandwidth (Mbps) second field in column	cos1BW	CoS 1
Bandwidth (Mbps) third field in column	cos2BW	CoS 2
Bandwidth (Mbps) fourth field in column	cos3BW	CoS 3
Bandwidth (Mbps) fifth field in column	cos4BW	CoS 4
Bandwidth (Mbps) sixth field in column	cos5BW	CoS 5
Bandwidth (Mbps) seventh field in column	cos6BW	CoS 6
Bandwidth (Mbps) eighth field in column	cos7BW	CoS 7

A.13.2 Booking Factor (%)

The Booking Factor parameter specifies the actual amount of bandwidth being booked on the physical link. For example, if a service requests 10 Mbps on CoS 0 and the booking factor on CoS 0 is 50%, then the actual amount of booked bandwidth is 5Mbs. The XML values are shown in Table A-4, “Booking Factor (%) parameters” (p. 3275) .

Note: This parameter is only available if service CAC is configured. See the *NSP System Administrator Guide* for information.

Table A-4 Booking Factor (%) parameters

Booking Factor (%) parameter name	XML string	Bandwidth being booked for:	Default value
Booking Factor (%) first field in column	cos0BookFactor	CoS 0	100
Booking Factor (%) second field in column	cos1BookFactor	CoS 1	100
Booking Factor (%) third field in column	cos2BookFactor	CoS 2	100
Booking Factor (%) fourth field in column	cos3BookFactor	CoS 3	100
Booking Factor (%) fifth field in column	cos4BookFactor	CoS 4	100
Booking Factor (%) sixth field in column	cos5BookFactor	CoS 5	100
Booking Factor (%) seventh field in column	cos6BookFactor	CoS 6	100
Booking Factor (%) eighth field in column	cos7BookFactor	CoS 7	100

A.13.3 Utilization Threshold (%)

The Utilization Threshold parameter specifies the percentage of bandwidth that, when exceeded, raises an alarm indicating that the used bandwidth has exceeded the maximum allowed bandwidth on the link. For example, if the threshold is set to 75% and the maximum allowed bandwidth is 20 Mbps, then an alarm will be raised when the used bandwidth value exceeds 15 Mbps. The XML values are shown in Table A-5, “Utilization Threshold (%) parameters” (p. 3276) .

Note: This parameter is only available if service CAC is configured. See the *NSP System Administrator Guide* for information.

Table A-5 Utilization Threshold (%) parameters

Utilization Threshold (%) parameter name	XML string	Bandwidth currently being used for:
Utilization Threshold (%) first field in column	cos0BWThreshold	CoS 0
Utilization Threshold (%) second field in column	cos1BWThreshold	CoS 1
Utilization Threshold (%) third field in column	cos2BWThreshold	CoS 2
Utilization Threshold (%) fourth field in column	cos3BWThreshold	CoS 3
Utilization Threshold (%) fifth field in column	cos4BWThreshold	CoS 4

Table A-5 Utilization Threshold (%) parameters (continued)

Utilization Threshold (%) parameter name	XML string	Bandwidth currently being used for:
Utilization Threshold (%) sixth field in column	cos5BWThreshold	CoS 5
Utilization Threshold (%) seventh field in column	cos6BWThreshold	CoS 6
Utilization Threshold (%) eighth field in column	cos7BWThreshold	CoS 7

A.14 Equipment Group parameters

A.14.1 Span

The Span parameter specifies whether span of control filtering is enabled.

Table A-6 Span parameter

Option	Description
Span Off (default, if span filtering is disabled on the User Preferences form)	Span of control filtering is disabled; objects in the View Access and Edit Access spans of the current user are displayed.
Span On (default, if span filtering is enabled on the User Preferences form)	Span of control filtering is enabled; only objects in the Edit Access spans of the current user are displayed.
User Preference	Span of control filtering is enabled or disabled, as configured on the User Preferences form.

A.15 IPsec VPN parameters

A.15.1 Delivery Service Interface Address

The Delivery Service Interface Address parameter specifies the IPv4 address, in dotted-decimal format, of the delivery service interface. The subnet for the parameter must be the same subnet as the A.15.3 “Local Gateway Address” (p. 3277) parameter.

A.15.2 ISA-Tunnel Group

The ISA-Tunnel Group parameter specifies the ISA-Tunnel group for the IPSEC VPN.

A.15.3 Local Gateway Address

The Local Gateway Address parameter specifies the IPv4 address, in dotted-decimal format, of the local NE of the IPsec tunnel. The subnet for the parameter must be the same subnet as the A.15.1 “Delivery Service Interface Address” (p. 3277) parameter.

A.15.4 Remote Gateway Address

The Remote Gateway Address parameter specifies the IPv4 address, in dotted-decimal format, of the remote gateway.

A.15.5 Secure Service Interface Address

The Secure Service Interface Address parameter specifies the IPv4 address for the secure service interface. The parameter is mandatory when the A.15.8 “Tunnel Type” (p. 3278) parameter is set to the Dynamic (Soft Client) option.

A.15.6 Static Route Address

The Static Route Address parameter specifies the IPv4 address, in dotted-decimal format, of the static route.

A.15.7 Static Route Prefix

The Static Route Prefix parameter specifies the prefix for the static route.

Table A-7 Static Route Prefix parameter

Option	Option description
24	Choose when the A.15.8 “Tunnel Type” (p. 3278) parameter is set to Dynamic (Site-to-Site) or Dynamic (Soft Client)
32	Choose when the A.15.8 “Tunnel Type” (p. 3278) parameter is set to Static

A.15.8 Tunnel Type

The Tunnel Type parameter specifies the tunnel type for the tunnel group.

The options are:

- Dynamic (Site-to-Site)
- Dynamic (Soft Client)
- Static

A.16 IES parameters

A.16.1 IES parameters overview

This section describes the parameters on the IES creation forms.

Note: This section also describes parameters common to the Service Template form and child forms. Service templates are intended for use in the NFM-P GUI. Nokia does not recommend using the XML API XML classes and methods associated with service templates.

A.16.2 Inherit from Network Element

The Inherit from Network Element check box specifies whether the interface inherits the value of the Operational State Transition Interval parameter from the NE.

The options are:

- Enabled (default)

- Disabled

When the check box is disabled, the Operational State Transition Interval parameter is configurable.

A.17 VLL parameters

A.17.1 VLL parameters overview

This section describes the parameters on the VLL service creation form and child forms.

Note: This section also describes parameters common to the Service Template forms. Service templates are intended for use in the NFM-P GUI. Nokia does not recommend using the XML API XML classes and methods associated with service templates.

A.17.2 Auto-Generate ID

The Auto-Generate ID parameter specifies the value used by each end of a service tunnel to identify the VC. The range is 0 to 4 294 967 295.

The options are:

- Enabled (default)
- Disabled

A.17.3 No Revert

The No Revert parameter specifies whether the VLL returns the primary spoke SDP to service after a failure. When the parameter is enabled, the VLL does not return the primary spoke SDP to service after a failure.

The options are:

- disabled (default)
- enabled

A.18 VPLS parameters

A.18.1 VPLS parameters overview

This section describes the parameters on the VPLS creation form and child forms.

Note: This section also describes parameters common to the Service Template form and child forms. Service templates are intended for use in the NFM-P GUI. Nokia does not recommend using the XML API XML classes and methods associated with service templates.

A.18.2 Clear Forced Switchover

The Clear Forced Switchover button allows you to clear a manual switchover from a redundant spoke SDP binding back to an active spoke SDP binding that was previously initiated by using the A.18.5 “Force Switchover” (p. 3280) button. You must clear any such manually forced switchovers by using the Clear Forced Switchover button after the active spoke SDP binding has been restored.

The system does not switch over automatically to other active spoke SDP bindings if this is not done, even if the redundant spoke SDP binding subsequently goes down.

The Active State read-only display indicates the change in status for the SDP binding when you press the Clear Forced Switchover button.

A.18.3 Disable Revert Time (Infinite)

The Disable Revert Time (Infinite) parameter specifies whether or not to disable the Revert Time (seconds) parameter indefinitely.

The options are:

- enabled
- disabled (default)

A.18.4 Enable

The Enable parameter specifies whether to enable or disable the configurability of the Enable Lease Populate parameter. The parameter is configurable when the Snooping parameter is enabled.

The options are:

- enabled
- disabled (default)

When the parameter is enabled, you can configure the Enable Lease Populate parameter. Disabling the parameter is equivalent to configuring the Enable Lease Populate parameter with a value of 0.

A.18.5 Force Switchover

The Force Switchover button allows you to force a switchover of the active SDP binding to a redundant SDP binding under an endpoint of the VPLS. When you do this, the redundant SDP binding becomes active and the Active State indicator displays the status change.

Note: You must clear a manually forced switchover using the A.18.2 “Clear Forced Switchover” (p. 3279) button after the active spoke SDP binding is restored. The NFM-P is unable to switch automatically to other active spoke SDP bindings if this is not done, even if the redundant spoke SDP binding subsequently goes down.

A.18.6 Maximum Number of Sources

The Maximum Number of Sources parameter specifies the maximum number of IGMP sources in each group that can be statically or dynamically learned. The default for all devices is 0, which means that no limit is imposed. Reducing the value of the parameter below the number of learned sources does not remove already-learned sources, but does prevent new sources from being learned. The range is 0 to 1000.

A.19 VPRN parameters

A.19.1 VPRN parameters overview

This section describes the parameters on the VPRN service creation forms.

Note: This section also describes parameters common to the Service Template form and child forms. Service templates are intended for use in the NFM-P GUI. Nokia does not recommend using the XML API XML classes and methods associated with service templates.

A.19.2 Enable DNS

The Enable DNS parameter specifies whether or not DNS is configured on a routing instance. This parameter must be enabled in order to configure primary, secondary, and tertiary DNS addresses on the routing instance.

The options are:

- disabled (default)
- enabled

A.19.3 Inherit from Network Element

The Inherit from Network Element check box specifies whether the interface inherits the value of the Operational State Transition Interval parameter from the NE.

The options are:

- Enabled (default)
- Disabled

When the check box is disabled, the Operational State Transition Interval parameter is configurable.

A.20 Bundles parameters

A.20.1 Show Only Compatible Channels

The Show Only Compatible Channels parameter specifies whether to filter the list of channels presented, by determining whether the channels can be configured as members of a multilink bundle.

The options are:

- Enabled (default)
- Disabled

A.21 Card Slot parameters

A.21.1 Administrative

The Administrative parameter is used to differentiate between an OS 6850 and OS 6850E NE for card slot configuration purposes at the administrative level.

The options are:

- OS 6850 (default)
- OS 6850E

A.21.2 Operational

The Operational parameter is used to differentiate between an OS 6850 and OS 6850E NE for card slot configuration purposes at the operational level.

The options are:

- OS 6850 (default)
- OS 6850E

A.22 Channel parameters

A.22.1 Ds3 Channel Payload Type

The Ds3 Channel Payload Type parameter specifies the type of payload that is configured for the channel. The parameter appears only for OC3 or OC12 ASAP ports when the Channel Type parameter is set to SONET Sts1 (Sdh Au3) and the Sts1 Channel Payload Type parameter is set to PDH Ds3.

The options are:

- None (default)
- DS1
- E1

A.22.2 Edit ATM button

Click on the Edit ATM button to configure the ATM Interface parameters for the channel.

A.22.3 Edit ILMI Link button

Click on the Edit ILMI Link button to configure the ILMI link parameters for the ATM interface.

A.22.4 Edit PPP button

Click on the Edit PPP button to configure the keep-alive Period and Drop Count parameters for the channel.

A.22.5 Keep Alive (seconds)

The Keep Alive (seconds) parameter specifies the interval, in seconds, used to send periodic keepalive packets. The receiver process expects to receive a keepalive packet at every keepalive interval. The link is declared down if the receiver process does not receive a keepalive packet within the timeout interval. The link is declared up when the number of continual keepalive packets received equals the up-count. The nodes at the two endpoints of the cHDLC link should be provisioned with the same values. The range is 0 to 300. The default is 10.

A.22.6 Restore Keep-Alive Defaults

Click on the Restore Keep-Alive Defaults button to restore the default values of the keep-alive parameters for the ILMI link.

A.22.7 STs1 Channel Payload Type

The STs1 Channel Payload Type parameter specifies the type of payload that is configured for the channel.

The options are:

- SONET VT15 (SDH Tu11)
- PDH Ds3 (default)

A.22.8 Time Slots per DS0 Channel Group

The Time Slots per DS0 Channel Group parameter specifies the timeslots from the selected TDM port to be assigned to the channel group. The range is 1 to 24 for DS1 channel groups and 1 to 31 for E1 channel groups. If you set the parameter option to “0”, a DS0 channel group is not created.

For DS1 channel types, the maximum number of timeslots is 24 (that is, TS1-TS24). For E1 channel types, the maximum number of timeslots is 31 (that is, TS2-TS32). Depending on the TDM port selected, the NFM-P automatically creates the DS0 channel groups with the appropriate type of timeslots.

Assigned DS0 timeslots are greyed out with a check mark against the assigned timeslot.

A.22.9 Vt15 Channel Payload Type

The Vt15 Channel Payload Type parameter specifies the type of payload that is configured for the channel. The parameter appears only for OC3 or OC12 ASAP ports when the Channel Type parameter is set to SONET Sts1 (Sdh Au3) and the Sts1 Channel Payload Type parameter is set to SONET VT15 (SDH Tu11). The only option is PDH Ds1.

