NOKIA

NSP
Network Services Platform

Network Resource Controller - Packet (NRC-P)
Network Resource Controller - Cross domain (NRC-X)
Network Services Director
Release 18.12

User Guide

3HE-14122-AAAD-TQZZA
Issue 1
December 2018

Legal notice

Contents

About this document

Purpose

The *NSP NSD and NRC User Guide* serves as an introduction to the NSD and NRC modules of the NSP, and provides detailed information regarding their operation.

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:

- Customer Documentation Welcome Page
- Technical support

How to comment

Documentation feedback

- Documentation Feedback

Part I: Getting started

Overview

Purpose

This volume describes new feature content for the NSD and NRC release, and describes the tenancy mechanism that governs authentication.

Contents

1 Safety information

1.1 Structure of safety statements

1.1.1 Overview

This topic describes the components of safety statements that appear in this document.

1.1.2 General structure

Safety statements include the following structural elements:

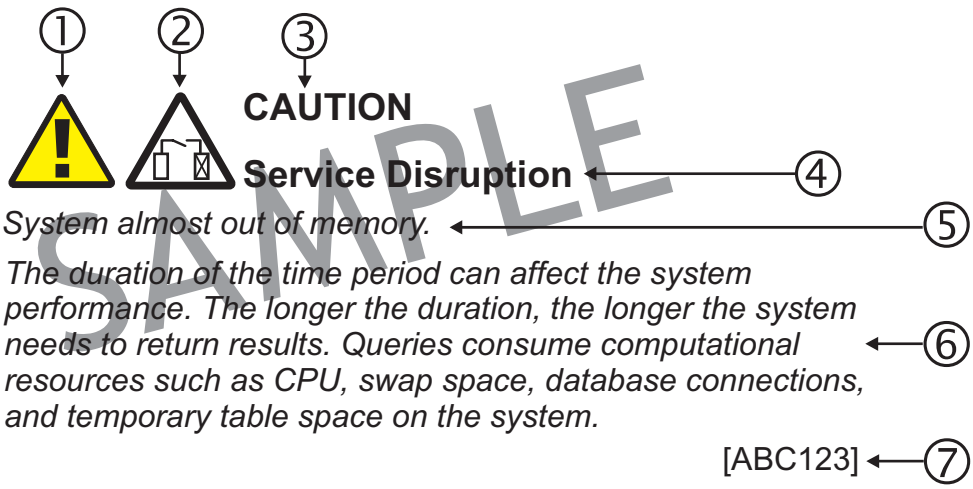

Item	Structure element	Purpose
1	Safety alert symbol	Indicates the potential for personal injury (optional)
2	Safety symbol	Indicates hazard type (optional)
3	Signal word	Indicates the severity of the hazard
4	Hazard type	Describes the source of the risk of damage or injury
5	Safety message	Consequences if protective measures fail
6	Avoidance message	Protective measures to take to avoid the hazard
7	Identifier	The reference ID of the safety statement (optional)

1.1.3 Signal words

The signal words identify the hazard severity levels as follows:

Signal word	Meaning
DANGER	Indicates the described activity or situation may, or will, represent a potential for a serious injury.
WARNING	Indicates the described activity or situation may, or will, cause equipment damage or serious performance problems.
CAUTION	Indicates the described activity or situation may, or will, cause service disruption.
NOTICE	Indicates the described activity or situation may, or will, cause minor performance problems.

2 What's new?

2.1 What's new in NSP Release 18

2.1.1 Introduction

This chapter highlights new features for NSP Release 18 and provides pointers into the documentation for more information. The *NSP NSD and NRC Release Description* provides Committed feature lists for all Release 18 deliveries.

2.1.2 Maintenance releases

Some maintenance releases may not be listed in this section, either because no new features are introduced or because the introduced features do not require documentation.

2.1.3 What's new in NSD/NRC Release 18.12

The following table lists the features added in Release 18.9 for the NSD and NRC modules.

Table 2-1 NSD/NRC 18.12 features

Feature	Description and reference
NRC-P features	
NSP-133179 — Latency Steering Optimizer - Support for multi-homed ASBR	This feature adds support for multi-homing to the current peer engineering Latency Steering Optimizer (LSO) application, which allows for the selection of exact egress peer link. See the *Latency Steering Optimizer* application.
NSPF-148215 — Providing Kafka notifications to NRC-X	Kafka notifications are provided to applications such as NRC-X.
NSPF-153424 — Latency-based LSP optimization enhancements	NSP path routing and optimization automatically computes an LSP path using specified objectives and constraints. See the *IP/MPLS Optimization* application.
NSPF-153451 — LSP optimization using live telemetry IP stats	When optimizing LSPs, the NSP can move LSPs off congested interfaces to interfaces that have less congestion using real time-measured values. See the *IP/MPLS Optimization* application.

Table 2-1 NSD/NRC 18.12 features (continued)

Feature	Description and reference
NSPF-154105 — NRC-P Sim - Additional simulation options	Additional functionality has been added to the NRC-P Sim. See the *IP/MPLS Simulation* application.
NSPF-154178 — PCE-Initiated LSPs - Hardening and path behavior	The NSP supports the configuration and the creation of SR-TE LSPs via PCEP. See the *IP/MPLS Optimization* application and 5.3.4 “PCE-initiated LSPs” (p. 42).
NSPF-155604 — Inbound Route Optimizer - Operator-triggered mode	This feature supports Operator-triggered mode within the Inbound Route Optimizer application, which allows inbound traffic to be controlled through the use of BGP communities that influence where incoming traffic needs to be received. See the *Inbound Route Optimizer* application.
NRC-X features	
NSPF-123374 — Diversity analytics	Diversity analysis in NRC-X now has the ability to analyze the groups of services if they are diverse to one another. See the *Cross Domain Coordinator* application.
NSPF-144269 — NRC-X bottom-up navigation	The NRC-X user can view a list of the IP services that are running over a physical link. See the *Cross Domain Coordinator* application.
[NSPF-156257 — 3D map enhancements	The NRC-X supports a 3D map and linear diagrams for Link Layer Interconnects. See the *Cross Domain Coordinator* application.
NSD features	
NSPF-125388 — Description field on services, endpoints, and tunnels	A Description field can be configured on services, service endpoints, and service tunnel objects within the Service Fulfillment application. See the *Service Fulfillment* application.
NSPF-140751 — Support for EVPN services on MDM-managed nodes	The Service Fulfillment application supports E-Line EVPN & E-LAN EVPN-based services on devices managed through MDM. See the *Service Fulfillment* application.

Table 2-1 NSD/NRC 18.12 features (continued)

Feature	Description and reference
NSPF-145026 — Workflow Manager application	The Workflow Manager (WFM) application is introduced in beta quality for network automation as part of the NSD module (requires the premium license). See the *Workflow Manager* application.
NSPF-145258 — MDM support for REST SBI	REST support has been added to MDM to be used as south-bound protocol. See the *NSP Developer Portal*.
NSPF-151281 — Brownfield discovery/resync for MDM service augmentation	This feature allows the Service Fulfillment application to discover values of augmented attributes from devices managed through MDM. See the *Service Fulfillment* application.
NSPF-151287 — Define port selection criteria on service templates for enhanced object filtering	Port Filter criteria can be configured on Service Templates managed in the Policy Management application. See the *Policy Management* application.
NSPF-151319 — Support service tunnels to non-system interfaces	Tunnel Endpoint Conditions entries can be configured on Tunnel Profiles in the Policy Management application. See the *Policy Management* application.
NSPF-152335 — MDC to support multiple changes in one submission	Operators are now able to perform multiple change operations (create, update, delete) as a single commit. See the *Modeled Device Configurator* application.
NSPF-153658 — Auto-creation of LSPs during L3 VPN service creation	The Service Fulfillment application automatically creates all static LSPs and SDPs required to configure and fulfill an L3 VPN service. See the *Service Fulfillment* application.
NSPF-154183 — Support for IES services (MDM)	The Service Fulfillment application supports the configuration of IES services on devices discovered through MDM (using NETCONF). See the *Service Fulfillment* application.

Table 2-1 NSD/NRC 18.12 features (continued)

Feature	Description and reference
NSPF-161017 — Cross-launching MDC	Operators are able to navigate directly from the Fault Management application to the alarmed object in the Modeled Device Configurator application. See the *Modeled Device Configurator* application.
NSPF-161042 — Support for augment and RPC YANG statements	MDM supports YANG augmentation, submodule referencing, RPC & Action operations from SROS and third party (Cisco and Juniper) YANG model definitions.
NSPF-161133 — MDC support for RESTCONF RPCs	The MDC RESTCONF API supports operations and actions defined in network device schemas, which are mounted on network device instances, based on draft-ietf-netmod-schema-mount. See the *NSP Developer Portal*.
NSPF-161135 — MDC support for RESTCONF YANG-PATCH (RFC 8072)	The MDC RESTCONF API supports yang-patch as defined in IETF RFC 8072. See the *NSP Developer Portal*.

2.1.4 What's new in NSD/NRC Release 18.9

The following table lists the features added in Release 18.9 for the NSD and NRC modules.

Table 2-2 NSD/NRC 18.9 features

Feature	Description and reference
NRC-X features	
NSPF-114718 — Multi-controller support infrastructure	NRC-X can interface with multiple transport controllers. See the *Cross Domain Coordinator* application.
NSPF-132571 — 3D topology rendering	LLI links can be viewed on a multi-layer map. See the *Cross Domain Coordinator* application.
NSPF-138093 — Support for separate NRC-T	An list of existing optical services can be viewed. See the *NSP Deployment and Installation Guide*.

Table 2-2 NSD/NRC 18.9 features (continued)

Feature	Description and reference
NSPF-147907 — More optical service details	NRC-X can interface with an NRC-T system that shares a host server with the NFM-T product. See the *Cross Domain Coordinator* application.
NSPF-148605 — LLI and cross domain link extensions	LLI links can be un-managed rather than deleted, preserving their underlying infrastructure, when the cross domain link they are riding is deleted. See the *Cross Domain Coordinator* application.
NSD features	
NSPF-114870 — Support for IES services (NFM-P)	Support for IES services configured on Nokia SR OS devices using NFM-P mediation. 8.7 “IES services” (p. 76)
NSPF-139909 — Support for L3 VPN services (uWave, Wavence, NFM-P)	Support for L3 VPN services on uWave nodes (Wavence) using NFM-P mediation. 8.8.3 “L3 VPN services on Wavence SM NEs” (p. 77) 8.8.4 “Composite L3 VPN service across multiple Wavence SM domains and one 7x50 SR domain” (p. 78)

2.1.5 What's new in NSD/NRC Release 18.6

The following table lists the features added in Release 18.6 for the NSD and NRC modules.

Table 2-3 NSD/NRC 18.6 features

Feature	Description and reference
NRC-P features	
NSPF-112619 — Next hop tracking	This feature allows for real-time tracking of next hops within the Autonomous System Optimizer application. See the *Autonomous System Optimizer* application's help page.
NSPF-128268 — Support for operator-directed routes	The NSP provides operators with the ability to modify existing LSP paths by providing an explicit set of interfaces (IP links). See the *IP/MPLS Optimization* application's help page.

Table 2-3 NSD/NRC 18.6 features (continued)

Feature	Description and reference
NSPF-131691 — Nominal path creation	This feature ensures that a path computed for the LSP via the controller (only works for delegated LSPs) is the only computed path. See the *IP/MPLS Optimization* application's help page.
NSPF-133696 — Ingress Peer Optimizer support for automatic mode	This feature provides policy abstraction, and network control automation (or semi-automation with real-time traffic monitoring) and correlation for controlling inbound traffic. See the *Ingress Peer Optimizer* application's help page.
NSPF-138865 — GA for RSVP LSP live bandwidth via telemetry optimization	Real-time stats indicating actual utilization and LSP bandwidth may be used for path optimization if the information is retrieved periodically and frequently. See the *NSP NSD and NRC Release Description* for more information.
NSPF-199390 — Simulating network import and link failure	NRC-P now provides the ability to simulate changes in the IP topology that was discovered by the IP/MPLS Optimization application. The simulation application is run in a separate VM and imports the IP topology and LSPs from the NRC-P IP/MPLS Optimization application. See the *IP/MPLS Simulation* application's help page.
NRC-X features	
NSPF-141045 — Additional support for optical service types	The NRC-X module has been extended to support additional optical service configurations within a Nokia optical network. Regarding control plane services, support has been added for L0 PRC, L1, L1 PRC, and MRN. In addition, the NRC-X module now supports management plane services (unprotected and OPSA). See the *Cross Domain Coordinator* application's help page.
NSPF-143208 — Link Layer Interconnect extensions	The NRC-X module correlates link layer interconnects between router ports with the underlying optical transport network. With this feature, further optical service types and parameters are supported in the correlation between optical and IP topology. See the *Cross Domain Coordinator* application's help page.

Table 2-3 NSD/NRC 18.6 features (continued)

Feature	Description and reference
NSPF-144272 — Cross Domain Coordinator application improvements	This feature includes various improvements of the Cross Domain Coordinator application, particularly when large networks are managed. See the *Cross Domain Coordinator* application's help page.
NSD features	
NSPF-114792 — Support of all service attributes not modelled on NSD using new MFM framework	This feature uses the new multi-vendor framework to allow for the flexible configuration of any attributes on any NE. 7.2 “MDM” (p. 61) 9.5 “Mediation Profiles” (p. 87)
NSPF-114795 — NSD: Flexible support of all Nokia service attributes	This feature extends the new multi-vendor framework to support model augmentation for NEs managed by NFM-P Nokia routers. 7.2 “MDM” (p. 61) 9.5 “Mediation Profiles” (p. 87)
NSPF-118414 — Support of new model-driven mediation framework on NSD	The NSD uses the new mediation framework for service fulfillment purposes. 7.2 “MDM” (p. 61) 9.5 “Mediation Profiles” (p. 87)
NSPF-118419 — NSD: Support of MC-LAG and PW redundancy (demo quality)	This feature introduces support of MC-LAG and pseudowire redundancy for the E-LINE service at demo quality. 8.4.4 “E-Line service with MC-LAG termination and pseudowire redundancy” (p. 71)
NSPF-119494 — NSD: IPv6 support for L3VPN services	This feature implements the support of IPv6 addresses for L3 VPN services. Users can now enter IPv6 addresses during the L3 VPN service creation. See the *Service Fulfillment* application.
NSPF-123335 — NSD: Service flexibility using MFM adaptor templates	This feature enables users to select what MFM adaptor template to use. The functionality supports E-LAN, E-Line, and L3 VPN services. 9.5 “Mediation Profiles” (p. 87)
NSPF-129726 — NSD: Service Save (GA quality)	The NSD extends the functionality to save services without deployment to all supported service types. 8.1.2 “Saving service configuration without deployment” (p. 65)

Table 2-3 NSD/NRC 18.6 features (continued)

Feature	Description and reference
NSPF-129910 — NSD: Combine TSP and TCP (Tunnel Selection and Creation Profile)	This feature streamlines the usage of different tunnel types and merges the functionality of the tunnel creation policy into the tunnel selection policy. See the *Policy Management* page of the help system.
NSPF-131107 — NSD: Allow group id assignment for pcc-initiated lsp	This feature introduces a new E-Line parameter (Diverse From) that allows users to specify an existing NSD-managed E-Line service from which the new E-Line service must be diverse. See the *Service Fulfillment* page of the help system.
NSPF-131530 — NSD: Support of brownfield SDP and LSP of any transport type (GA quality)	This feature delivers support of LSPs created in the NFM-P (brownfield LSPs) of any type at general-availability quality. 7.1.5 “Brownfield LSP and SDP tunnels” (p. 59)
NSPF-134294 — NSD: Improve scale and performance	This feature improves the NSD scalability and performance in multiple areas.
NSPF-138099 — Ability to define Service ID for services created through REST API	This feature enables NSD users to configure a service ID for the L3 VPN, C-Line, E-LAN and E-Line services using the REST API.
NSPF-140792 — Link maintenance mode toggle via REST API and UI	The System IP MPLS Configuration policy allows you to configure the maintenance mode of an IP link. 9.4.1 “System IP MPLS Configuration” (p. 86)
NSPF-141446 — GQP functionality with MDM	This feature allows the NSD to create services and assign Generic QoS Policies (GQP) when the mediation layer is NFM-P or MDM, or both. 9.2.5 “Generic QoS Policies” (p. 84)
NSPF-141747 — LLDP Link Discovery	This feature enables the NSD to discover Physical Links based on LLDP information, and to populate the discovered links to the common store so that other applications can access them. 7.2.3 “LLDP link discovery” (p. 61)
NSPF-142152 — Deprecation of the LAG Service	This feature implements the deprecation and removal of the LAG service from the Service Fulfillment application and of the LAG policy from the Policy Management application.
NSPF-142848 — NSD: Handling the failed service creation due to validation	This feature enables the Service Fulfillment application to display specific error messages for service form fields during the service configuration.