A.23 Common equipment navigation tree parameters

A.23.1 Telnet Session button

Click on the Telnet Session button to open a Telnet session with the route. You can use the Telnet session to communicate directly with the managed objects.

A.24 Daughter Card and Daughter Card Slot parameters

A.24.1 Assigned MCM Card Type

The Assigned MCM Card Type parameter specifies the MDA carrier module type to be configured for the daughter card slot. The options depend on the settings of the unconfigurable Supported Daughter Card Types and Allowed Daughter Card Types parameters. You can configure the parameter when the In MDA Carrier Module Slot parameter is enabled.

A.25 Device parameters

A.25.1 Custom Property 1

The Custom Property 1 parameter specifies any custom user text you want to attach to an NE. You can add up to 80 characters of custom text. There is no default.

Note: This parameter name reflects the default name of the parameter on the NFM-P GUI. The name of this parameter can be changed using the System Preferences form.

A.25.2 Custom Property 2

The Custom Property 2 parameter specifies any custom user text you want to attach to an NE. You can add up to 80 characters of custom text. There is no default.

Note: This parameter name reflects the default name of the parameter on the NFM-P GUI. The name of this parameter can be changed using the System Preferences form.

A.25.3 Custom Property 3

The Custom Property 3 parameter specifies any custom user text you want to attach to an NE. You can add up to 80 characters of custom text. There is no default.

Note: This parameter name reflects the default name of the parameter on the NFM-P GUI. The name of this parameter can be changed using the System Preferences form.

A.25.4 Host Address

The Host Address parameter specifies the Cflowd collector host address. Specify an IPv4 address in dotted-decimal format

A.25.5 Ignore Timestamps

The Ignore Timestamps parameter specifies whether entries with unchanged last change timestamps are processed.

The options are:

- Enabled
- Disabled (default)

A.25.6 Port Number

The Port Number parameter specifies the Cflowd collector UDP port. The range is 1 to 65 535. The default is 2055.

A.25.7 Resource Group ID

CAUTION

Service Disruption

Changing the Resource Group ID parameter setting has serious consequences.

Do not configure the parameter without first contacting your Nokia technical-support representative.

The Resource Group ID parameter specifies the group to which the NE is assigned for the management of NFM-P system resources.

A.25.8 Template Re-transmit (seconds)

The Template Re-transmit (seconds) parameter specifies how often, in s, Cflowd sends Cflowd template definitions to collectors. The parameter is configurable when the Version parameter is set to version-9 or version-10. The range is 10 to 600. The default is 600.

A.25.9 Vendor-Specific ICMP Extensions

The Vendor-Specific ICMP Extensions parameter specifies whether vendor-specific ICMP functionality is enabled on the 7450 ESS, 7705 SAR, 7750 SR, and 7950 XRS device types.

The options are:

- Enabled
- Disabled (default)

A.26 Gateway parameters

A.26.1 Dual Stack Preference Cplane

The Dual Stack Preference Cplane parameter specifies the IP version that is used for the control plane for the S5/S8 path.

The options are:

- IPv4
- IPv6 (default)

A.26.2 Dual Stack Preference Uplane

The Dual Stack Preference Uplane parameter specifies the IP version that is used for the user plane for the S5/S8 path.

The options are:

- IPv4
- IPv6
- Control Plane

A.27 ISA-AA Group parameters

A.27.1 Collector Port

The Collector Port parameter specifies the UDP port of the AA Cflowd collector. The range is 1 to 65 535. The default is 4739.

A.27.2 Host Address

The Host Address parameter specifies the Cflowd collector host address. Specify an IPv4 address in dotted-decimal format

A.27.3 Template Re-transmit

The Template Re-transmit parameter specifies how often, in s, AA Cflowd sends Cflowd template definitions to collectors. The range is 10 to 600. The default is 600.

A.28 LAG parameters

A.28.1 Auto-Generate

The Auto Generate parameter specifies whether to automatically configure an Actor Administration Key parameter.

The options are:

- Enabled (default)
- Disabled

A.28.2 Hash

The Hash for Ethernet LAG enabled parameter specifies service-level hashing for Ethernet services on the Wavence SM.

The options are:

- L1
- L2
- L3

A.28.3 Show Only Compatible Ports

The Show Only Compatible Port Numbers parameter specifies whether to filter the list of ports presented, by determining whether the ports can be configured as LAGs.

The options are:

- Enabled (default)
- Disabled

When the parameter is enabled, several checks are performed.

Which of the following checks are performed depends on the type of NE:

- the port mode is network
- the port cannot be a mobile port
- the port cannot be POS
- the port cannot belong to another LAG
- there cannot be more than eight ports in a LAG
- the port cannot be bound to a Layer 3 interface
- the ports must have the same speed
- the port must be set to full duplex and auto negotiation must be turned off

A.28.4 View the newly created Port Termination

The View the newly created Port Termination parameter specifies that the Properties form for the newly created LAG opens when you close the configuration form.

A.28.5 Wait to Restore (ms)

The Wait to Restore (ms) parameter specifies the amount of time to wait after a fault clears before a traffic interface is restarted in the LAG. This parameter is specific to Wavence NEs. The range is 100 to 8000. The default is 1000.

A.29 MME parameters

A.29.1 Abort MME Load Balance button

The Abort MME Load Balance button cancels the load balancing operation.

A.29.2 Lock MME button

The Lock MME button prevents the MME from accepting new UE connections.

A.29.3 Unlock MME button

The Unlock MME button returns an MME instance to an unlocked state, allowing the MME to accept new UE connections.

A.30 Port parameters

A.30.1 Configured Alarms

The Configured Alarms parameter specifies the alarms that are monitored by the interface.

Table A-8 Configured Alarms (cfgAlarms): OTU Alarms

Configurable OTU Alarms	
OTU Backward Defect Indication	OTU Bit Error Rate Signal Degrade
OTU Bit Error Rate Signal Fail	OTU Alarm Indication Signal
Loss of Multi-frame	Loss of OTU Framing
Loss of Signal	Loss of Clock
ODU Alarm Indication Signal	Uncorrectable FEC errors
FEC Rx/Tx Mode Mismatch	FEC Signal Degrade
FEC Signal Failure	OTU Backward Incoming Alignment Error
OTU Incoming Alignment Error	OTU Trace ID Mismatch
PM Trace ID Mismatch	OPU PSI Payload Type Mismatch
OPU PSI Trace Mismatch	PM Backward Defect Indication
ODU Locked	ODU Open Connection Indication

Table A-9 Configured Alarms (cfgAlarms): Wave Tracker Alarms

Configurable Wave Tracker Alarms	
Power Control Low limit reached	Power Control High limit reached
Power Control Degrade	Power Control Failure
Encoder Degrade	Encoder Failure

Table A-10 Configured Alarms (ampcfgAlarms): Optical Amplifier Alarms

Configurable Optical Amplifier Alarms	
Amplifier Module communication failure	Amplifier Loss of output power
Amplifier Loss of input optical power	Amplifier Module Case temperature low
Amplifier Module Case temperature high	Amplifier Pump temperature
Amplifier Pump over-current	—

Table A-11 Configured Alarms (tdcmcfgAlarms): Optical Tunable Dispersion Compensation Module Alarms

Configurable Optical Tunable Dispersion Compensation Module Alarms	
Tdcm module communication failure	Tdcm EEPROM invalid
Tdcm thermal control temperature limit	Tdcm thermal control unlocked
Tdcm module temperature low	Tdcm module temperature high
Tdcm not ready	—

A.30.2 Expected Rx Bytes

The Expected Rx Bytes parameters specify the expected type of receiver (Rx) Trail Trace Identifier (TTI) in the OTU. The parameters appear only when the A.30.3 “Expected Rx Mode” (p. 3289) parameter is set to Bytes.

The value is a hexadecimal number.

Table A-12 Expected Rx Bytes parameters

Parameter	Specifies the TTI for
Expected Rx Bytes (pmTtiExp)	PM
Expected Rx Bytes (smTtiExp)	SM
Expected Rx Bytes (psiTtiExp)	PSI

A.30.3 Expected Rx Mode

The Expected Rx Mode parameters specify the expected type of TTI in the OTU overhead.

Table A-13 Expected Rx Mode parameters

Parameter	Specifies the TTI for	Options
Expected Rx Mode (pmTtiExpMode)	PM	Auto (default) String Bytes
Expected Rx Mode (smTtiExpMode)	SM	
Expected Rx Mode (psiTtiExpMode)	PSI	

A.30.4 Expected Rx String

The Expected Rx String parameters specify the expected receiver (Rx) TTI in the OTU overhead. The parameters are only configurable when the A.30.3 “Expected Rx Mode” (p. 3289) parameter is set to String.

The value is 0 to 192 bytes.

Table A-14 Expected Rx String parameters

Parameter	Specifies the TTI for
Expected Rx String (pmTtiExp)	PM
Expected Rx String (smTtiExp)	SM
Expected Rx String (psiTtiExp)	PSI

The auto-populated default sequence is the name of the NE, the IOM number, the MDA slot number, the port number, and the DWDM channel number.

A.30.5 J0 Byte

The J0 Byte parameter specifies a numeric value for a SONET section trace. This value is inserted at the source and is checked against the value expected by the receiver. The parameter is configurable when the Framing parameter is set to SONET and the SONET Section Trace Mode parameter is set to Byte. The range is a hexadecimal number from 00 to FF. The default is 01.

A.30.6 Receiver

These parameters specify the TTI in the received OPU overhead. The auto-populated default sequence is the name of the NE, the IOM number, the MDA slot number, the port number, and the DWDM channel number.

Table A-15 Receiver parameters

Parameter	Specifies the TTI for	Options
Receiver (pmTtiRx)	PM	0 to 192 bytes
Receiver (smTtiRx)	SM	0 to 192 bytes
Receiver (psiPayloadTypeRx)	PSI	0 to 254 bytes

A.30.7 Transmitter Bytes

The Transmitter Bytes parameters specify the type of transmit (Tx) TTI in the OTU. These parameter appears only when the A.30.8 “Transmitter Mode” (p. 3290) parameter is set to Bytes.

The value is a hexadecimal number.

Table A-16 Transmitter Bytes parameters

Parameter	Specifies the TTI for
Transmitter String (pmTtiTx)	PM
Transmitter String (smTtiTx)	SM
Transmitter String (psiTtiTx)	PSI

A.30.8 Transmitter Mode

The Transmitter Mode parameters specify the type of TTI in the OTU overhead.

Table A-17 Transmitter Mode parameters

Parameter	Specifies the TTI for	Options
Transmitter Mode (pmTtiTxMode)	PM	Auto (default) String Bytes
Transmitter Mode (smTtiTxMode)	SM	
Transmitter Mode (psiTtiTxMode)	PSI	

A.30.9 Transmitter String

These parameters allow specify the type of transmit TTI in the OTU overhead. These parameters are only configurable when the A.30.8 “Transmitter Mode” (p. 3290) parameter is set to String.

The value is 0 to 192 bytes.

Table A-18 Transmitter String parameters

Parameter	Specifies the TTI for
Transmitter String (pmTtiTx)	PM
Transmitter String (smTtiTx)	SM
Transmitter String (psiTtiTx)	PSI

The auto-populated default sequence is the name of the NE, the IOM number, the MDA slot number, the port number, and the DWDM channel number.

A.31 Shelf parameters

A.31.1 Interface Name

Click on the Select button beside the Interface Name parameter and select a timing reference.

A.31.2 Port or Channel Name

Click on the Select button beside the Port or Channel Name parameter and select a timing reference. When configuring Timing Reference One on a 7210 SAS-D ETR, only ports 1 to 4 can be selected. When configuring Timing Reference Two on a 7210 SAS-D ETR, only ports 5 and 6 can be selected. Ports 7 to 10 can be selected for any timing reference.

A.31.3 Quality Level Override

This parameter specifies the quality level used for SETS input selection. This value overrides any value received by the SSM process of that reference.

Table A-19 Quality Level Override parameters

Parameter	Description	Options
Quality Level Override (firstTimingReferenceQualityLevel)	Specifies the quality level used for the first timing reference for SETs input selection.	None Prs Stu St2 Tnc St3e St3 Prc Ssua Ssub Sec Eec1 Eec2
Quality Level Override (secondTimingReferenceQualityLevel)	Specifies the quality level used for the second timing reference for SETs input selection.	
Quality Level Override (bitsQualityLevel)	Specifies the quality level used for the reference for SETs input selection.	
Quality Level Override (ptpQualityLevel)	Specifies the quality level used for the reference for SETs input selection.	

A.31.4 Radio Name

The Radio Name parameter specifies specifies a description for the MW Link Member. The range is 0 to 32 characters.