Table 2-3 NSD/NRC 18.6 features (continued)

Feature	Description and reference
NSPF-144277 — Maps Scale Improvements	This features introduces miscellaneous improvements to the framework that supports the network maps in the Service Fulfillment application.

2.1.6 What's new in NSD/NRC Release 18.3

The following table lists the features added in Release 18.3 and described in the NSD and NRC customer documentation.

Table 2-4 NSD/NRC 18.3 features

Feature	Description and reference
NRC-X features	
NSPF-120646 — Multi-layer correlation of protection/resiliency	Within the Cross Domain Coordinator application, when IP-Optical Correlation is chosen from the drop down menu, a new column displays the optical protection type. See the *Cross Domain Coordinator* application.
NSPF-122916 — Link layer interconnect with diversity constraints	Within the Cross Domain Coordinator application, a user can request link layer interconnects, which provide connectivity between router ports over an optical transport network, as well as the corresponding links between router ports and the client ports of the optical switches. See the *Cross Domain Coordinator* application.
NSPF-126030 — NRC-X GEO/DR implementation	The NRC-X module can be installed and deployed in 1+1 mode. See the *NSP Deployment and Installation Guide*
NSPF-129987 — IP/optical correlation improvements	Within the Cross Domain Coordinator application, a new navigation concept replaces the existing tabs See the *Cross Domain Coordinator* application.
NSPF-141045 — NRC-X application name change	The application enabled by the NRC-X module is renamed to Cross Domain Coordinator. See the *Cross Domain Coordinator* application.
NSD features	

Table 2-4 NSD/NRC 18.3 features (continued)

Feature	Description and reference
NSPF-134562 NSD: Enhancements to L3 VPN services across Wavence and 7x50 SR domains	This feature introduces new functionality to support the provisioning of composite L3 VPN services across one Wavence SM domain and one 7x50 SR domain. 8.8.4 “Composite L3 VPN service across multiple Wavence SM domains and one 7x50 SR domain” (p. 78)

3 Common concepts

3.1 Overview

3.1.1 Introduction

This chapter describes concepts that are common to all NSD and NRC modules.

3.1.2 Applications

NOTICE

View of network can be affected

The NSD and NRC applications view of the network can be affected whenever activities are drawing heavily on CPU and memory usage.

This can happen when a large number of services are being created, modified, or deleted via the NSD and NRC modules' REST APIs.

The NSD and NRC modules provide functionality using browser-based applications. Each of these applications use the standard NSD and NRC REST security mechanisms for authentication and authorization, so every request sent to the server contains the provided session key. All applications are HTML5–based, and supported on the latest version of Google Chrome. Use the following URL to access the NSP launchpad, from which you can launch all supported applications:

https://*server*

Where *server* is the hostname or IP address of your installed NSD and NRC server.

Localized language support

All NSD and NRC applications support 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 NSD and NRC modules support localized language settings using predefined strings, and do not translate data to different languages.

3.1.3 REST APIs

The NSD and NRC modules also provide functionality using northbound RESTful APIs that expose a simplified view of the network. This view is constructed from the internal model, which is stored in the Topology Database. The APIs support queries, service creation requests, and many additional functions.

To view and interact with the APIs online, go to one of the following URLs:

- https://*server*:8543/sdn/doc

- https://*server*:8543/task-scheduler/doc
  where *server* is the hostname or IP address of your installed NSD and NRC server

More detailed information about using the NSP's REST API's can be found on the *NSP Developer Portal*.

3.1.4 Help system

NSD and NRC users have access to a help system that provides information about performing tasks within the NSP's various applications, including those that serve the NSD and NRC modules. This help system can be accessed from the NSP launchpad by clicking on the More... button and selecting Help from the contextual menu, or at the following URL:

https://*server*/Help/index.html

Where *server* is the hostname or IP address of your installed NSD and NRC server.

3.1.5 Additional components

The NSD and NRC modules rely on the following additional components to provide end-to-end functionality:

- *Topology Database* — The Topology Database contains a representation of the network in the form of a highly abstract, multi-layer graph. The graph is stored in a Neo4j database.
- *Network Mediation* — The Network Mediation component is responsible for populating the Topology Database with the network information and for deployment of network configuration. It is comprised of the generic plugin framework, as well as the mediation plugins that operate inside these. Plugins may interact with the network through Element Manager Systems (EMS) such as the NFM-P, and/or standard communication protocols such as PCEP, BGP-LS, or OpenFlow. The NSD and NRC modules support the deployment of network tunnels, services, and, potentially, tunnels.
- *Service Connection Manager* — The Service Connection Manager is responsible for finding appropriate tunnels for services.
- *Algorithm Framework* — The Algorithm Framework is the component that provides a run time environment for the invocation and execution of both routing and optimization algorithms.
- *Network Deployment* — The NSD and NRC modules support the deployment of network tunnels, services, and, potentially, tunnels. This means that some plugins and mediation framework may support the “push to the network” function that involves the mapping and conversion of the Topology Database entities to the network objects.
- *Security* — Security is the component that handles sign-in, encryption, logging of operator actions, and network events.
- *Relational Database* — A PostgreSQL database that contains all non-topological information requiring persistence. This includes policies, templates, etc.
- *Global Cache* — The Global Cache enables the NSD and NRC modules to track resources being used by the network, including the resources of services that originate from the NFM-P. In order for the NSD to discover such services, they must have their “NSD-managed” flag enabled within the NFM-P. Once this is done, the usage of VLAN IDs, L3 VPN Route Distinguishers (RD), and L3 VPN Route Targets (RT) can be tracked across NFM-P/NSD managed networks. When the NSD requests one of these resources, the Global Cache verifies their availability before

assignment. Only freed resources are considered available for usage. All services created using the NSD will be validated for resource usage, and therefore will not infringe upon the resources of existing services.

3.1.6 Security

SSL provides encryption on the following interfaces:

- The northbound REST interface that accepts requests from the GUI client and OSS systems
- The internal communication channels from the SDN application to the Policy Server
- The southbound interface to the NFM-P (only if NFM-P has SSL enabled)

SSL is enabled by default on all northbound and internal interfaces and no additional configuration is required, as the installer will automatically generate keystores to be used on those interfaces. Keystores generated automatically at installation contain a generated, self-signed certificate shared by all NSD and NRC instances. Custom keystores can also be pre-generated by the user and provided to the installer. These can contain either a self-signed certificate, or a security certificate signed by a certificate authority (CA).

Note: If a pre-generated keystore containing a self-signed certificate is used, the user will only have to manually accept the certificate when they first launch the web GUI and connect to the server. If a pre-generated keystore is *not* provided to the installer, then the certificate must be manually accepted the first time that each server becomes master of the cluster.

For information about retroactively enabling SSL, or generating keystores, see the *NSP Deployment and Installation Guide*.

4 Tenants and roles

4.1 Introduction

4.1.1 Tenants

A tenant represents a group of users with the same network management role. In the Service Fulfillment application, the system allows you to perform the service configuration tasks that are part of the privileges associated with your tenant type. You can create, delete, update and assign tenants through the NSD and NRC REST API. For details, see the *NSP Developer portal*

4.1.2 Roles

Each tenant that uses the NSD and NRC system is assigned a role. The tenant role specifies the type of access that users have to their tenant resources. A user can perform only the operations that are authorized by their role.

A user can be assigned multiple roles, but assumes the assigned role with the highest priority. The roles are prioritized as follows:

Role	Permissions	Priority
Admin	Modifies and manipulates any object within the NSD and NRC modules	1
Operator	Performs read and write operations on assigned network resources	2
User	Performs read-only operations on assigned network resources	3

4.1.3 Tenants and Customer IDs

The NSD and NRC system enforces the association of a tenant to a Customer ID. You must assign a Customer ID to a tenant during the tenant creation. If you do not specify a Customer ID when creating a tenant, then the NSD assigns the default Customer ID value of 1 to the tenant. A tenant can be associated with one Customer ID, and a Customer ID can be assigned to one tenant. However this rule does not apply to the default Customer ID value (1), which can be associated with multiple tenants.

Note: If the default Customer ID does not exist, then the creation of an NSP tenant without providing a Customer ID or with a Customer ID set to 1results in a tenant with the Customer ID 1. The NSD rejects any other Customer ID value provided to the NSP tenant creation operation.

The NSD can use only Customer IDs that are already defined in the NFM-P. The NSD checks if the Customer ID used to create a tenant exists in the NFM-P, and rejects the operation if the Customer ID is not already defined in the NFM-P.

A Customer ID can be deleted only from the NFM-P, and only if no services are associated with that

Customer ID. The NFM-P sends a notification when a Customer ID is deleted, and the NSD assigns the default Customer ID value 1 to all tenants that were associated with a deleted Customer ID.

You can modify the Customer ID of a tenant in the NSD, but only if no services are associated with the tenants. If you remove the defined Customer ID from a tenant, the NSD assigns the default Customer ID value to the tenant.

Limitation

If you change the tenant of an existing service, the existing customer association does not change.

4.1.4 Tenant association with service templates

The NSD supports the assignment of a template to one or more tenants, and provides the option to assign a template to all tenants.

Supported templates

The NSD supports the template assignment to tenants for the C-Line, E-Line, E-LAN and L3 VPN service templates, regardless of the underlying policies assigned to that service template. A service template can be assigned to multiple tenants, and a tenant can be associated with multiple services templates. During the service creation process, a tenant can select and apply only the service templates that are assigned to the tenant.

The NSD does not allow tenants access to the underlying policies assigned to a service template(such as the Tunnel Selection and Endpoint QoS policies).

Template assignment to all tenants

When you select All Tenants template assignment option, the NSD assigns the template to all current and future tenants. Select the All Tenants check box only if you want the template to be available to all the current tenants, as well as to all the tenants that will be created on your system.

Note: If you want to assign the template only to all current tenants, you must select the tenants individually and not use the All Tenants check box.

Admin tenant

The All Tenants option assigns the template to all tenants, both current and future, and also to the admin user. However, the admin user must be able to create a service template only for the admin user, and not for any other tenants. Therefore, you must create a specific tenant for the admin user, and then to select that tenant when creating the service template. If the tenant created for the admin user creates a service template and does not assign it to any tenants, then the service template is assigned by default only to the admin user tenant.

4.1.5 Tenant management REST APIs

You can configure tenants only through the NSD and NRC modules' REST API. To assist with the management of tenants and associate them with the appropriate customers and services, the following API calls are available:

- Query all customers

- Find a customer by customer-id
- Update tenant
- Find all ports in the network that are assigned to a tenant

Note: When the REST API call is executed to find all ports in the network that are assigned to a tenant, LAGs will not appear in the response unless the tenant is assigned to the physical layer of the LAG. This is because an access LAG has three layers - Physical, IP IGP, and IP service - whereas a normal access port simply has the IP Service layer.

4.2 To view services associated with a tenant

4.2.1 Steps

1

In the Service Fulfillment application, click on the Inventory tab.

2

Click on Tenants. The system displays the list of configured tenants on your system.

3

Click on a tenant entry in the list. The Info panel displays the tenant details.

4

Click VIEW SERVICES. The system displays the list of services associated with the tenant.

5

Click on a service to view service details.

END OF STEPS

4.3 To manage tenants associated with a port

4.3.1 Steps

1

In the Service Fulfillment application, click on the Inventory tab.

2

Click on Ports. The system displays the list of configured ports.

3

Click on a port entry in the list. The Info panel displays the ports details, including the number of tenants associated with the port.

4

Click on the Manage Tenants button for the port. The Manage Tenants form opens.

5

Manage the tenants associated with the port. You can add or remove tenants for the port, as required.

If you do not need to manage the tenants associated with the port, click CANCEL.

6

Click SAVE.

END OF STEPS

Part II: NSD and NRC modules

Overview

Purpose

This volume serves as an introduction to the NSD and NRC modules of the NSP, which form a carrier software-defined networking (SDN) platform that unifies service automation with network optimization, allowing network operators to deliver on-demand network services cost-effectively and with scalability. Using the NSD and NRC modules, an operator can define, provision, and activate network services across networks that span multiple layers (Layer 0 to Layer 3), services, and infrastructures (physical and virtual), as well as equipment from multiple vendors. The NSD and NRC modules are scalable, and based on standard protocols with multiple APIs.

The NSD and NRC modules act as bridge between the IT and network worlds. Upstream, they provide standard APIs, object models, and abstractions to IT OSS applications. Downstream, they manage network complexity by translating simple service requests into commands that program physical and virtual network elements. This is done automatically, across both IP and optical boundaries, and across multiple network vendors.

Contents

5 NRC-P

5.1 Overview

5.1.1 Introduction

The Network Resource Controller – Packet (NRC-P) module centralizes path computation and network optimization in order to leverage a whole network view and make the best possible decision for each request. For IP networks, this is done with a packet PCE. Centralized path computation elements are opened up to application and policy control, and to specialized algorithms. For instance, path computation can be enhanced to take link congestion into account. You can also make better use of your network assets and keep SLAs high by using KPIs and metrics to trigger optimization policies. You can also do all this in a multi-tenant way where each tenant – or in essence, each business unit - has their own abstracted view of the network and their own policies for maintaining service quality and assurance.

5.2 NRC-P (flow collector)

5.2.1 Introduction

The Network Resource Controller – Packet (NRC-P) can be deployed in a flow collector configuration that implements SDN-based, traffic-steering-related protocols, and applications. On the southbound side, the NRC-P uses flow-based protocols such as OpenFlow and BGP FlowSpec, and routing protocols such as BGP for route injection.

The NRC-P supports OpenFlow protocol Specification version 1.3.1 on both the controller and on Nokia 7750, 7450, and 7950 routers. The NRC-P is able to discover previously-configured OpenFlow switches on these routers, and can be used to add OpenFlow rules to their flow tables, or to delete flows from these tables altogether. Nokia VSR-NRC OpenFlow Experimenters are supported, and can redirect traffic to alternate next hops using plugins.

The NRC-P supports the following applications when used in a flow collector role:

- Autonomous System Optimizer
- Ingress Peer Optimizer
- Latency Steering Optimizer
- Traffic Steering Controller

For information about accessing flow collector NRC-P functionality through these applications, see 3.1.4 “Help system” (p. 22).

For information about accessing flow collector NRC-P functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

5.2.2 NRC-P flow collector deployment

The following figure shows a typical NRC-P flow collector deployment on virtual machines:

Figure 5-1 NRC-P flow collector deployment on virtual machines

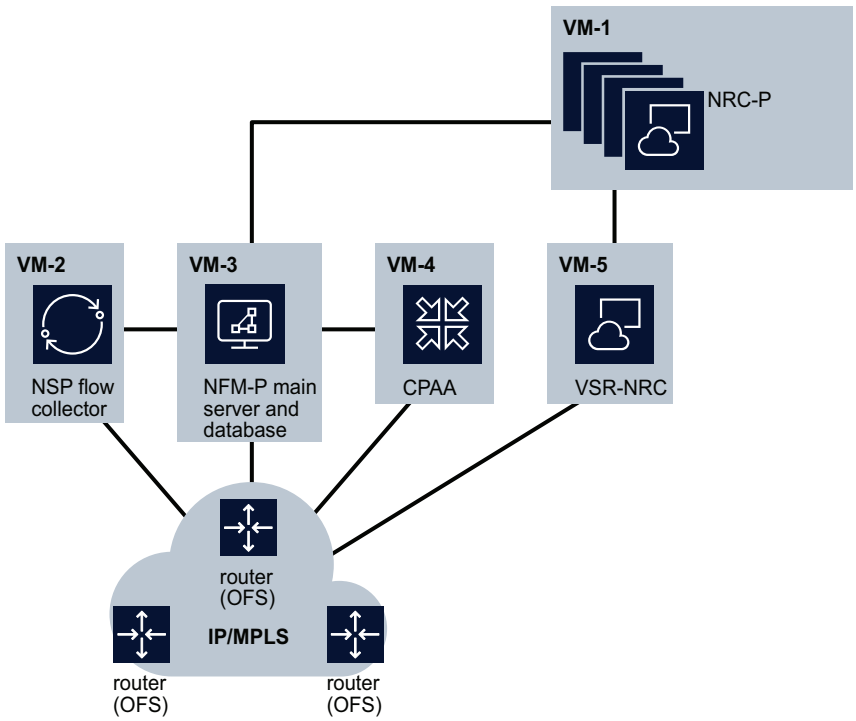

NFM-P

Routers monitored by the NRC-P are discovered from NFM-P network topology. In order for the NRC-P to receive information about the ports of these monitored routers, a TCA policy must be configured on the NFM-P. Each monitored port must be added to this policy. TCA Rules must also be created, with thresholds that represent the NSP's utilization bands: 20%, 40%, 60%, and 80%. These threshold values should include both rising and falling thresholds. An initial threshold can be set at 1% to allow for port utilization to be observed on low usage ports.