A.32 TWAMP parameters

A.32.1 Conn Idle Time Periodic Threshold (seconds)

The Conn Idle Time Periodic Threshold (Seconds) parameter specifies the configurable threshold for the elapsed time, in seconds, for a TWAMP message to be received on this control connection. When the value of this parameter exceeds the value configured for Inactivity Timeout (Seconds), the connection will be closed.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.2 Conn Session Count Periodic Threshold

The Conn Session Count Periodic Threshold parameter specifies the configurable threshold for the number of test sessions conducted by the TWAMP server, for the connection specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.3 Conn Test Packets Rx Periodic Threshold

The Conn Test Packets Rx Periodic Threshold parameter specifies the configurable threshold for the number of TWAMP test packets received by the TWAMP server, for the connection specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.4 Conn Test Packets Tx Periodic Threshold

The Conn Test Packets Tx Periodic Threshold parameter specifies the configurable threshold for the number of TWAMP test packets sent by the TWAMP server, for the connection specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.5 Conn Test Sess Completed Periodic Threshold

The Conn Test Sess Completed Periodic Threshold parameter specifies the configurable threshold for the number of test sessions completed by the TWAMP server, for the connection specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.6 Conn Test Sess Rejected Periodic Threshold

The Conn Test Sess Rejected Periodic Threshold parameter specifies the configurable threshold for the number of test sessions rejected by the TWAMP server, for the connection specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.7 Srv Pfx Conn Count Periodic Threshold

The Srv Pfx Conn Count Periodic Threshold parameter specifies the configurable threshold for the number of control connections currently managed by the TWAMP server, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.8 Srv Pfx Conns Rejected Periodic Threshold

The Srv Pfx Conns Rejected Periodic Threshold parameter specifies the configurable threshold for the number of control connection requests which have been rejected by the TWAMP server, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.9 Srv Pfx Session Count Periodic Threshold

The Srv Pfx Session Count Periodic Threshold parameter specifies the configurable threshold for the number of currently in-progress TWAMP test sesssions, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.10 Srv Pfx Test Packets Rx Periodic Threshold

The Srv Pfx Test Packets Rx Periodic Threshold parameter specifies the configurable threshold for the number of TWAMP test packets received by the TWAMP server.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.11 Srv Pfx Test Packets Tx Periodic Threshold

The Srv Pfx Test Packets Tx Periodic Threshold parameter specifies the configurable threshold for the number of TWAMP test packets sent by the TWAMP server.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.12 Srv Pfx Test Sess Abort Periodic Threshold

The Srv Pfx Test Sess Abort Periodic Threshold parameter specifies the configurable threshold for the number of test sessions aborted by the TWAMP server, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.13 Srv Pfx Test Sess Completed Periodic Threshold

The Srv Pfx Test Sess Completed Periodic Threshold parameter specifies the configurable threshold for the number of test sessions completed by the TWAMP server, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.32.14 Srv Pfx Test Sess Rejected Periodic Threshold

The Srv Pfx Test Sess Rejected Periodic Threshold parameter specifies the configurable threshold for the number of test sessions rejected by the TWAMP server, for the prefix specified by the index values.

The range is 0 to 9 223 372 036 854 775 807. The default is 0.

A.33 Common Manage menu parameters

A.33.1 Common Manage menu parameters overview

This section describes the parameters that are common to the NFM-P Manage menus forms.

Note: This section also describes parameters common to the Service Template form and child forms. Service templates are intended for use in the NFM-P GUI. Nokia does not recommend using the XML API XML classes and methods associated with service templates.

A.33.2 Allow Binding Of Templates Not Associated With Any Subscriber

The Allow Binding Of Templates Not Associated With Any Subscriber parameter specifies whether you can bind templates even when they do not have a subscriber association.

The options are:

- enabled (default)
- disabled

A.33.3 SAP Administrative State

The SAP Administrative State parameter specifies whether the SAP is administratively enabled.

The options are:

- Up (default)
- Down

A.34 Customers parameters

A.34.1 Apdex Thresholds - Unacceptable

The Apdex scores below this threshold are unacceptable quality, the parameter range is 0.00 to 1.00 in 0.01 increments. The default value is 0.5.

A.34.2 Apdex Thresholds - Poor

The Apdex scores below this threshold are poor quality, the parameter range is 0.00 to 1.00 in 0.01 increments. The default value is 0.7.

A.34.3 Apdex Thresholds - Fair

The Apdex scores below this threshold are fair quality, the parameter range is 0.00 to 1.00 in 0.01 increments. The default value is 0.85.

A.34.4 Apdex Thresholds - Good

The Apdex scores below this threshold are good quality, the parameter range is 0.00 to 1.00 in 0.01 increments. The default value is 0.94.

A.34.5 MOS Thresholds - Unacceptable

The MOS scores below this threshold are unacceptable quality, the parameter range is 0.00 to 5.00 in 0.01 increments. The default value is 1.0.

A.34.6 MOS Thresholds - Poor

The MOS scores below this threshold are poor quality, the parameter range is 0.00 to 5.00 in 0.01 increments. The default value is 2.0.

A.34.7 MOS Thresholds - Fair

The MOS scores below this threshold are fair quality, the parameter range is 0.00 to 5.00 in 0.01 increments. The default value is 3.0.

A.34.8 MOS Thresholds - Good

The MOS scores below this threshold are good quality, the parameter range is 0.00 to 5.00 in 0.01 increments. The default value is 4.0.

A.35 IPsec VPN parameters

A.35.1 Security Policy ID

The Security Policy ID parameter specifies an ID for the IPsec security policy on the VPRN site. The range is 1 to 8192. The default is 0.

A.35.2 Tunnel Type

The Tunnel Type parameter specifies the tunnel type for the tunnel group.

The options are:

- Dynamic (Site-to-Site)
- Dynamic (Soft Client)
- Static

A.36 LSPs parameters

A.36.1 Auto Select Hop-less Path

The Auto Select Hop-less Path parameter specifies whether MPLS paths are explicitly specified during LSP creation or the LSP uses a completely loose path and the NFM-P selects the MPLS paths for the LSP.

Table A-20 Auto Select Hop-less Path parameter

Option	Option description
Disabled (default)	Specifies that you must manually choose or create the MPLS paths during LSP creation based on the source NE. IGP is used to identify the next hop.
Enabled	Specifies a loose path. The NFM-P chooses the MPLS paths for the LSP. If no hopless path to the destination exists, the NFM-P creates one.

Note: If the A.36.1 “Auto Select Hop-less Path” (p. 3296) parameter is enabled, the NFM-P automatically creates a hopless MPLS path for the dynamic LSP. This path will be named “HOP-LESS PATH <system_IP_adress>-<ID>”, where the system_IP_address is the source NE address and the ID is the unique identifier for the MPLS. If you want to create an MPLS path and use the A.36.1 “Auto Select Hop-less Path” (p. 3296) parameter to assign it to the LSP, it must follow this naming convention.

A.36.2 Exclude Node in ERO

The Exclude Node in ERO parameter specifies whether or not a device should be excluded from the manual bypass LSP path.

When this parameter is enabled, you can enter an IPv4 address for the Exclude Node IP Address parameter.

The options are:

- enabled
- disabled (default)

A.36.3 View the newly created Bypass Only Lsp

The View the newly created Bypass Only Lsp parameter specifies whether you want to view the configuration information about the newly created LSP. The Bypass Only LSP child form displays the service tunnel configuration information.

The options are:

- Enabled
- Disabled (default)

A.36.4 View the newly created Dynamic LSP

The View the newly created Dynamic LSP parameter specifies whether you want to view the configuration information about the newly created LSP. The Dynamic LSP child form displays the service tunnel configuration information.

The options are:

- Enabled
- Disabled (default)

A.37 Mirror Services parameters

A.37.1 Mirror Services parameters overview

This section describes the unique parameters on the Manage Mirror Service Templates form and child forms.

Note: The mirror service templates are intended for use in the NFM-P GUI. Nokia recommends that you do not use the XML API XML classes and methods that are associated with the service templates.

A.37.2 Template Description

The Template Description parameter specifies a description for the template type. The range is an interface name of 0 to 80 characters. The default is an empty string.

A.38 MPLS Paths parameters

A.38.1 Insert Hop

Click the Insert Hop button to insert a hop in the MPLS provisioned path. A hop can be an interface on an Nokia-managed NE or a third-party device.

A.38.2 Specify Site

The Specify Site parameter specifies the method for configuring the destination site of the MPLS path.

Table A-21 Specify Site parameter

Option	Option description	Dependencies
Manually	Specifies that you must select the destination site for the MPLS path by manually entering the IP address.	—
By Selection	Specifies that you must select the destination site for the MPLS path by: • listing and choosing the network element site ID or IP address using the Select buttons • manually entering the IP address	

A.38.3 Starting Network Element

Click the Select button to specify the starting network element for the MPLS path.

A.39 Services parameters

A.39.1 Number of SAPs

The Number of SAPs parameter displays the total number of SAPs present under a service site.

A.39.2 What type of interface would you like to create?

The What type of interface would you like to create? parameter specifies the:

- type of multicast interface that is added to an existing IES
- type of protocol that is applied to an existing IES or IES SAP

Table A-22 What type of interface would you like to create? parameter

Object	Options
IES	PIM (default) IGMP

Table A-22 What type of interface would you like to create? parameter (continued)

Object	Options
IES or IES SAP	OSPFv2 (default) OSPFv3 RIP ISIS

A.40 Service Tunnel parameters

A.40.1 No VLAN VC Ethertype

The No VLAN VC Ethertype parameter specifies whether a VLAN VC Ethertype parameter can be specified. When the No VLAN VC Ethertype parameter is enabled the VLAN VC Ethertype parameter cannot be set.

The options are:

- enabled (default)
- disabled

A.40.2 Order

The Order parameter specifies whether the group is ordered or unordered.

The options are:

- unordered (default)
- ordered

A.40.3 View the newly created tunnel

The View the newly created Dynamic LSP parameter specifies whether you want to view the configuration information about the newly created LSP. The Dynamic LSP child form displays the service tunnel configuration information.

The options are:

- Enabled
- Disabled (default)

A.40.4 VPLS

The VPLS parameter specifies whether Ethernet ring element interconnection is enabled through VPLS. The VPLS parameter is configurable when the Type parameter is set to Non Virtual Link. When the VPLS parameter is set to True, the Ethernet Ring ID parameter is set to 4294967295. The default is False.

A.41 Templates parameters

A.41.1 Command Type

The Command Type parameter specifies the command type for the template.

The options are:

- Create (default)
- Modify

When you choose the Create option, the template is used to create new objects. When you choose the Modify option, the template is used to modify an existing object.

A.41.2 Generate First (Base) Version

The Generate First (Base) Version parameter specifies whether a Velocity UI header is generated in the XML API configuration template script at template creation.

The options are:

- Enabled (default)
- Disabled

A.41.3 Generate Velocity Properties

The Generate Velocity Properties parameter specifies whether a Velocity UI header is generated in the XML API configuration template script.

The options are:

- Enabled (default)
- Disabled

A.41.4 Mode

The Mode parameter is used to specify the state of the template.

The options are:

- Draft (default)
- Released

A.41.5 Show Created Object

The Show Created Object parameter specifies whether to open the properties form for the object that you are creating from the template.

The options are:

- Enabled
- Disabled (default)

A.41.6 Type

The Type parameter specifies a user-defined description of the XML API configuration templated category. The range is 0 to 255 characters.

A.42 VLAN group and path parameters

A.42.1 Destination Network Element

The Destination Network Element parameter specifies the IP address of the Wavence SM at the end of the VLAN path. There is no default.

A.42.2 Has Tail

The Has Tail parameter specifies whether a tail node is attached to a ring element.

The options are:

- Disabled (default)
- Enabled

A.42.3 Head End

The Head End parameter specifies the IP address of the Wavence SM at the head end of the VLAN group.

A.42.4 Is Point to Multipoint

The Point to Multipoint parameter specifies whether the VLAN path is a P2P (dot1q) VLAN path or a P2MP (dot1q) VLAN path.

The options are:

- Disabled -- P2P (default)
- Enabled -- P2MP

Note: The P2P (dot1q) VLAN services is for users that need to restrict a dot1q VLAN service to two sites.

The P2MP (dot1q) VLAN services is for users who need a dot1q VLAN service with two or more sites.

Services must be deleted to change from P2P to P2MP and P2MP to P2P.

A.42.5 Link

The Link parameter specifies the link for the VLAN path to traverse.

A.42.6 Minimum Bandwidth (kbps)

The Minimum Bandwidth (kbps) parameter specifies the guaranteed bandwidth that is available across all network ports that have to be selected during creation of a VLAN path. This parameter is optional.

A.42.7 Network Element

The Network Element parameter specifies which hops the VLAN paths traverses.

A.42.8 Starting Network Element

The Starting Network Element parameter specifies the IP address of the Wavence SM at the start of the VLAN path.

A.43 7705 SAR Fabric parameters

A.43.1 Rate To MDA (destRateTo1<card#>

The Rate To MDA parameter specifies the rate, in kb/s, to each daughter card slot. You can configure the rates individually for each daughter card slot on the 7705 SAR. There is a rate for each slot identified by IOMslot/daughter card slot; for example, Rate to MDA 1/1, Rate to MDA 1/2, and so on. You can configure each of the rates in the range 1 to 1 000 000 kb/s and 200 000 kb/s is the default value.

A.44 Wavence NE QoS parameters

A.44.1 DSCP

The DSCP (DiffServ) parameter specifies a QoS classification that uses the DS field of the IPv4 or IPv6 packet header defined in RFC2474. The priority values in the DSCP (Differentiated Service Code Point) field are not directly mapped to queues but are mapped to the internal forwarding classes.

The Priority values are mapped to eight forwarding classes from 0 (lowest priority) to 7 (highest priority). When an incoming packet is not IPv4 or IPv6, the packet is assigned to the lowest priority queue.

A.45 Application Assurance parameters

A.45.1 Add Factor

The Add Factor parameter specifies that the Factor Value parameter value is to be added to each application filter entry ID to be renumbered.