In order for the NRC-P to receive TCA notifications and real-time statistics, the following text must be added to the *opt/nsp/nfmp/server/nms/config/nms-server.xml* file:

```
<registry
enabled="true"
zkConnectionString="zookeeper_address"
/>
```

Where *zookeeper_address* is the IP address of the machine where the NSP's instance of Zookeeper is installed. In the case of redundant NSP deployments, two IP addresses separated by a semicolon must be provided.

Note: An NFM-P server restart is required in order for this configuration to take effect.

See the *NSP NFM-P User Guide* for more information.

See the *NSP Deployment and Installation Guide* for more information about TCAs.

NSP flow controller

The NSP flow controller aggregates and relays statistics to the NFM-P. In order for the NRC-P to receive statistics from its monitored routers, the NSP flow controller must be configured to communicate with the NFM-P, and the ports of each router monitored by the NRC-P must have Cflowd collection enabled.

Once the NSP flow controller has been installed, the following configuration commands must be performed from the server CLI-based samconfig utility:

```
samconfig -m flow
<flow> configure category sys
<flow configure> back
<flow> apply
<flow> exit_all
```

Filters must also be configured to show the monitored routers.

See the *NSP NFM-P Installation and Upgrade Guide* for more information.

vCPAA

The vCPAA is used to retrieve BGP prefixes for Autonomous Systems (AS) monitored by the NRC-P. In order for the NRC-P to monitor these ASs, the vCPAA must be integrated with the NFM-P that is relaying network information to the NRC-P. This vCPAA must also be configured with a BGP administrative domain.

In order for the NRC-P to monitor the ASs discovered by the vCPAA, the BGP section of the */opt/nsp/configure/config/nrcf.conf* file must be populated as follows:

```
bgp
{
        # BGP Autonomous system number of CPAA router
        cpaa_autonomous_system_number = CPAA_AS_number

        # BGP prefix filter id used for fetching prefixes.
        prefix_filter_id = 65535

        # BGP prefix fetch timeout (milli seconds).
        prefix_fetch_timeout = 60000

        # BGP AS subnet info refresh timer(hours)
        as_subnet_refresh_timer = 24
}
```

where *CPAA_AS_number* is an integer representing the vCPAA autonomous system number

Note: The retrieval of BGP AS subnets is based on the local AS of the BGP route, as seen by the vCPAA. When retrieving AS subnets, the NSP will modify the specified BGP prefix filter list for the requested local AS. The BGP ASs and the vCPAA ASs must match.
See the *NSP NFM-P Control Plane Assurance Manager User Guide* for more information.

VSR-NRC

In an NRC-P flow collector deployment, the VSR-NRC serves as an OpenFlow Controller (OFC). The OFC is used to push flows to the OpenFlow Switches (OFSes) under its control, as directed by the NRC-P.

Note: In order for the NRC-P to communicate with the OFC, the *openflow* parameter in the */opt/nsp/configure/config/sros-vms.conf* file must be set to true.

Note: The OFC CLI tree is visible on hardware, but cannot be used for NRC-P OpenFlow functions.

For more information about configuring the VSR-NRC for use with NRC-P, see the *NSP Deployment and Installation Guide*.

5.2.3 OpenFlow controller rules

The following rules are supported for the NRC-P OpenFlow Controller:

- **Forward action** — The OFC can program forward action when a specific flow is to be forwarded using regular router forwarding.
- **Redirect to GRT instance or VRF instance** — A router supports redirection of IPv4 or IPv6 traffic arriving on an L3 interface to a different routing instance (GRT or VRF).
- **Redirect to SDP** — For traffic arriving on a VPLS interface, routers support PBF to steer traffic over a VPLS SDP within the same service. The OFC leverages this functionality and programs the PBF steering action for H-OFS instances with switched-defined-cookie enabled.
- **Redirect to SAP** — For traffic arriving on a VPLS interface, routers support PBF to steer traffic over another VPLS SAP in the same service. The OFC leverages this functionality and programs the PBF steering action for H-OFS instances with switched-defined-cookie enabled.

For encoding information pertaining to these rules, see the *Router Configuration Guide R15.0.R5*.

5.2.4 Express peering

Modern web-based are becoming increasingly latency- and bandwidth-sensitive, making internet traffic highly unpredictable. Various events can generate sudden traffic spikes in the network which can impact the end user's QoE. Nokia Express peering (part of the NRC-P module) is an automated approach to internet peering that addresses rapidly changing traffic patterns more effectively by integrating real-time traffic visibility with instant SDN software control of network resources and application traffic flows — which enables automatic traffic engineering at the peering routers and network interconnections.

The goal of Express peering is to simplify operations, improve service delivery with the intent to reduce costly manual misconfigurations as much as possible. Express peering focuses on better traffic congestion detection and resolution and/or decreasing end-to-end latency to improve performance and QoE — particularly for time-sensitive applications.

Express peering functionality is accessible within the following NRC-P applications:

Note: Detailed information - including procedures - are provided for each application using the help system that is accessed from the NSP launchpad. See 3.1.4 “Help system” (p. 22) for more information.

Ingress peer optimizer

The Ingress Peer Optimizer (IPO) application provides policy abstraction, and network control automation (or semi-automation with real-time traffic monitoring), and correlation for controlling inbound traffic. The NSP integrates both routing programmability and control, routing protocol-based traffic visibility, automatic congestion detection and reporting, and finally, peering topology map with traffic visibility. All of these elements, integrated into a single NSP application, provide carriers and enterprises with a powerful tool to proactively and operate and understand inbound traffic that is destined or transiting their network in real-time. This allows them to apply appropriate policies, either manually or through automatic computational algorithm, at the right time and at the right location.

The IPO allows the operator to perform the following tasks:

- Create a network optimization slice where optimization will be contained within a slice, or across slices
- Set the application to operate in visibility-only mode, or full optimization mode
- Dynamically monitor the traffic on EBGP peering links and routers for high traffic utilization
- Automatically detect congestion situations
- Automatically report BGP community real-time bandwidth information
- The IPO shifts traffic from one peering point to another available link using BGP policies (depending on optimization mode - manual or automatic)
- Optionally, select specific routes (which correspond to customer traffic flows) from the auto-discovered BGP table to be steered to specific auto-discovered LSPs

The IPO integrates both peering and internal BGP topology maps with traffic link utilization reporting and BGP community reporting obtained from Deepfield. In automatic optimization mode, the IPO automatically constructs the appropriate BGP policies to be injected when congestion is detected on specific links.

The IPO also supports the ability to revert to default routing states for a given peering relationship by shifting traffic back to the initial state when congestion is cleared, or when it is safe to return to the preferred route.

Egress peer engineering with Latency steering optimizer

The Latency Steering Optimizer application allows for traffic from an entry point router to a specific destination to be routed onto a path with the lowest latency. Latency is considered on the whole path based on both the internal latency over the LSPs and the latency from the border to the destination.

NE points are added for routers which are candidates to be used for the path to the destination. Destination Sites are also added to a representation of the destination for steered traffic. A Destination Site contains Destination Points, which are comprised of subnets. These subnets define

the destination bound traffic and serve as a representative probe address for latency measurement. Connections between NE points are discovered automatically based on existing LSPs. Connections to destination points must be created manually. Latency on connections can be manually injected or automatically updated by consuming results of OAM tests within NFM-P.

When the NSP system is restarted, the Latency Steering Optimizer has a five minute hold off timer. This allows sufficient time for all OFSes in the network to be fully discovered in the Latency Steering Optimizer before performing potentially disruptive optimization. For example, the view of the network may not initially include an undiscovered OFS, which would be omitted from the latency calculation. This hold off timer is only enforced in the condition that a previously reachable OFS is not discovered within the given period. If all OFSes are discovered before this time is reached, network optimization will proceed as normal.

When the following events occur, the Latency Steering Optimizer application will not optimize for a period of sixty seconds:

1. A new LSP is discovered between two NE Points
2. The state of an existing LSP changes from down to up
3. A new, internal NE Point is created
4. The latency value of a connection changes
5. A previous steering action failed to deploy to the OFS

Note: Events occurring during this sixty second period are queued and will be processed at the end of the period.

Within the Latency Steering Application, steering and flow deployment begins as soon as latency is detected on paths between source NE points and destination subnets. This means that all internal and external connections along these paths will have latency values applied. If the paths include multiple LSPs, flow deployment will occur on all NE points along the paths from which the LSPs are sourced.

Autonomous system optimizer

The Autonomous System Optimizer application is used to steer traffic on monitored routers, on a per-destination-AS-basis, to alternate next hops. Steering per destination AS implies that steering will be performed for all prefixes associated with a given destination AS. The NRC-P will automatically correlate the destination AS number to the set of prefixes associated with it. Steering is accomplished using the NRC-P OpenFlow controller, by automatically adding an OpenFlow flow rule per destination subnet. This allows the user to offload high traffic usage from the uplinks onto alternate paths on a per-AS-basis.

Users can monitor traffic distribution on a set of uplinks so that link congestion and/or high bandwidth utilization can be identified per AS. The traffic monitoring is accomplished by collecting flow statistics, per AS, on Nokia 7750, 7450, and 7950 routers. These flow statistics are then communicated to the collector with IPFIX record encoding.

The application allows users to identify the set of top bandwidth consumption per destination AS, while the set of destination subnets associated with a given AS are automatically identified. Threshold Crossing Alarms (TCAs) on the monitored links can be tracked and a user can plot both real-time and historical port utilization.

5.3 NRC-P (PCE)

5.3.1 Introduction

The Network Resource Controller - Packet (NRC-P) leverages centralized, intelligent network control capabilities so that operators can rapidly adapt to changing demand and traffic patterns and run their networks more efficiently. The NRC-P accepts path connection requests from the NSD, from OSS and orchestration systems, and from physical/virtual network elements. The NRC-P calculates optimal paths through the network for a given set of business and technical constraints by leveraging centralized views of all available assets/topologies and their current state.

The NRC-P module is based on a Path Computation Element (PCE) architecture that integrates standard protocols such as PCEP to open up path computation to external control. This allows PCCs to be enhanced with various path optimization algorithms that ensure optimal path placement across the network. The NRC-P is stateful in nature and will maintain an up-to-date Traffic Engineering Database (TED), as well as the current RSVP-based label switched paths (LSP) and the segment routing path (SRP) state. It tracks RSVP BW and manages BW for the Segment-Routed TE paths as a unified state.

The NRC-P supports the IP/MPLS Optimization application.

For information about accessing PCE NRC-P functionality through these applications, see 3.1.4 “Help system” (p. 22).

For information about accessing PCE NRC-P functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

5.3.2 NRC-P PCE deployment

The following figure shows a typical NRC-P PCE deployment on virtual machines.

Figure 5-2 NRC-P PCE deployment on virtual machines

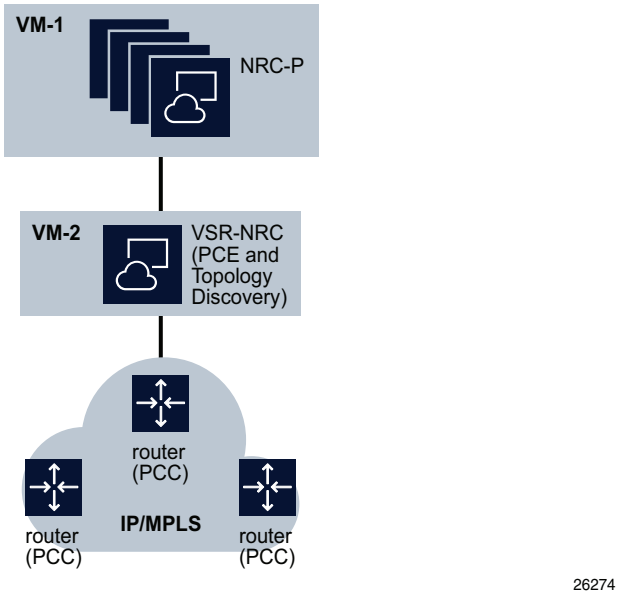

VSR-NRC

The NRC-P is the PCE, and contains the logic to calculate paths. The VSR-NRC is a component of the NRC-P, but does not calculate any paths. The VSR-NRC terminates PCEP connections and conveys path request messages from PCCs to the NRC-P. The NRC-P computes the requested path and responds to the VSR-NRC, which conveys the response to the PCCs. The communication between VSR-NRC and PCCs is accomplished using the PCEP protocol.

In order for the NRC-P to compute paths, it must discover the IGP topology. Topology discovery can be performed by peering the VSR-NRC directly in the IGP or using BGP-LS. If using IGP, the VSR-NRC must have full visibility of the topology. For multi-area topologies, this means that the VSR-NRC must be connected to every area, or to the ABRs/(L1/L2s) via IGP (OSPF or ISIS) adjacencies. If using BGP-LS, the VSR-NRC must be peered with a BGP speaker, ABRs/(L1/L2s) that are BGP speakers, or a Router Reflector that is peered to a BGP speaker in each IGP area. In order for BGP-LS discovery to be successful, each BGP speaker must support BGP-LS.

Note: Only the VSR-NRC supports topology discovery for the NRC-P. Do not use any other devices, such as the vCPAA, because they are not supported.

For more information about configuring the VSR-NRC for use with NRC-P, see the *NSP Deployment and Installation Guide*.

5.3.3 PCC-initiated LSPs

The NSD and NRC modules support the creation of PCC-initiated Segment-Routed TE LSPs, as well as PCC-initiated RSVP LSPs. The NSD sends a service creation request, through NFM-P, to the PCC router(s) endpoints which are part of the service. An empty LSP path is created on each of

the PCC router(s). The PCC router(s) then send an LSP path request to the NRC-P (PCE). The NRC-P (PCE) computes the LSP path and sends the path to the PCC. The PCC then installs the computed path and informs the NRC-P (PCE) that the path is ready. This is reported to NSD, where the path is attached to the service.

5.3.4 PCE-initiated LSPs

The NRC-P supports the creation of PCE-initiated Segment-Routed TE LSPs. Operators can specify the LSP parameters and PCC address using an LSP creation form within the IP/MPLS Optimization application, or by using the NSD and NRC modules' API. Operators can also select a path profile to associate to the LSP path. There is also an NBI. PCE-initiated LSPs are deployed through PCEP.

In order to create a PCE-initiated LSP, the following commands must be executed on the node:

1. **config>router>pcep>pcc**

   max-srte-pce-init-lsps <max-number>

   Where *max-number* is a number between 0 and 8191, which can only be modified when config>router>pcep>pcc is shutdown.

2. **config>router>mpls# lsp-template template-name pce-init-p2p-srte {default | template-id } default-path {pathname}**

   path "P"

   no shutdown

   exit

   lsp-template "test" pce-init-p2p-srte template-id default

   default-path "P"

   cspf

   pce-report enable

   no shutdown

   exit

3. **config>router>mpls>**

   pce-initiated-lsp sr-te

   no shutdown

5.3.5 RSVP LSPs

Any PCC node intending to request a path computation from the NRC-P must first set the PCE computation option in the LSP definition. The PCC then assigns a unique PLSP-ID to the LSP. This uniquely identifies the LSP within a PCEP session and is maintained for the lifetime of the LSP. The PLSP-ID is also associated to the tunnel and path ID.

Once the PLSP-ID is assigned, the PCC sends a PCReq message to the NRC-P PCE, requesting a path for the LSP. This request includes the LSP parameters in the METRIC object, the LSPA object, and the Bandwidth object. It also includes the LSP object with the selected PLSP-ID. The NRC-P is now able to compute a new path, to check the bandwidth, and to return the path in a PCRep

message with the computed Explicit Router Object (ERO) in the ERO object. It also includes the LSP object with the unique PLSP-ID, the METRIC object with the computed metric value (if any), and the Bandwidth object.

The NRC-P does not keep track of the LSP yet. At this point, it has simply returned the ERO. The PCC has yet to confirm that the path was signaled. If the path was locally signaled, and the local TE database (TEDB) was updated, the NRC-P receives the updates via BGP-LS and update its TEDB.

For stateful operation, which allows the NRC-P to track the LSP path and bandwidth (among other constraints), the PCE report option must be set in the LSP definition. When this option is set, the PCC sends both a PCRpt message to update the NRC-P with the state of UP, and the Record Route Object (RRO) object as confirmation. The RRO object now includes the LSP object with the unique PLSP-ID. With this, the NRC-P is able to display the LSP, as well as its hops and constraints. The RRO also contains information about the protection that is enabled on the signaled path. Therefore, the NRC-P is aware of the protection at the hops, but not aware of the detour/bypass tunnel details. If a local failure causes the LSP on the PCC to switch to a detour or bypass, a PCE report is sent to the NRC-P, and the NRC-P becomes aware that the LSP is using a detour or bypass.

Note: In the VSR-NRC, the PCE reporting option can either be set globally, or on a per LSP basis.