A.45.2 Multiply Factor

The Multiply Factor parameter specifies that each application filter entry ID to be renumbered is to be multiplied by the Factor Value parameter value.

A.45.3 Port Value Type

The Port Value Type parameter specifies the destination port value type.

The options are:

- Single (default)

- Range

A.45.4 Port Value Type

The Port Value Type parameter specifies the source port value type.

The options are:

- Single (default)
- Range

A.46 Access Egress parameters

A.46.1 Port Redirect Queue Group

Table A-23 Port Redirect Queue Group parameter

Parameter	OSS equivalent name	Description
Port Redirect Queue Group	portRedirectGroupQueue	Specifies that packets are redirected to a queue in a SAP egress port queue group. The forwarding class queue ID is explicitly specified in the SAP egress QoS policy.
Port Redirect Queue Group (for HSMDA queues)	portRedirectHsmdaQueuePortGrpQ	Specifies that packets are redirected to an HSMDA queue in a SAP egress port queue group. The forwarding class HSMDA queue ID is explicitly specified in the SAP egress QoS policy.

The choices for either version of this parameter are:

- disabled (default)
- enabled

A.46.2 Traffic Control

The Traffic Control parameter specifies the type of object used in the forwarding class.

The options are:

- Use Policer
- Use Queue Group
- Use Queue (default)

The Use Policer option specifies whether a policer is used in the forwarding class. You can assign a policer to a forwarding class for Unicast traffic. When you assign a policer, you must select a Policer ID in the Policers panel.

The Use Queue Group option specifies whether a queue from an egress queue group template is used in the forwarding class. You must select an egress queue from the specified queue group template policy in the Queue panel.

The Use Queue option specifies whether a queue configured locally within the policy is used in the forwarding class. You must select a Queue ID in the Queue panel.

A.47 Access Ingress parameters

A.47.1 Use Policer

The Use Policer parameter specifies whether a policer is used in the forwarding class.

You can assign a policer to a forwarding class for the following traffic types:

- Unicast
- Multicast
- Broadcast
- Unknown

When you assign a policer for one or more of the traffic types, you must select a policer ID for each traffic type in the Policers panel.

A.47.2 Use Queue Group

The Use Queue Group parameter specifies whether a queue from an ingress queue group template is used in the forwarding class, and if so, the type of queue.

The options are:

- Unicast Queue
- Multicast Queue
- Broadcast Queue
- Unknown Queue

When you enable one or more of the options, you must select a ingress queue from the specified queue group template policy in the Queues panel. If you do not enable the parameter options, the ingress queues that are listed for each queue type are queues that are configured locally within the policy.

A.48 Auto Tunnels parameters

A.48.1 Order

The Order parameter specifies whether the group is ordered or unordered.

The options are:

- unordered (default)
- ordered

A.48.2 View the newly created tunnel

The View the newly created Dynamic LSP parameter specifies whether you want to view the configuration information for the newly created LSP. The Dynamic LSP child form displays the service tunnel configuration details.

The options are:

- Enabled
- Disabled (default)

A.49 Common Policies menu parameters

A.49.1 Default

The Default parameter specifies whether a parameter is set to the default value or can be configured.

The options are:

- disabled
- enabled (default)

When the Default parameter is set to enabled, you cannot configure the associated parameters.

A.49.2 Destination IP Prefix

The Destination IP Prefix parameter specifies a list of IP prefixes to use as destinations for the filter policy. When the parameter is enabled, click on the Select button and choose a prefix list.

A.49.3 Log ID

The Log ID parameter specifies the filter log for the filter entry. Click on the Select button beside the parameter to choose a filter, or type in a value. The range is 101 to 199. The default is 0, which means that the parameter is not configured.

A.49.4 New Entry ID

The New Entry ID parameter specifies the new identifier of the filter entry. The parameter value becomes the Entry ID of the filter entry. A filter policy compares the contents of a packet to each filter entry until it finds a match, beginning with the lowest Entry ID. The range is 1 to 65535. There is no default.

A.49.5 Override CIR

The Override CIR parameter specifies the override value for the administrative committed information rate for the queue or the policer. The default is MAX which specifies the maximum available CIR.

The CIR override is configured in kbps if the Rate Type for the queue or policer was originally set to kbps. The CIR override is configured as a Percent (%) if the Rate Type for the queue or policer was originally set to either Percent Port Limit (queues only) or Percent Local Limit.

The range depends on the type of object being configured.

Table A-24 Override CIR parameter

Object type	Range
Queue	The range is 0 to 100 000 000 kbps (0 to 100%), depending on the line rate of the object to which the policy is applied. A value of -2 specifies that no override is applied. A value of -1 specifies that the maximum override is applied.
Policer	The range is 0 to 20 000 000 kbps (0 to 100%), depending on the line rate of the object to which the policer is applied. A value of -2 specifies that no override is applied. A value of -1 specifies that the maximum override is applied.

A.49.6 Override CIR Adaptation

The Override CIR Adaptation parameter specifies the override value for the method used by the device to derive the operational CIR setting when the queue is provisioned in hardware

You can also choose the Default option, which specifies that the operational CIR for the queue is the rate closest to the rate specified by the CIR (kbps) parameter.

Table A-25 Override CIR Adaptation parameter

Option	Option description	Dependencies
Closest (default)	The operational CIR for the queue is the rate closest to the rate specified by the CIR (kbps) parameter.	—
Min	The operational CIR for the queue is equal to or greater than the administrative rate specified by the CIR (kbps) parameter.	
Max	The operational CIR for the queue is equal to or less than the administrative rate specified by the CIR (kbps) parameter.	

A.49.7 Override Committed Burst Size

The Override Committed Burst Size parameter specifies the override value for the Committed burst pool size for a queue.

The range is -1 to 4 194 304, with -1 indicating default. Default specifies that the device calculates committed burst pool size based on the CIR and the PIR. The default is default.

A.49.8 Override High Priority Reserved

The Override High Priority Reserved parameter specifies the override value for the percentage of buffer pool space for the queue that is only used by high-priority packets.

The range is -1 to 100, with -1 indicating default. The default is default, which means that the device calculates the high priority reserved pool size based on the CIR and the PIR values.

A.49.9 Override Maximum Burst Size

The Override Maximum Burst Size parameter specifies the override value for the maximum burst pool size for a queue.

The range is -1 to 4 194 304, with -1 indicating default. Default specifies that the device calculates the committed burst pool size based on the CIR and PIR.

A.49.10 Override Packet Offset

The Override Packet Offset parameter specifies the number of offset bytes to add to (or remove from) a packet handled by a policer. A positive number adds bytes. A negative number removes bytes. The range is -128 and -32 to 31. The -128 value indicates that no override is applied. The default is -1.

A.49.11 Override PIR

The Override PIR parameter specifies the override value for the administrative peak information rate for the queue or the policer. The default is MAX which specifies the maximum available PIR.

For policers and non-HSMDA queues, the PIR override is configured in kbps if the Rate Type for the queue or policer was originally set to kbps. The PIR override is configured as a Percent (%) if the Rate Type for the non-HSMDA queue or policer was originally set to either Percent Port Limit (queues only) or Percent Local Limit.

The range depends on the type of object being configured.

Table A-26 Override PIR parameter

Object type	Range
HSMDA queue	The range is 0 to 100 000 000 kbps, depending on the line rate of the object to which the policy is applied. A value of -2 specifies that no override is applied. A value of -1 specifies that the maximum override is applied.
Queue	The range is 0 to 100 000 000 kbps (0 to 100%), depending on the line rate of the object to which the policy is applied. A value of -2 specifies that no override is applied. A value of -1 specifies that the maximum override is applied.
Policer	The range is 0 to 20 000 000 kbps (0 to 100%), depending on the line rate of the object to which the policer is applied. A value of -2 specifies that no override is applied. A value of -1 specifies that the maximum override is applied.

A.49.12 Override PIR Adaptation

The Override PIR Adaptation parameter specifies the override value for the method used by the device to derive the operational PIR setting when the queue is provisioned in hardware

You can also choose the Default option, which specifies that the operational PIR for the queue is the rate closest to the rate specified by the PIR (kbps) parameter.

Table A-27 Override PIR Adaptation parameter

Option	Option description	Dependencies
Closest (default)	The operational PIR for the queue is the rate closest to the rate specified by the PIR (kbps) parameter.	—
Min	The operational PIR for the queue is equal to or greater than the administrative rate specified by the PIR (kbps) parameter.	
Max	The operational PIR for the queue is equal to or less than the administrative rate specified by the PIR (kbps) parameter.	

A.49.13 Override Port Average Overhead

The Override Port Average Overhead parameter specifies the override value for the average percentage that the offered load to a queue is expected to expand during the frame encapsulation process before sending traffic on queues that egress a SONET or SDH port or channel.

The range is 0 to 100. Default specifies that the egress QoS policy for the queue is applied.

A.49.14 Override Queue CIR Weight

The Override Queue CIR Weight parameter specifies the override value for the weight that should be assigned to this queue by the parent scheduler among all the entities feeding into the parent when the traffic is conforming to the committed rate.

The range is -2 to 100. Default specifies that the egress or ingress QoS policy for the queue is applied. A value of '0' specifies that the queue will not receive bandwidth for the 'within-cir' pass on its parent scheduler.

A.49.15 Override Queue Weight

The Override Queue Weight parameter specifies the weight that needs to be used by the scheduler to which this queue would be feeding.

The range is -2 to 100. Default specifies that the egress or ingress QoS policy for the queue is applied.

A.49.16 Override Summed CIR

The Override Summed CIR parameter specifies the override value for the summed CIR.

The options are:

- true
- false (default)

A.49.17 Value

The following table lists the variations in the Value parameter for the Advanced Configuration Policy. These values will be assigned to whatever entity the Advanced Configuration Policy is applied to, unless the associated Default parameter is enabled.

Table A-28 Value parameter

Applicability	XML reference	Description	Range
Above Offered Cap (kbps)	aboveOfferedCapRate	Specifies the explicit number of kilobits per second that should be used as the limit to the applicable entity's fair share increase to the operational PIR. This is capped by the entity's configured PIR.	0 to 100 000 000, and also -1. Default = 0 -1 represents the maximum value.
Above Offered Cap (%)	aboveOfferedCapPirPercent	Specifies the percentage of the administrative PIR that should be used as the limit to the applicable entity's fair share increase to the operational PIR. This is capped by the entity's configured PIR.	0 to 100 Default = 100
Granularity (of Bandwidth Distribution) (kbps)	bwdGranularityRate	Specifies the explicit number of kilobits per second that should be used as the rounding step value for the applicable entity's administrative PIR.	0 to 100 000 000 Default = 0
Granularity (of Bandwidth Distribution) (%)	bwdGranularityPirPercent	Specifies the granularity percentage of the applicable entity’s administrative PIR that should be used as the rounding step value.	0 to 100 Default = 0
Granularity (of Measured Offered Rate) (kbps)	omGranularityRate	Specifies the explicit number of kilobits per second that should be used as the offered rate change sensitivity value for the applicable entity's administrative PIR.	0 to 100 000 000 Default = 0
Granularity (of Measured Offered Rate) (%)	omGranularityPirPercent	Specifies the granularity percentage of the applicable entity’s administrative PIR that should be used as the threshold sensitivity to the offered rate.	0 to 100 Default = 0
Measured Offered Rate Increase (kbps)	childAdminRate	Specifies the explicit number of kilobits per second of the applicable entity’s administrative PIR that should be added to the entity's offered rate. This is capped by the entity's configured PIR.	0 to 100 000 000 Default = 0
Measured Offered Rate Increase (%)	childAdminPirPercent	Specifies the percentage of the applicable entity’s administrative PIR that should be added to the entity's offered rate. This is capped by the entity's configured PIR.	0 to 100 Default = 0

Table A-28 Value parameter (continued)

Applicability	XML reference	Description	Range
Max Decrement (kbps)	maxDecrementRate	Specifies the explicit number of kilobits per second of the applicable entity's administrative PIR that should be used as the decrement limit to the offered rate change.	0 to 100 000 000 Default = 0
Max Decrement (%)	maxDecrementPirPercent	Specifies the percentage of the applicable entity's administrative PIR that should be used as the decrement limit to the offered rate change.	0 to 100 Default = 100

A.50 Format and Range Policies parameters

A.50.1 Auto Assign By Default

The Auto Assign By Default parameter specifies whether the Auto-Assign ID check box is enabled by default.

The options are:

- enabled
- disabled

A.50.2 Auto Assignment Enabled

The Auto Assignment Enabled parameter specifies whether the Auto-Assign ID parameter is enabled by default.

The options are:

- enabled
- disabled

A.50.3 Copy Text From Position

The Copy Text From Position parameter specifies the source attribute text that is copied. This parameter is configurable when you choose the Auto-Filled option for the text block format parameters. The range is 1 to 1000 characters.

A.50.4 Default Value

The Default Value parameter specifies the default of the text field. The parameter is configurable when the text block format option is Text Parameter. The range is 0 to 250 characters.

A.50.5 Mask

The Mask parameter specifies the text that is entered for a parameter value when the parameter value is configured by a format policy. For example, when a format policy is created for a service description, you can configure the text that automatically populates the description field. You cannot enter values that do not comply with the mask. The range is 0 to 250 characters.

The following table lists the special formatting characters available for the Mask parameter.

Table A-29 Special formatting characters

Character	Source object property options
‘	Used to escape any of the special formatting characters
#	Any valid number
U	Any letter, where lowercase characters are mapped to uppercase characters
L	Any letter, where uppercase characters are mapped to lowercase characters
A	Any letter or number
?	Any letter
*	Anything
H	Any hexidecimal character: • 0 to 9 • a to f • A to F

A.50.6 Maximum

The Maximum parameter specifies the highest number in the ID range for a range policy. The range is 0 to 1000000999999. The default is automatically set when the range policy is created. You can click on the Reset to Default button to reset the value.

A.50.7 Max. Length

The Max. Length parameter specifies the maximum length of a text field parameter that is configured in a format policy text block. The range is 1 to 250 characters.

A.50.8 Minimum

The Minimum parameter specifies the lowest number that can be used in the ID range for a range policy. The range is 0 to 1000000999999. The default is automatically set when the range policy is created. You can click on the Reset to Default button to reset the value.

A.50.9 Min. Length

The Min. Length parameter specifies the minimum length of a field text parameter that is configured in a format policy text block. The range is 0 to 250 characters.

A.50.10 Object Type

The Object Type parameter specifies the object class that is configured in the format or range policy. Choose an object type by clicking on the Select button.

A.50.11 Priority

The Priority parameter specifies the sequence in which multiple instances of a policy are sorted. The range is 1 to 1000 (lowest priority). The list of policies returned by the policy that matches a query is listed using the priority value.

A.50.12 Property Name

The Property Name parameter specifies the parameter to which the format or range policy applies.

The options are:

- Name
- Service Name
- Description
- Service ID
- ID
- Interface ID
- Outer Encapsulation Value

A.50.13 Read Only

The Read Only parameter specifies when the A.50.9 “Min. Length” (p. 3311) and A.50.7 “Max. Length” (p. 3311) parameters are read-only. The Read Only parameter is configured in the Text Format option of the Text Block Format form.

The options are:

- Enabled
- Disabled (default)

A.50.14 Source Object Name

The Source Object Name parameter specifies the object type to be used as the source to create the text string. This parameter is set during Auto-Filled Parameter configuration.

A.50.15 Source Property Name

The Source Property Name parameter specifies the object attributes to be used as the source to create the text string. This parameter is configured when the Auto-Filled Parameter option is configured.

A.50.16 Tooltip Text

The Tooltip Text parameter allows the operator to create a description for the parameter value to inform users that a format policy is applied to the parameter. The range is 0 to 1000 characters.

A.50.17 Through To Position

The Through To Position parameter copies the text of the source parameter up to the position that is configured by this parameter. This parameter is configurable when the auto-filled option is selected for the text block format and the A.50.18 “Unlimited” (p. 3312) parameter is disabled. The range is 1 to 1000.

A.50.18 Unlimited

The Unlimited parameter specifies the amount of source attribute text to be copied. The parameter is configurable when the auto-filled option is selected for the text block format. When the parameter is enabled, the source attribute text is copied from the A.50.3 “Copy Text From Position” (p. 3310) parameter value to the end. When the parameter is disabled you can copy the source parameter text to a specified position.

The options are:

- enabled
- disabled

A.51 HSMDA WRED Slope parameters

A.51.1 HSMDA High Slope

The High Slope parameter specifies how the high-priority packets, also known as in-profile packets, access the shared portion of the buffer pool. The probability of the buffer pools discarding packets is determined by the RED slope. The RED slope is indicated on the client GUI in a graph.

The Start Depth, Max Depth, and Max Prob. parameters specify the values of the RED slope.

In general, packets that are in-profile fall within the configured CIR and PIR ranges. This means that the packets are more likely to gain access to the shared buffer pool, and not be discarded because of a buffer overflow.

A.51.2 HSMDA Low Slope

The Low Slope parameter specifies how the low-priority packets, also known as out-of-profile packets, access the shared portion of the buffer pool. The probability of the buffer pools discarding packets is determined by the RED slope. The RED slop is indicated on the client GUI in a graph.

The Start Depth, Max Depth, and Max Prob. parameters specify the values of the RED slope.

In general, packets that are out-of-profile fall outside configured CIR and PIR ranges. This means that the packets should not be allowed to cause a buffer overflow in the shared buffer pool, and are more likely to be discarded because of a buffer overflow.

A.52 Ingress Multicast Path Management parameters

A.52.1 Administrative BW (kbps)

The Administrative BW parameter specifies the multicast channel's administrative bandwidth in kilo-bits per second. The range is from 0 to 40 000 000 kbps. The default is 0.

The XML property name for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-30 Administrative BW (kbps) parameter variations

Parameter configuration level	XML property name
Channel bundle	defaultAdminBw
Channel range	adminBw
Channel override	adminBw

A.52.2 Black Hole Rate (kbps)

The Black Hole Rate parameter specifies at which current rate a channel (including channels within a channel range or bundle) should be placed in the black-hole state (packets are dropped). This value can only be set when the A.52.3 “BW Decision” (p. 3314) parameter is set to Dynamic. The range is 0 to 40 000 000. The default is 0, which means never.

The XML property name for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-31 Black Hole Rate (kbps) parameter variations

Parameter configuration level	XML property name
Channel bundle	defaultBlackHoleRate
Channel range	blackHoleRate
Channel override	blackHoleRate

A.52.3 BW Decision

The BW Decision parameter specifies how the multicast ingress path manager determines the amount of bandwidth required by a multicast channel, including channels within a channel range or bundle.

The XML property name and default value for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-32 BW Decision parameter variations

Parameter configuration level	XML property name	Default value
Channel bundle	defaultBwDecision	Dynamic
Channel range	bwDecision	unspecified
Channel override	bwDecision	unspecified

A.52.4 Explicit Path

The Explicit Path parameter specifies an explicit ingress switch fabric multicast path for the channels (including those in a range or a bundle). If the parameter is set to None, the Multicast Path Manager dynamically places channels on the most optimal switch fabric paths.

The options are:

- None (default)
- Ancillary
- Primary
- Secondary

The XML property name for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-33 Explicit Path parameter variations

Parameter configuration level	XML property name
Channel bundle	defaultExplicitPath
Channel range	explicitPath
Channel override	explicitPath

A.52.5 Falling Delay (seconds)

The Falling Delay parameter specifies the value the bandwidth manager uses as a threshold to hold on to the previous highest bandwidth until the delay time has expired, while operating in dynamic bandwidth mode. This allows the bandwidth manager to ignore momentary drops in channel bandwidth. This value can only be set if A.52.3 “BW Decision” (p. 3314) is set to Dynamic. A value of 0 specifies that a non-zero value of the parent is applied.

The XML property name for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-34 Falling Delay (seconds) parameter variations

Parameter configuration level	XML property name	Default value	Range	Parent
Channel bundle	defaultFallingDelay	30	10 to 3600s	—
Channel range	fallingDelay	0	0 | 10 to 3600	Channel bundle
Channel override	fallingDelay	0	0 | 10 to 3600	Channel bundle Channel range

A.52.6 Preference Level

The Preference Level parameter specifies the relative preference level for multicast channels. The preference of a channel (including those in a bundle or range) specifies its relative importance over other multicast channels. Eight levels of preference are supported: 0 through 7. Preference value 7 indicates the highest preference level.

The XML property name and default value for this parameter varies depending on where it is set, namely at the bundle, channel range, or channel override level.

Table A-35 Preference Level parameter variations

Parameter configuration level	XML property name	Default value
Channel bundle	defaultPrefLevel	7
Channel range	prefLevel	0
Channel override	prefLevel	0

A.53 Named buffer pool parameters

A.53.1 Default Reserved CBS

The Default Reserved CBS parameter specifies whether the default percent for reserved CBS is used.

The options are:

- Enabled
- Disabled

A.54 NAT Policy parameters

A.54.1 Default

The Default parameter specifies whether the IPFIX Collector Source Address parameter is included in IPFIX messages. Enabling this parameter allows you to configure the A.54.3 “IPFIX Collector Source Address” (p. 3316) parameter.

The options are:

- Enabled
- Disabled (default)

A.54.2 IPFIX Collector Address

The IPFIX Collector Address parameter specifies the IPv4 address of the IPFIX collector to which IP flow information is sent. Specify an IPv4 address in dotted-decimal format.

A.54.3 IPFIX Collector Source Address

The IPFIX Collector Source Address parameter specifies the IPv4 address from which IP flow information is sent to an IPFIX collector. When this parameter is not set, the IPFIX collector identifies sources by the IP address from which the IPFIX message is sent; if the source uses a proxy, this is the IP address of the proxy. Use this parameter to specify a different IP address. This parameter cannot be configured unless the A.54.1 “Default” (p. 3316) parameter is enabled. Specify an IPv4 address in dotted-decimal format.

A.54.4 IPFIX Export Policy

The IPFIX Export Policy parameter specifies the IPFIX export policy for the NAT policy. Click on the Select button to choose an IPFIX export policy.

A.54.5 MTU

The MTU parameter specifies the MTU size, in bytes, for IPFIX messages sent to the IPFIX collector. The range is 512 to 9212. The default is 1500.

A.54.6 Refresh Timeout

The Refresh Timeout parameter specifies the refresh rate of the IPFIX source, in hours, minutes, and seconds. After this time, the IPFIX source resends the IPFIX template information to the IPFIX collector. The range is 0 seconds to 24 hours, 59 minutes, and 59 seconds. The default is 10 minutes.

A.54.7 Router Instance Type

The Router Instance Type parameter specifies the type of router instance that is used as a collector source.

The options are:

- Base (default)
- VPRN

A.55 Network parameters

A.55.1 Use Queue Group

The Use Queue Group parameter specifies whether a queue from an egress queue group template is used in the egress forwarding class of a network policy.

The options are:

- Enabled
- Disabled (default)

When you set the Use Queue Group to Enabled, you must enter the egress queue ID specified in the queue group template policy in the Queue panel when applied to an IP interface.

A.56 Policer Control parameters

A.56.1 Cumulative MBS Contribution

The Cumulative MBS Contribution parameter specifies the maximum amount of cumulative buffer space (in bytes) allowed for a specific priority level by a policer control policy. The range is 1 to 134217728. The default is the maximum value.

A.56.2 Fixed MBS contribution

The Fixed MBS Contribution parameter specifies if the cumulative buffer space is fixed for a specific priority level. When this parameter is set to True for a priority level, the system treats the defined Cumulative MBS Contribution value as an explicit MBS definition for the priority level.

A.57 RADIUS Based Accounting parameters

A.57.1 Enable

The Enable parameter specifies whether the interval for updating subscriber host information is enabled.

The options are:

- enabled
- disabled (default)

When the Enable parameter is set to disabled, the subscriber host information is not updated. When the parameter is set to enabled, you can configure the Value (minutes) parameter.

A.58 Residential Subscriber parameters

A.58.1 Assign Aggregate Rate Limit

The Assign Aggregate Rate Limit parameter specifies whether or not to use the aggregate rate limit to you set using the associated Aggregate Rate Limit (kbps) parameter.

The options are:

- Disabled (default)
- Enabled

A.58.2 Keep-Alive - Default

The Keep-Alive - Default parameter specifies whether default setting are used for the Keep-Alive-Interval (seconds), Keep-Alive-Timeout (seconds), and Keep-Alive-Retries (seconds) parameters. The Default parameter must be disabled in order to specify non-default settings for the Keep-Alive parameters.

The options are:

- Disabled
- Enabled (default)

A.58.3 Message-Retransmit - Default

The Message-Retransmit - Default parameter specifies whether default setting are used for the Message-Retransmit-Retries and Message-Retransmit-Timeout (seconds) parameters. The Default parameter must be disabled in order to specify non-default settings for the Message-Retransmit parameters.

The options are:

- Disabled
- Enabled (default)

A.58.4 No Constraint

The No Constraint parameter specifies whether the Mandatory Bandwidth (kbps) or Unconstrained Bandwidth (kbps) parameters are set to their default values, or can be configured.

The options are:

- disabled
- enabled (default)

When the No Constraint parameter is set to enabled, you cannot configure the associated parameters.

A.58.5 Use Ingress L2TP DSCP

The Use Ingress L2TP DSCP parameter specifies whether the DSCP value found in the L2TP header of ingress traffic of PPP subscribers is used for egress classification. When this parameter is enabled, egress traffic is assigned to a forwarding class based on the L2TP DSCP information.

The options are:

- enabled
- disabled (default)

A.59 Routing parameters

A.59.1 View the newly created Policy Statement

The View the newly created Policy Statement parameter specifies whether the properties form for the policy is displayed when the current configuration form is closed.

The options are:

- Enabled
- Disabled (default)

A.60 Service PW Template parameters

A.60.1 Egress Filter Type

The Egress Filter Type parameter specifies the type of egress filter to be used.

The options are:

- No filter defined (default)
- Egress IP
- Egress IPv6
- Egress MAC

A.60.2 Ingress Filter Type

The Ingress Filter Type parameter specifies the type of ingress filter to be used.

The options are:

- No filter defined (default)
- Ingress IP
- Ingress IPv6
- Ingress MAC

A.61 Time of Day parameters

A.61.1 Priority

The Priority parameter specifies the priority of the time of day suite entry. If there are overlapping time range entries within a time of day suite entry, the time range entry with the highest priority is run first. The range is 1 to 10, where 1 is the highest priority, and 10 is the lowest. The default is 5.

The Priority parameter value must be unique within a same policy type.

A.62 WRED Slope parameters

A.62.1 High Slope

The High Slope parameter specifies how the high-priority packets, also known as in-profile packets, access the shared portion of the buffer pool. The probability of the buffer pools discarding packets is determined by the RED slope. The RED slope is indicated on the client GUI using a graph.