The PCC can also delegate control of the LSP to the NRC-P for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once the NRC-P is controlling the LSP, the operator can manually re-signal/re-optimize the LSP. Re-signalling routes the LSP using its original constraints, while re-optimizing routes the LSP using an optimization algorithm. The NRC-P also re-routes LSPs automatically on resource failures, or when calculating disjoint paths.

Note: When the PCC has delegated control of the LSP to the NRC-P, any change to the LSP definition (such as changes in constraints) requires the PCC to first revoke the delegation via the PCE report option, and then to issue a new request to the NRC-P.

Secondary path behavior

The PCC sends PCE requests for standby secondary paths. A new PLSP-ID is used for these paths over the PCEP session, and is associated to the LSP path ID and the LSP tunnel ID. When a secondary path is not in standby, the PCE request is not sent until the primary path is down, or in FRR. However, if the path is delegated to the NRC-P, this results in a PCE update from the NRC-P. The LSP may switch to the secondary path in the interim, but will switch back to the primary path as soon as possible.

The NRC-P maintains the active path in case both the primary and secondary paths are signaled, and also when the primary path is down. The NRC-P also maintains the shared explicit behavior when the primary and secondary paths share common link resources.

The NRC-P also indicates the active path between the primary and secondary pair.

FRR notification

Fast re-route (FRR) is signaled locally, with locally-created detour tunnels. These tunnels are not reported to the NRC-P, and therefore the NRC-P is not aware of the detours and bypass. However, the types of node and/or link protection are communicated to the NRC-P via the PCE report.

Note: All the RSVP-TE LSPs created by the NSD NRC have FRR enabled by default. The FRR method used is “facility”.

RSVP LSP bandwidth management

The NRC-P manages the LSP bandwidth consumption on the TE links for both stateless and stateful PCC configurations. In a stateless configuration, the NRC-P receives TE updates from the network as LSPs are signaled, thereby mimicking the TE DB bandwidth consumption on the nodes. This allows for accurate LSP path computation without maintaining state on the NRC-P. In a stateful case, wherein the reports are sent to the NRC-P from the PCC, the bandwidth is again communicated by the PCC to the NRC-P via the bandwidth object. Here, the NRC-P will reconcile the TE update with the specific LSP bandwidth update via the report. Therefore, the NRC-P maintains full LSP state along with the consumption on the TE links for these LSPs only.

It is possible that existing brownfield LSPs will not request paths from the NRC-P, and therefore, will have no state on the NRC-P. The NRC-P will not show these LSP reservations on the TE links. For a mixture of LSPs that are PCE-reported and non-PCE-reported, the NRC-P will track and show the actual TE consumption on a TE link in addition to the LSP reservation for PCE-reported LSPs.

Although bandwidth is not tracked until reported, bandwidth is reserved for one (1) minute when a request is made. Therefore, if multiple requests are made in quick succession, subsequent requests will be impacted, even though reports have not yet been received.

If telemetry is enabled on the LSP, the greater of the requested bandwidth or the telemetry-measured bandwidth will be used. Telemetry tracking can be enabled under “path profile” for the specific LSP.

5.3.6 Segment-routed TE LSPs

Any PCC node intending to request a path computation from the NRC-P must first set the PCE computation option in the LSP definition. The PCC then assigns a unique Path LSP-ID (PLSP-ID) to the LSP. This uniquely identifies the LSP within a PCEP session and is maintained for the lifetime of the LSP. The PLSP-ID is also associated to the tunnel and path ID.

Once the PLSP-ID is assigned, the PCC sends a PCReq message to the NRC-P PCE, requesting a path for the LSP. This request includes the LSP parameters in the SRP object, the METRIC object, the LSPA object, and the Bandwidth object. It also includes the LSP object with the selected PLSP-ID. The NRC-P will reserve bandwidth for the path to be returned, but will not keep track of the operational status or other requirements for the LSP yet. At this point, bandwidth is consumed and an ERO is returned. The PCC has yet to confirm that the path was signaled. If the path was locally signaled, and the local TEDB has been updated, the NRC-P will receive a REPORT from the PCC and the updates via BGP-LS and update its TEDB. If the PCC fails to send a report, after a period of time the bandwidth reserved will be released from the NRC-P. The path computed by the NRC-P is specified explicitly with the next hop interfaces and the adjacency SIDs encoded in the SR ERO sub-object.

When the PCE report option is set in the LSP definition, the PCC sends both a PCRpt message to update the NRC-P with the state of UP, and the RRO object as confirmation. The RRO object now includes the LSP object with the unique PLSP-ID. With this, the NRC-P is able to display the LSP, as well as its hops and constraints. The RRO also contains information about the protection that is enabled on the signaled path. Therefore, the NRC-P is aware of the protection at the hops, but not aware of the detour/bypass tunnel details. If a local failure causes the LSP on the PCC to switch to a detour or bypass, a PCE report is sent to the NRC-P, and the NRC-P becomes aware that the LSP is using a detour or bypass.

Note: In the VSR-NRC, the PCE reporting option can either be set globally, or on a per LSP basis.

The PCC can also delegate control of the LSP to the NRC-P for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once the NRC-P is controlling the LSP, the operator can manually re-signal/re-optimize the LSP. Re-signalling routes the LSP using its original constraints, while re-optimizing routes the LSP using an optimization algorithm. The NRC-P also re-routes LSPs automatically on resource failures, or when calculating disjoint paths.

Note: When the PCC has delegated control of the LSP to the NRC-P, any change to the LSP definition (such as changes in constraints), requires the PCC to first revoke the delegation via the PCE report option, and then issue a new request to the NRC-P.

Bandwidth management

A bandwidth value that is specified on an LSP has no significance on the PCC/router because the SR TE does not maintain any state on the intermediate or destination routers. Therefore, no bandwidth tracking is done in the local TE DB. The bandwidth has to be tracked by the NRC-P if the LSP is configured to report bandwidth. Bandwidth tracking on the NRC-P is done only after a valid PCE report message is generated by the PCC. The NRC-P tracks the bandwidth reservation for SR TE LSPs separate from RSVP TE LSPs.

Note: A loose hop SR LSP whose bandwidth is specified and computed locally will not be tracked by the NRC-P, even with the PCE report option enabled. The NRC-P only tracks SR TE LSP paths computed by the NRC-P itself.

Although bandwidth is not tracked until reported, bandwidth is reserved for one (1) minute when a request is made. Therefore, if multiple requests are made in quick succession, subsequent requests will be impacted, even though reports have not yet been received.

Failure detection

The head end router for an SR TE path, or an SR path, has no indication when a downstream link failure has impacted traffic for that SR TE or SR path. For a stateless and stateful application without PCE control, the SR TE tunnel on the head end router will remain up, as it receives no notification from the control plane either locally, or via NRC-P. For an LSP with delegated control to the NRC-P, the NRC-P will react to the topology change and issue a new ERO update to the PCC via PCE update.

5.3.7 Anycast and loopback for LSPs

The NRC-P supports path computation requests that include anycast or loopback addresses as destinations.

When inter-domain with multiple instances on routers are supported, the NRC-P can specify a loose hop ERO with anycast loopbacks as intermediate hops. This allows for the generation of an inter-domain ERO between domains when domain boundary routers have anycast loopbacks configured.

The ERO generation is controlled via a path profile with a new ERO specification option field. If the specification is *anycast preferred*, then the inter-domain computed path will consist of border routers which have the anycast configuration as loopback addresses with identical anycast SIDs. If the specification is *loose hop preferred*, then the inter-domain computed path will consist of the best loose hop border routers with node SIDs.

Note: Anycast SIDs are node SIDs that are associated to the loopback addresses instead of the system address. In SROS, there is no specific designation for anycast SIDs.

Note: The ERO specification default is the complete path with Adjacency SIDs, however, in the inter-domain cases, the number of Adjacency SIDs will most likely exceed the MSD.

Note: When the *anycast preferred* ERO specification is used and the inter-domain border routers do not have anycast SIDs, the best loose hop node SID among the inter-domain border routers will be selected.

5.3.8 IRO object

The NRC-P supports the IRO object specification within a PCC request. The NRC-P computes a CSPF path from the source to the IRO object, and another CSFP path from the IRO object to the destination. If the second CSPF path visits any of the nodes in first CSPF path, the path computation fails.

When used with a path profile that contains the bidirectional disjoint specification, a forward LSP and its matching reverse LSP must share the same IRO configuration. This means that the list of addresses in the IRO path must be the same, but their order reversed. This is because the disjoint algorithm is natively bidirectional strict. If the reverse LSP contained IROs that did not exist in the forward path, no path would be found, because it would no longer be bidirectional strict.