The Start Avg., Max Avg., and Max Prob. parameters specify the values of the RED slope.

In general, packets that are in-profile fall within the configured CIR and PIR ranges. This means the packets are more likely to gain access to the shared buffer pool, and not be discarded because of a buffer overflow.

A.62.2 Low Slope

The Low Slope parameter specifies how the low-priority packets, also known as out-of-profile packets, access the shared portion of the buffer pool. The probability of the buffer pools discarding packets is determined by the RED slope. The RED slop is indicated on the client GUI using a graph.

The Start Avg., Max Avg., and Max Prob. parameters specify the values of the RED slope.

In general, packets that are out-of-profile fall outside configured CIR and PIR ranges. This means the packets should not be allowed to cause a buffer overflow in the shared buffer pool, and should be more likely to be discarded because of a buffer overflow.

A.62.3 Max Avg.

(maxAverage)

The Max Avg. parameter specifies the maximum average of the RED slope position. The parameter indicates that the packet discard probability has increased to 1, and packets are discarded. The range is 0 to 100%. The default is 75% buffer pool utilization before the packet discard probability is 100% for the low slope and 90% for the high slope. The parameter must be greater than or equal to the Start Avg. parameter.

For example, when out-of-profile packets arrive on an interface configured with a low slope policy that has the parameter set to 75% and the Start Avg. parameter set to 50%, out-of-profile packets may start to be discarded by the buffer when it reaches 50% full, and all packets are dropped when the buffer reaches 75% full. The probability of the packets being discarded before the Max Avg. parameter is reached is determined by the Max Prob. parameter.

A.63 Routing Instance parameters

A.63.1 Infinite

The Infinite parameter specifies whether or not the static port forward will have an infinite lifetime.

The options are:

- Enabled (default)
- Disabled

A.63.2 Notify When Pool Is Depleted

The Notify When Pool Is Depleted parameter specifies whether an alarm is raised when there are no IP addresses available in the IP address pool.

The options are:

- Enabled (default)
- Disabled

A.63.3 Suppress Duplicate Error

The Suppress Duplicate Error parameter specifies whether an error message appears when the NFM-P attempts to synchronize a static port forward that already exists on the backup router.

The options are:

- Enabled
- Disabled (default)

A.64 Interface parameters

A.64.1 Inherit from Network Element

The Inherit from Network Element check box specifies whether the interface inherits the value of the Operational State Transition Interval parameter from the NE.

The options are:

- Enabled (default)
- Disabled

When the check box is disabled, the Operational State Transition Interval parameter is configurable.

A.64.2 View the newly created Network Interface

The View the newly created Network Interface parameter specifies whether the properties form for the interface is displayed when the current configuration form is closed.

The options are:

- Enabled
- Disabled (default)

A.64.3 What type of interface would you like to create?

The What type of interface would you like to create? parameter specifies the type of network interface that you are creating.

The options are:

- MPLS (default)
- OSPFv2
- OSPFv3
- RIP
- ISIS
- LDP

A.65 IS-IS parameters

A.65.1 Remove Key button

The Remove Key button deletes the MD5 key.

A.66 L2TP parameters

A.66.1 Load Balance Method

The Load Balance Method parameter specifies whether the NFM-P performs load balancing on an LT2P tunnel group on a per-session, or per-tunnel basis.

Table A-36 Load Balance Method parameter

Option	Description
No	Do not perform load balancing. This option is available only for tunnel profiles.
Per Session	Perform load balancing on a per-session basis.
Per Tunnel	Terminate all L2TP sessions with the same local tunnel ID on one MS-ISA, and perform load balancing on a per-tunnel basis.

A.67 MPLS parameters

A.67.1 View the newly created MPLS path

The View the newly created MPLS path parameter specifies whether you want to view the configuration information for the newly created MPLS path using the MPLS Path form.

The options are:

- Enabled
- Disabled (default)

A.68 Network Domain parameters

A.68.1 Interface Association Count

The Interface Association Count parameter specifies the number of interfaces associated with the domain.

A.68.2 Routing Instance ID

The Routing Instance ID displays the ID of the routing instance in the site.

A.68.3 SDP Association Count

The SDP Association Count lists the number of SDP and tunnels associated with the domain.

A.68.4 Site ID

The Site ID is the ID in which the network domain was configured.

A.68.5 Network Interfaces tab

The Network Interfaces tab lists the details of all the interfaces associated with the network domain. You can enter details in each of the options or select an option from the dropdown menu to filter your criteria.

A.68.6 Service Tunnels tab

The Service Tunnels tab lists the details of all service tunnels and SDP associated with the network domain. You can enter details in each of the options or select an option from the dropdown menu to filter your criteria.

A.69 OSPF parameters

A.69.1 Change Password button

You can use the Change Password button to change the text password used to validate OSPF messages between neighbors. The button is enabled when the Authentication Type parameter is set to the Simple Password option.

A.70 RSVP parameters

A.70.1 Down Threshold (%)

The Down Threshold (%) parameters (1 through 16) specify the Down Threshold level percentages for reserved bandwidth per interface. Any reserved bandwidth change per interface is compared to the configured threshold levels. An IGP TE update is triggered if the configured threshold levels are crossed, as a result of either an LSP setup or teardown. Threshold levels configured for the NE at the RSVP level can also be inherited by all configured RSVP interfaces under the NE.

You can configure from one to sixteen Down Threshold percentage levels. If you configure one or more thresholds to a non-default value, then the configured Down Thresholds will be rearranged in descending order. The range is -1, 0 to 100, where -1 indicates that the threshold level is disabled.

Table A-37 Down Threshold (%) parameters

Down Threshold (%) parameter name	XML string	Default value
Down Threshold 1 (%)	teThresholdLevelDown1	100
Down Threshold 2(%)	teThresholdLevelDown2	99
Down Threshold 3 (%)	teThresholdLevelDown3	98
Down Threshold 4 (%)	teThresholdLevelDown4	97
Down Threshold 5 (%)	teThresholdLevelDown5	96
Down Threshold 6(%)	teThresholdLevelDown6	95

Table A-37 Down Threshold (%) parameters (continued)

Down Threshold (%) parameter name	XML string	Default value
Down Threshold 7 (%)	teThresholdLevelDown7	90
Down Threshold 8 (%)	teThresholdLevelDown8	85
Down Threshold 9 (%)	teThresholdLevelDown9	80
Down Threshold 10(%)	teThresholdLevelDown10	75
Down Threshold 11 (%)	teThresholdLevelDown11	60
Down Threshold 12 (%)	teThresholdLevelDown12	45
Down Threshold 13 (%)	teThresholdLevelDown13	30
Down Threshold 14 (%)	teThresholdLevelDown14	15
Down Threshold 15 (%)	teThresholdLevelDown15	0
Down Threshold 16 (%)	teThresholdLevelDown16	-1

A.70.2 Up Threshold (%)

The Up Threshold (%) parameters (1 through 16) specify the Up Threshold level percentages for reserved bandwidth per interface. Any reserved bandwidth change per interface is compared to the configured threshold levels. An IGP TE update is triggered if the configured threshold levels are crossed, as a result of either an LSP setup or teardown. Threshold levels configured for the NE at the RSVP level can also be inherited by all configured RSVP interfaces under the NE.

You can configure from one to sixteen Up Threshold percentage levels. If you configure one or more thresholds to a non-default value, then the configured Up Thresholds will be rearranged in ascending order. The range is -1, 0 to 100, where -1 indicates that the threshold level is disabled.

Table A-38 Up Threshold (%) parameters

Up Threshold (%) parameter name	XML string	Default value
Up Threshold 1 (%)	teThresholdLevelUp1	0
Up Threshold 2(%)	teThresholdLevelUp2	15
Up Threshold 3 (%)	teThresholdLevelUp3	30
Up Threshold 4 (%)	teThresholdLevelUp4	45
Up Threshold 5 (%)	teThresholdLevelUp5	60
Up Threshold 6(%)	teThresholdLevelUp6	75
Up Threshold 7 (%)	teThresholdLevelUp7	80
Up Threshold 8 (%)	teThresholdLevelUp8	85
Up Threshold 9 (%)	teThresholdLevelUp9	90
Up Threshold 10(%)	teThresholdLevelUp10	95
Up Threshold 11 (%)	teThresholdLevelUp11	96

Table A-38 Up Threshold (%) parameters (continued)

Up Threshold (%) parameter name	XML string	Default value
Up Threshold 12 (%)	teThresholdLevelUp12	97
Up Threshold 13 (%)	teThresholdLevelUp13	98
Up Threshold 14 (%)	teThresholdLevelUp14	99
Up Threshold 15 (%)	teThresholdLevelUp15	100
Up Threshold 16 (%)	teThresholdLevelUp16	-1

A.71 Static Routes parameters

A.71.1 Enable CPE Check

The Enable CPE parameter allows configuration of the CPE check parameters when it is enabled. The CPE check parameters are hidden when it is disabled. The default is disabled.

A.72 Accounting Policies parameters

A.72.1 Counters

The Counters parameter in the Egress panel of the CustomQueueConfig or CustomOverrideConfig form specifies the ingress statistics counters or override queue counters to monitor. See the Egress Counters parameter in this section for the parameter options.

A.72.2 Counters

The Counters parameter in the Ingress panel of the CustomQueueConfig or CustomOverrideConfig form specifies the ingress statistics counters or override queue counters to monitor. See the Ingress Counters parameter in this section.

A.73 Auto-Provision Profiles parameters

A.73.1 Network Element Type

The Network Element Type parameter specifies the type of NE that supports auto provisioning.

A.73.2 Network Element Version Information

The Network Element Version Information parameter specifies the release of the NE that supports auto provisioning.

A.73.3 View the newly created Auto-Provisioning

The View the newly created Auto-Provisioning parameter specifies whether you need to view the configuration details of the newly created Auto-Provision.

The options are:

- Enabled
- Disabled (default)

A.74 Bulk Operations parameters

A.74.1 Batch Size

The Batch Size parameter specifies the size of each batch. The range is 100 to 5000. The default is 2000.

A.74.2 Batch Status

The Batch Status parameter displays the bulk change status of all of the batch items. The value is not configurable.

If all of the batch items have the same bulk change status, the Batch Status parameter displays that status. If one or more of the batch items has a different bulk change status, the Batch Status parameter displays the Mixed status.

Table A-39 Batch Status descriptions

Status	Description
Cancelled	The operation has been manually stopped.
DB Failures	A database error occurred with the batch change.
Deployment Failure [1]	Deployment to the node failed.
Exception	An internal error occurred during the batch execution.
Execution Failures	An error occurred with the operation.
In Progress	The operation is executing.
Mixed	One or more of the batch items has a different bulk change status.
No Change	No objects were changed.
Not Applicable	—
Not Executed	The operation has never been executed before.
Not in User Span	Objects were not in the user span.
Object Not Found	The object to modify no longer exists in the NFM-P.
Queued	The operation is queued to execute.
Successful	The operation completed successfully.

Notes:

1. This status does not update automatically if or when the deployment failure clears.

A.74.3 Batch Status Summary

The Batch Status Summary parameter displays the status of all of the batches. The value is not configurable.

If all of the batches have the same batch status, the Batch Status Summary parameter displays that status. If one or more of the batches has a different batch status, the Batch Status Summary parameter displays the Mixed status.

See Table A-39, “Batch Status descriptions” (p. 3327) for a list of generated statuses and their descriptions.

A.74.4 Changed

The Changed parameter displays the number of objects that were modified in the batch. The value is not configurable.

A.74.5 Continue on Failure

The Continue on Failure parameter specifies whether the operation stops if a failure occurs or if the operation continues until the operation is complete regardless of failures.

The options are:

- Enabled
- Disabled (default)

A.74.6 Creator

The Creator parameter displays the name of the NFM-P user that created the last operation. The value is not configurable.

A.74.7 Duration

The Duration parameter specifies how long the last operation took to execute. The value is not configurable.

A.74.8 Execution Status

The Execution Status parameter displays the status of the bulk change operation. The value is not configurable.

Table A-40 Execution Status descriptions

Status	Description
Cancelled	The operation has been manually stopped.
Completed	The operation has been executed.

Table A-40 Execution Status descriptions (continued)

Status	Description
Generating...	Batch generation is in progress.
Generation Complete	Batch generation is complete.
In Progress	The operation is executing.
No Deployers Available	All available deployers are busy, or the deployment failed.
Not Applicable	—
Not Executed	The operation has never been executed.
Queued	The operation is queued for execution.

A.74.9 Failures

The Failures parameter displays the number of objects in the batch that failed to be changed. The value is not configurable.

A.74.10 Last Total Changed

The Last Total Changed parameter displays the number of objects that were changed by the last executed operation. The value is not configurable.

A.74.11 Not Changed

The Not Changed parameter displays the number of objects in the batch that did not change. The value is not configurable.

A.74.12 Not Found

The Not Found parameter displays the number of objects in the executed batches that could not be changed because they no longer exist in the NFM-P. The value is not configurable.

A.74.13 Not in Span

The Not in Span parameter displays the number of objects in the executed batches that could not be changed because they are not in the user span of control. The value is not configurable.

A.74.14 Object Type

The Object Type parameter specifies the object class to which the bulk change applies. Choose an object type using the pull-down list.

A.74.15 Operation ID

The Operation ID parameter specifies a unique numeric identifier for the operation. The parameter is configurable when the Auto-Assign ID parameter is disabled. The range is 1 to 65 535. The default is 0, which means that a value is not specified.

A.74.16 Range

The Range parameter displays the range of objects that are included in the batch. The range is not configurable.

A.74.17 Time Last Started

The Time Last Started parameter specifies the time that the last operation execution started. The value is not configurable.

A.74.18 Time Last Finished

The Time Last Finished parameter specifies the time that the last operation execution finished. The value is not configurable.

A.75 Card Migration Event Manager parameters

A.75.1 Auto Reboot

The Auto Reboot parameter specifies whether the NFM-P reboots a target NE automatically after the migration tasks on the NE are complete. The parameter is configurable at the card migration event level, which applies to each target NE, and at the target NE level, which applies only to the selected NEs.

The options are:

- Enabled
- Disabled (default)

A.75.2 New Type

The New Type parameter specifies the new IOM or MDA type, depending on which panel contains the parameter, and on the Current Type value in the panel.

The parameter options for IOMs are:

- No Change (default)
- 2 x XP MDA IOM 3

The following table lists the parameter options for MDAs. The default is No Change.

Table A-41 New Type parameter

Current Type	New Type
10 x 1-Gig Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
5 x 1-Gig Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX

Table A-41 New Type parameter (continued)

Current Type	New Type
1 x 10-Gig Ethernet	4 x 10Gig Extended Performance XFP 2 x 10Gig Extended Performance XFP 1 x 10Gig Extended Performance SFP
20 x 100 Ethernet Fx	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
20 x 10/100/1000 Ethernet Tx	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
2 x 10-Gig Ethernet XFP	2 x 10Gig Extended Performance XFP 4 x 10Gig Extended Performance XFP
20 x 10/100/1000 Ethernet SFP	20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
1 x 10-gig Ethernet XFP	4 x 10GigExtended Performance XFP 2 x 10Gig Extended Performance XFP 1 x 10Gig Extended Performance XFP
5 x 10/100/1000	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX
10 x 10/100/1000 Ethernet SFP	10 x 1Gig Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance SFP 20 x 10/100/1000 Ethernet Extended Performance TX

A.76 Copy/Move SAPs parameters

A.76.1 Continue on individual Failure

The Continue on individual Failure parameter specifies the action to be taken if a SAP copy or move operation fails. When the parameter is disabled, the operation stops after the first failure and all successful SAP copy or move operations are reversed. When the parameter is enabled, the operation continues until all SAP copy or move operations that can be completed successfully are finished.

The options are:

- Disabled (default)
- Enabled

A.76.2 Inner Encap Value End

The Inner Encap Value End parameter specifies the ending value of the inner encapsulation label. The range is 1 to 65 535 (default is 1) for ATM channels. Otherwise, the range is 0 to 4095 (default is 0).

A.76.3 Inner Encap Value Start

The Inner Encap Value Start parameter specifies the starting value of the inner encapsulation label. The range is 1 to 65 535 (default is 1) for ATM channels. Otherwise, the range is 0 to 4095 (default is 0).

A.76.4 Outer Encap Value End

The Outer Encap Value End parameter specifies the end value of the outer encapsulation label. The range is 0 to 4094. The default is 4094.

A.76.5 Outer Encap Value Start

The Outer Encap Value Start parameter specifies the start value of the outer encapsulation label. The range is 0 to 4094. The default is 0.

A.76.6 Service Type

The Service Type parameter specifies the type of service for the source interface SAPs.

When the Current Mode parameter is L2 Access Interface, the options are:

- All L2 Services (default)
- Apipe
- Epipe
- Ipipe
- VPLSs
- MVPLS

When the Current Mode parameter is L3 Access Interface or L3 Subscriber Interface SAP, the options are:

- All L3 Services (default)
- IES
- VPRN

A.77 NE Sessions parameters

A.77.1 Append to file

The Append to file parameter specifies whether the NFM-P appends the console output to the file specified by the A.77.7 “Log File Location” (p. 3334) parameter.

The parameter is configurable when the A.77.9 “Send Console To a File” (p. 3334) parameter is enabled.

Table A-42 Append to file parameter

Option	Description
disabled (default)	The NFM-P overwrites the file specified by the A.77.7 “Log File Location” (p. 3334) parameter using the console output.
enabled	The NFM-P appends the console output to the file specified by the A.77.7 “Log File Location” (p. 3334) parameter.

A.77.2 Background color

The Background color parameter specifies the background color of the console window. Click on the Set color button beside the parameter to choose a color using a palette form. The default is black.

A.77.3 Bold

The Bold parameter specifies whether the console window displays text in boldface type.

The options are:

- disabled (default)
- enabled

A.77.4 Font Name

The Font Name parameter specifies the typeface of the text in the console window.

The options are:

- Monospaced (default)
- DialogInput
- Lucida Console
- Lucida Sans Typewriter

A.77.5 Foreground color

The Foreground color parameter specifies the color of the text in the console window. Click on the Set color button beside the parameter to choose a color using a palette form. The default is white.

A.77.6 Italic

The Italic parameter specifies whether the text in the console window is italicized.

The options are:

- disabled (default)
- enabled

A.77.7 Log File Location

The Log File Location parameter specifies the path and name of the file that is to contain the console output. The NFM-P creates the file if it does not exist.

You can configure the parameter using one of the following methods.

- Manually type a file name.
- Click on the Change button and use the form that opens to specify a file.
- Click on the Set Default button to restore the default path and file name.

The default is *install_dir*/nms/log/client/*user_name*/cli_output.txt on a RHEL client station, and *install_dir*\nms\log\client*user_name*\cli_output.txt on a Windows client station

where

install_dir is the client installation directory, typically /opt/nsp/client on RHEL, and C:\nsp\client on Windows

user_name is the RHEL or Windows user name

The parameter is configurable when the A.77.9 “Send Console To a File” (p. 3333) parameter is enabled.

A.77.8 Minimum number of scrolling lines

The Minimum number of scrolling lines parameter specifies the maximum number of text lines that the console scroll buffer can contain. When the number of lines exceeds the parameter value, the NFM-P removes the oldest line from the buffer and the line is no longer displayed in the console window.

A.77.9 Send Console To a File

The Send Console To a File parameter specifies whether the NFM-P records the console output in a file.

The options are:

- disabled (default)
- enabled

A.78 Schedules parameters

A.78.1 Change Current User To

The Change Current User To parameter specifies the NFM-P user to which the NFM-P scheduled task is to be assigned. The options are all configured user accounts. The default is admin, which is the default NFM-P account.

A.78.2 Enable

The Enable parameter specifies whether to enable or disable the configurability of the Delay Time (seconds) parameter.

The options are:

- Enable
- Disable (default)

A.78.3 Run Every Days

The Run Every Days parameter specifies whether the NFM-P schedule runs every *n* days, where *n* is a numerical value that you set using the up and down arrows. The parameter is configurable when the Frequency parameter is set to Per Day. The range is 1 to 2 147 483 647. The default is 1.

A.78.4 Run Every Hours

The Run Every Hours parameter specifies whether the NFM-P schedule runs every *n* hours, where *n* is a numerical value set using the up and down arrows. The parameter is configurable when the Run Every Hour parameter is disabled. The range is 1 to 2 147 483 647. The default is 1.

A.78.5 Run Every Minutes

The Run Every Minutes parameter specifies that the NFM-P schedule runs every *n* minutes, where *n* is a numerical value set using the up and down arrows. The parameter is configurable when the Run Every Minute parameter is disabled.

A.78.6 Run Every Months

The Run Every Months parameter specifies that the NFM-P schedule runs every *n* months, where *n* is a numerical value set using the up and down arrows. The parameter is configurable when the Frequency parameter is set to Per Month. The range is 1 to 2 147 483 647. The default is 1.

A month is based on a 30-day interval.

A.78.7 Run Every Seconds

The Run Every Seconds parameter specifies that the NFM-P schedule runs every *n* seconds, where *n* is a numerical value set using the up and down arrows. The parameter is configurable when the Run Every Second parameter is disabled. The range is 1 to 2 147 483 647. The default is 1.

A.78.8 Time Alignment Setting

The Time Alignment Setting parameter specifies the base minute with which the NFM-P schedule aligns. The parameter is used with the A.78.5 “Run Every Minutes” (p. 3335) parameter. The range is 0 to 59. The default is 0.

A.79 Scripts parameters

A.79.1 Answer

The Answer parameter specifies the answer to a command sent by the system in response to a question intervention tag.

A.80 Scripts parameters

A.80.1 Default Value

The Default Value parameter specifies the default value that is associated with the parameter tag or intervention tag in the CLI script. The range is 0 to 255 characters. The characters ‘ ( ) ? / are not allowed.

A.80.2 Label

The Label parameter specifies the CLI script parameter tag in a CLI script. The range is 1 to 255 characters. Only alphanumeric characters, underscores, and dashes are allowed.

A.80.3 Question

The Question parameter specifies a question that is expected in response to a command in a script. The system responds to the question with the answer specified by the A.79.1 “Answer” (p. 3336) parameter.

A.80.4 Type

The Type parameter specifies the category of the script. The range is 0 to 255 characters.

A.81 Server Performance Statistics parameters

A.81.1 Accounting Stats Failure Periodic Threshold

The Accounting Stats Failure Periodic Threshold parameter specifies the number of accounting statistics failures that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. Accounting statistics failures can be caused by conditions such as failed file transfers. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.2 Accounting Stats Pending Periodic Threshold

The Accounting Stats Pending Periodic Threshold parameter specifies the number of accounting statistics that are pending to be processed that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.3 Accounting Stats Processed Periodic Threshold

The Accounting Stats Processed Periodic Threshold parameter specifies the number of accounting statistics processed that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3336) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.4 Accounting Stats Total Periodic Threshold

The Accounting Stats Total Periodic Threshold parameter specifies the number of accounting statistics received that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3336) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.5 Alarm Total Periodic Threshold

The Alarm Total Periodic Threshold parameter specifies the number of alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3336) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.6 Average Processing Time Periodic Threshold

The Average Processing Time Periodic Threshold parameter specifies the average processing time for AA statistics files during the interval specified by the A.81.8 “Collection Interval” (p. 3336) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.7 Cleared Periodic Threshold

The Cleared Periodic Threshold parameter specifies the number of cleared alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3336) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.8 Collection Interval

(pollingInterval)

The Collection Interval parameter specifies the frequency with which the statistics counters for the selected statistics class are logged. The default is 15 minutes.

A.81.9 Condition Periodic Threshold

The Condition Periodic Threshold parameter specifies the number of condition alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.10 Critical Periodic Threshold

The Critical Periodic Threshold parameter specifies the number of critical alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.11 Dropped Backpressure Periodic Threshold

The Dropped Backpressure Periodic Threshold parameter specifies the number of traps dropped because of backpressure during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. Backpressure conditions occur when the NFM-P server is very busy processing other traps. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.12 Dropped Duplicate Periodic Threshold

The Dropped Duplicate Periodic Threshold parameter specifies the number of duplicate SNMP traps dropped during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.13 Dropped Full Resync Periodic Threshold

The Dropped Full Resync Periodic Threshold parameter specifies the number of SNMP traps that are not processed because of a pending full resync during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.14 Dropped Not Managed Periodic Threshold

The Dropped Not Managed Periodic Threshold parameter specifies the number of SNMP traps associated with an unmanaged NE dropped during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.15 Dropped Out Of Sequence Periodic Threshold

The Dropped Out Of Sequence Periodic Threshold parameter specifies the number of SNMP traps with non-sequential trap IDs dropped during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.16 File Accounting Stats Pending Periodic Threshold

The File Accounting Stats Pending Periodic Threshold parameter specifies the number of accounting files transferred from NEs that are waiting to be processed. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.17 File Accounting Stats Processed Periodic Threshold

The File Accounting Stats Processed Periodic Threshold parameter specifies the number of accounting statistics files that have been read and processed into statistics records. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.18 File Accounting Stats Total Periodic Threshold

The File Accounting Stats Total Periodic Threshold parameter specifies the number of accounting statistics files that have been transferred from NEs. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.19 Incoming Periodic Threshold

The Incoming Periodic Threshold parameter specifies the number of SNMP traps received during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the SNMP Trap counter reaches this number, a threshold-crossing alarm is raised.

A.81.20 Indeterminate Periodic Threshold

The Indeterminate Periodic Threshold parameter specifies the number of indeterminate alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.21 Info Periodic Threshold

The Info Periodic Threshold parameter specifies the number of log entries for info alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.22 Major Periodic Threshold

The Major Periodic Threshold parameter specifies the number of major alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.23 Maximum Processing Time Periodic Threshold

The Maximum Processing Time Periodic Threshold parameter specifies the maximum processing time for AA statistics files during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.24 Minimum Processing Time Periodic Threshold

The Minimum Processing Time Periodic Threshold parameter specifies the minimum processing time for AA statistics files during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.25 Minor Periodic Threshold

The Minor Periodic Threshold parameter specifies the number of minor alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.26 Scheduled Polling Stats Pending Periodic Threshold

The Scheduled Polling Stats Pending Periodic Threshold parameter specifies the number of scheduled SNMP statistics awaiting processing that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.27 Scheduled Polling Stats Processed Periodic Threshold

The Scheduled Polling Stats Processed Periodic Threshold parameter specifies the number of processed scheduled SNMP statistics that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.28 Scheduled Polling Stats Total Periodic Threshold

The Scheduled Polling Stats Total Periodic Threshold parameter specifies the number of scheduled SNMP statistics received that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.29 Scheduled Resync Failure Periodic Threshold

The Scheduled Resync Failure Periodic Threshold parameter specifies the number of scheduled node resync failures that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.30 Scheduled Resync Processed Periodic Threshold

The Scheduled Resync Processed Periodic Threshold parameter specifies the number of scheduled processed node resyncs that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.31 Scheduled Resync Received Periodic Threshold

The Scheduled Resync Received Periodic Threshold parameter specifies the number of scheduled node resyncs that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.32 Scheduled Stats Failure Periodic Threshold

The Scheduled Stats Failure Periodic Threshold parameter specifies the number of scheduled SNMP statistics processing failures that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. Statistics failures may occur if an NE is unreachable or if collection takes longer than the collection interval time. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.33 Total Processing Time Periodic Threshold

The Total Processing Time Periodic Threshold parameter specifies the total processing time for AA statistics files during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.34 Unscheduled Resync Failure Periodic Threshold

The Unscheduled Resync Failure Periodic Threshold parameter specifies the number of unscheduled node resync failures that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.35 Unscheduled Resync Processed Periodic Threshold

The Unscheduled Resync Processed Periodic Threshold parameter specifies the number of unscheduled processed node resyncs that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.36 Unscheduled Resync Received Periodic Threshold

The Unscheduled Resync Received Periodic Threshold parameter specifies the number of unscheduled node resyncs that the Node Resync counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.37 Used Heap Memory Periodic Threshold

The Used Heap Memory Periodic Threshold parameter specifies the amount of heap memory currently in use that the NFM-P Memory counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.38 Unscheduled Polling Stats Pending Periodic Threshold

The Unscheduled Polling Stats Pending Periodic Threshold parameter specifies the number of unscheduled SNMP statistics awaiting processing that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.39 Unscheduled Polling Stats Processed Periodic Threshold

The Unscheduled Polling Stats Processed Periodic Threshold parameter specifies the number of processed unscheduled SNMP statistics that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.40 Unscheduled Polling Stats Total Periodic Threshold

The Unscheduled Polling Stats Total Periodic Threshold parameter specifies the number of unscheduled SNMP statistics received that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.41 Unscheduled Stats Failure Periodic Threshold

The Unscheduled Stats Failure Periodic Threshold parameter specifies the number of unscheduled SNMP statistics processing failures that the Stats Collection counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.42 Used Non Heap Memory Periodic Threshold

The Used Non Heap Memory Periodic Threshold parameter specifies the amount of non-heap memory currently in use that the NFM-P Memory counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.81.43 Warning Periodic Threshold

The Warning Periodic Threshold parameter specifies the number of warning alarms that the Alarm Rate counter records during the interval specified by the A.81.8 “Collection Interval” (p. 3337) parameter. When the counter reaches this number, a threshold-crossing alarm is raised.

A.82 Statistics Manager parameters

A.82.1 Default Polling Interval (seconds)

The Default Polling Interval (seconds) parameter specifies the default interval, in seconds, for the polling interval for real-time statistics in the Statistics Plotter form.

A.82.2 Maximum Data Retention Time (seconds)

The Maximum Data Retention Time (seconds) parameter specifies the number of seconds to keep statistics data in the Statistics Plotter form.

A.82.3 Retention Time (hours)

The Retention Time (hours) parameter specifies the time, in hours, that the database stores the statistics associated with a statistics policy.

A.82.4 Statistics Type

The Statistics Type parameter specifies the type of logged data that is displayed in the search list.

Table A-43 Statistics Type parameter

Option	Description
Current Data (default)	The most recent available statistics record
Statistics Policy	The statistics policy configured for the statistics class
Statistics Record	All available records for the statistics class

Note: Following an upgrade to NSP Release 24.8, all historical CPU Utilization Statistics will be deleted from the NFM-P system/database.

A.82.5 Threshold Reporting State

The Threshold Reporting State parameter specifies whether to generate threshold alarms when the number of log entries exceeds the Max Log Records parameter value. The options are:

- Up (default)
- Down

A.83 Service Test Manager parameters

A.83.1 Count

The Count parameter specifies the number of frames to be sent during an OmniSwitch ping. The range is 0 to 2 147 483 647. The default is 6.

A.83.2 Include Falling Threshold

The Include Falling Threshold parameter specifies whether to generate a threshold-crossing alarm for an OAM test if the value falls below a specified level.

The options are:

- Enabled
- Disabled (default)

A.83.3 Interval (seconds)

The Interval (seconds) parameter specifies the polling interval that is used for the ping. The range is 1 to 10 000. The default is 1.

A.83.4 Maximum Hop

The Maximum Hop parameter specifies the maximum number of hops in the path to the target OmniSwitch. The range is 0 to 2 147 483 647. The default is 5.

A.83.5 Source Site ID

The Source Site ID parameter specifies the IP address of the service tunnel origin. Specify an IP address in dotted-decimal format. The default is 0.0.0.0, which means that the parameter is not configured.

A.83.6 Test Interval (minutes)

The Test Interval (minutes) parameter specifies the time delay, in minutes, between contiguous tests. The range is 10 to 1500. The default is 10.

A.83.7 Timeout (seconds)

The Timeout (seconds) parameter specifies the time, in s, to wait for a message response. All reply messages that arrive after the specified time are silently discarded. The range is 1 to 100. The default is 5.

A.83.8 Timeout (minutes)

The Timeout (seconds) parameter specifies the number of minutes to wait for the test suite completion. The range is 1 to 600. The default is 15.

A.84 Time Range Entry Assignment parameters

A.84.1 End Date

The End Date parameter specifies the end date and time of the time range entry assignment analysis. You can use the up and down arrows to change the year, month, day, hour, or minute. The default value is the current system time of the NFM-P server.

A.84.2 Search by Time Of Day Entry Type

The Search by Time Of Day Entry Type parameter specifies whether the time range entry assignment analysis tool uses the A.7.2 “Password Change Required” (p. 3264) parameter as an input criterion.

The options are:

- enabled
- disabled (default)

A.84.3 Start Date

The Start Date parameter specifies the start date and time of the time range entry assignment analysis. You can use the up and down arrows to change the year, month, day, hour, or minute. The default value is the current system time of the NFM-P server.

A.84.4 Time Of Day Entry Policy Type

The Time Of Day Entry Policy Type parameter specifies the type of time of day entry policy that the time range entry assignment analysis tool uses as an input criterion. The options are:

- Ingress IPV6 Filter
- Egress MAC Filter
- Egress QoS Scheduler Policy
- Ingress IP Filter
- Egress IPV6 Filter
- Ingress QoS Scheduler Policy
- Egress QoS Policy
- Ingress QoS Policy

A.84.5 Time Range Entry Container Type

The Time Range Entry Container Type parameter specifies the type of policy string that the time range entry assignment analysis tool uses as an input criterion.

The options are:

- L2 Access Interface
- L3 Access Interface
- Aggregation Suite
- Time Of Day Suite

A.85 CPAM parameters

A.85.1 Destination FEC

The Destination FEC parameter specifies the IP address and subnet mask of the destination router or link in a CSPF calculation. Enter an IPv4 address in dotted decimal format or use the Select button to list and choose a destination router or link.

A.85.2 Ethernet Error

The Ethernet Error parameter specifies whether the audit verifies the error rate of the IP link against a specified threshold. The options are:

- Enabled (default)
- Disabled

A.85.3 First IP

The First IP parameter specifies the IP address of the source router. Enter the IPv4 address of the router in dotted-decimal format or use the Select button to list and choose a router.

A.85.4 Group IP

The Group IP parameter specifies the IP address of a multicast group to be highlighted.

A.85.5 In Profile Drop

The In Profile Drop parameter specifies whether the audit verifies the in profile drop rate for FC queues 1 to 8 against a specified threshold. The options are:

- Enabled
- Disabled (default)

A.85.6 Interface Error

The Interface Error parameter specifies whether the audit verifies the error rate of the network interface against a specified threshold. The options are:

- Enabled (default)
- Disabled

A.85.7 Least Fill Min Threshold %

The Least Fill Min Threshold % parameter specifies the minimum threshold of least-utilized links that can be used in the ECMP tie-breaking process.

The Least Fill Min Threshold % parameter is configurable only when the A.85.14 “Request Type” (p. 3347) parameter is set to Least_Fill.

A.85.8 Line - Severely Errored Seconds

The Line - Severely Errored Seconds parameter specifies whether the audit verifies the counter associated with the number of severely errored framing seconds found by a SONET/SDH line in the current 15 minute interval. The options are:

- Enabled
- Disabled (default)

A.85.9 Mask

The Mask parameter specifies the subnet mask length for the multicast group prefix. This parameter works in conjunction with the Multicast Group Prefix parameter. The range is 0 to 32. The default is 0.

A.85.10 Multicast Group Prefix

The Multicast Group Prefix parameter specifies the prefix for a multicast group.

A.85.11 Out of Profile Drop

The Out of Profile Drop parameter specifies whether the audit verifies the out of profile drop rate for FC queues 1 to 8 against a specified threshold. The options are:

- Enabled
- Disabled (default)

A.85.12 Prefix Length

The Prefix Length parameter, when combined with an IP address value, specifies the local IP prefix. The range is 0 to 32. The default is 24.

A.85.13 Protocol

The Protocol parameter specifies the multicast protocol of the routers to be highlighted. The options are:

- PIM
- IGMP (default)
- PIM or IGMP

A.85.14 Request Type

The Request Type parameter specifies how CSPF selects a path in the tie-breaking process for ECMP if there are several equal cost paths. The parameter options are described in the following table:

Table A-44 Request Type parameter

Option	Description	Dependencies
All	Traffic is sent over all of the matching paths.	—
Random	Traffic is sent over one randomly selected path.	—
Least_Fill	Traffic is sent over the path with the least-utilized links. The CPAM considers the value of the A.85.7 “Least Fill Min Threshold %” (p. 3346) only when this option is selected.	—

A.85.15 Required Bandwidth

The Required Bandwidth parameter specifies the minimum available bandwidth of the calculated CSPF. The range is 0, or 1 to 4 294 967 295. A value of zero means that the CSPF calculation includes no minimum bandwidth requirements.

A.85.16 Second IP

The Second IP parameter specifies the IP address of the destination router. Enter the IPv4 address of the router in dotted-decimal format or use the Select button to list and choose a router. On an ISIS topology map, an IPv6 router address can be entered in colon-hexadecimal format.

A.85.17 Section - Severely Errored Framing Seconds

The Section - Severely Errored Framing Seconds parameter specifies whether the audit verifies the counter associated with the number of severely errored framing seconds found by a SONET/SDH section in the current 15 minute interval. The options are:

- Enabled
- Disabled (default)

A.85.18 Section - Severely Errored Seconds

The Section - Severely Errored Seconds parameter specifies whether the audit verifies the counter associated with the number of severely errored seconds found by a SONET/SDH section in the current 15 minute interval. The options are:

- Enabled
- Disabled (default)

A.85.19 Source IP

The Source IP parameter specifies the IP address of the source router or link in a CSPF calculation. Enter an IPv4 address in dotted decimal format or use the Select button to list and choose a source router or link.

A.85.20 SRLG Strict

The SRLG Strict parameter specifies whether the CSPF highlight is calculated if the CPAM does not find path with SRLG. If the SRLG Strict parameter is enabled, CSPF tries to find paths with excluded SRLG. If no paths are found, CSPF returns an empty result.

If the SRLG Strict parameter is disabled, CSPF tries to find paths with excluded SRLG. If no paths are found, CSPF returns paths that ignore the SRLG constraint, if any.

The options are:

- Enabled
- Disabled (default)

A.85.21 SRLG Value

The SRLG Value parameter specifies the SRLG value that must be excluded from a CSPF calculation. The range is 1 to 4 294 967 295.

A.85.22 Stats Capture Delay

The Stats Capture Delay parameter specifies the delay, in seconds, between first capture and the second capture of the statistics. The range is 15 to 300. The default is 30.

A.85.23 Total Drop

The Total Drop parameter specifies whether the audit verifies the drop rate for FC queues 1 to 8 against a specified threshold. The options are:

- Enabled (default)
- Disabled

A.85.24 Threshold

The Threshold parameter specifies the threshold against which the audit parameter is compared. The files that enable customization of the default Threshold values can be found in /opt/nsp/nfmp/server/nms/config/audit/params on an NFM-P main server. The following table lists the defaults for each audit parameter:

Table A-45 Threshold parameter default values

Audit parameter	Default	Measure
A.85.2 “Ethernet Error” (p. 3345)	0.5	%
A.85.5 “In Profile Drop” (p. 3345)	2.5	%
A.85.6 “Interface Error” (p. 3346)	0.5	%
A.85.8 “Line - Severely Errored Seconds” (p. 3346)	0	seconds
A.85.11 “Out of Profile Drop” (p. 3346)	2.5	%
A.85.18 “Section - Severely Errored Seconds” (p. 3348)	0	seconds
A.85.17 “Section - Severely Errored Framing Seconds” (p. 3347)	0	seconds
A.85.23 “Total Drop” (p. 3348)	2.5	%
A.85.25 “Utilization” (p. 3349)	97.5	%

A.85.25 Utilization

The Utilization parameter specifies whether the audit verifies the utilization rate of the network interface or IP link against a specified threshold. The options are:

- Enabled (default)
- Disabled