5.3.9 Algorithms

The NRC-P uses path computation algorithms to identify optimal paths within the network.

STAR algorithm

The NRC-P provides a load-balancing and optimal-path-placement algorithm, known as the STAR algorithm. This algorithm uses an internal metric, calculated from the current value of the TE bandwidth reservation, to route the CSPF paths. Every path that is allocated on a TE link changes the internal metric for both the link and the overall path. Initially, all links have the same star weight, or metric, so the first path requests for CSPF traversal will choose the shortest path that satisfies all constraints. If there are multiple paths that satisfy the user constraints, then a path will be chosen randomly. This behavior is the same for normal CSPF.

Subsequent requests will choose paths that possess the least star weight, thereby ignoring the path

that the normal CSPF algorithm would have chosen. The calculation of the star weight is based on a formula that uses the current link reservation. The user constraints are still satisfied. This balances the overall network utilization.

The STAR algorithm is invoked per LSP by associating that LSP to a path profile. The path profile template is defined in the NRC-P and requires setting the objective to use STAR WEIGHT. The path profile is specified with the LSP definition and is conveyed to the NRC-P via a PCE request message.

Disjoint optimal path computation algorithm

The NRC-P provides support for disjoint path computation between a source destination pair and between two pairs of sources and destinations. Applications can use this algorithm to provide no-impact redundancy for a service offering. The algorithm provides node/link and SRLG types of disjoint path computations. The algorithm can also re-optimize an existing path if a second path request asks to be disjoint from the existing path. The ability to treat a pair of paths as mutually disjoint requires associating a path profile ID to the path request. In addition, a path group ID specification is also essential to implicitly identify the path pair from other path pairs. The disjoint optimal path calculation algorithm can also compute paths that are bidirectionally symmetric, to ensure that forward and reverse traffic use the hops while being disjoint.

Note: The NRC-P can only compute bi-directionally symmetric forward or reverse paths. For an RSVP LSP with primary and secondary path specification, the profile is applied to both paths. For example, if there are two RSVP LSPs between the respective distinct sources and destinations, the primary path of LSP 1 will be mutually disjoint from the primary path of LSP2, and vice versa for secondary paths. The algorithm cannot be applied to ensure the primary and secondary paths between the same source and destination pair are mutually disjoint.

Global concurrent optimization algorithm

The NRC-P also supports optimizing the paths of existing LSPs by applying an optimization algorithm. This algorithm extracts the current resource availability on the current topology and re-routes the selected LSP paths such that the overall network consumption is minimized. The result is to utilize more network links, but also reduce the consumption on the links. LSPs must be delegated to the NRC-P and must be pre-selected. Profiles do not have to be associated to the paths in order to use this algorithm. The LSPs to be optimized are selected manually on from the NSD application.

Note: The LSPs that have a profile with the disjoint option enabled are excluded.

5.3.10 Multi-domain path computation

The NRC-P supports path computation across multiple IGP instances. These instances are discovered as admin domains with stitching points on the common ASBR routers. The path traversal algorithm uses a flat graph and computes the shortest path based on the required metric. Any optimization limiting domain traversals is not considered. Both Segment-Routed TE paths and RSVP TE paths are supported and deployed. Existing constraints such as the Max SID label depth apply.

5.3.11 Northbound interface for topology retrieval

NRC-P supports both an IETF-based NBI and a proprietary NBI model with extended attributes for topology retrieval. This is in addition to the existing IETF-based NBI. Using this NBI, northbound applications can obtain the IP and TE topology from the NSP, including additional information (such as area number/level instance and node/link/prefix SIDs). Northbound applications and controllers, such as the NRC-X, can also modify the following TE attributes: SRLG, TE metric, IGP metric, Latency, and admin group.

5.4 NRC-P Sim

5.4.1 Overview

The NRC-P Sim is a traffic engineering tool that can be used by network engineers to design a new network, or optimize and simulate failures in an existing network that is imported into the tool. A network topology can be imported into NRC-P Sim from an NRC-P system.

The NRC-P Sim is installed with its own nspOS on a separate platform from the NSD and NRC modules. The platform requirements for a NRC-P Sim deployment are the same as for a deployment of NSD and NRC modules, except where indicated.

The NRC-P Sim supports the IP/MPLS Simulation application.

For information about accessing NRC-P Sim functionality through these applications, see 3.1.4 “Help system” (p. 22).

For information about accessing NRC-P Sim functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

5.4.2 NRC-P Sim deployment

The following figure shows a typical NRC-P Sim deployment on virtual machines.

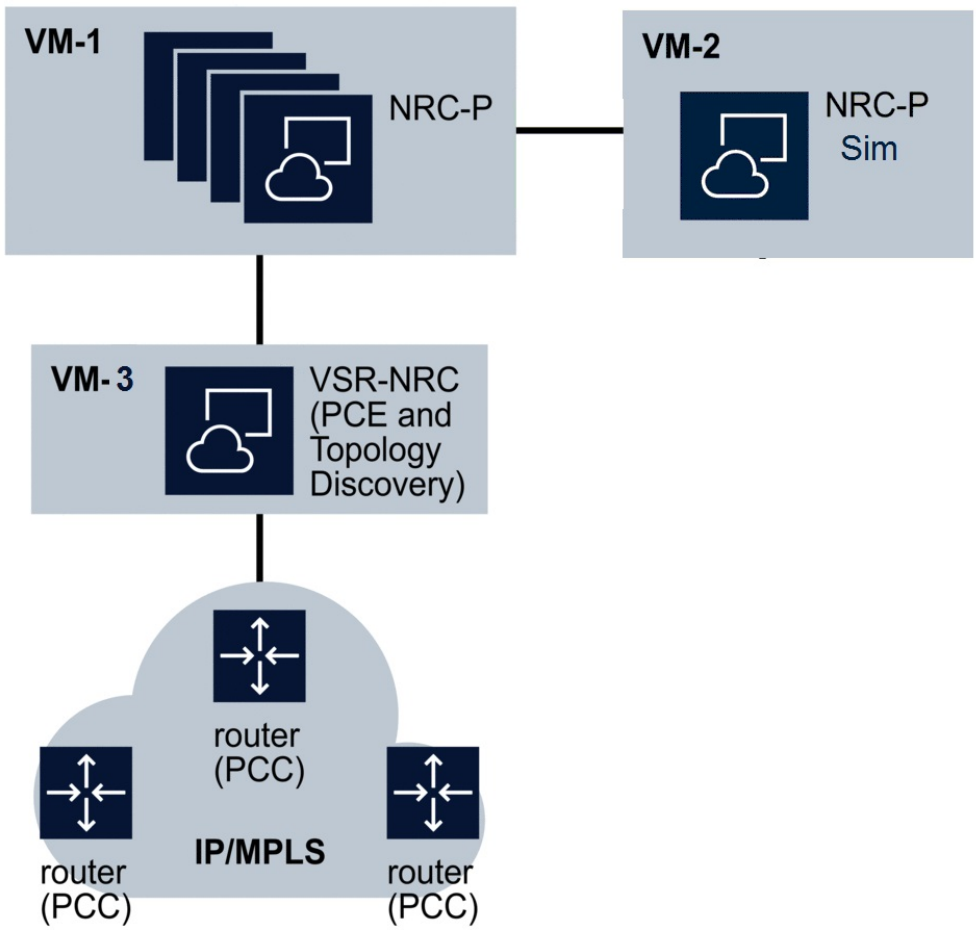

NRC-P

In this type of deployment, the NRC-P acts as the PCE, and contains the logic to calculate paths.

VSR-NRC

The VSR-NRC is a part of the NRC-P, but does not calculate any paths. The VSR-NRC terminates PCEP connections and conveys path request messages from PCCs to the NRC-P. The NRC-P computes the requested path and responds to the VSR-NRC, which conveys the response to the PCCs. The communication between VSR-NRC and PCCs is accomplished using the PCEP protocol.

In order for the VSR-NRC to discover the IGP topology, it must be peered with IGP routers. It can then discover IGP link-state topologies using either IGP or BGP-LS. If using IGP, the VSR-NRC must have full visibility of the topology. For multi-area topologies, this means that the VSR-NRC must be connected to every area, or to the ABRs/(L1/L2s) via IGP (OSPF or ISIS) adjacencies. If

using BGP-LS, the VSR-NRC must be peered with a BGP speaker, ABRs/(L1/L2s) that are BGP speakers, or a Router Reflector that is peered to a BGP speaker in each IGP area. In order for BGP-LS discovery to be successful, each BGP speaker must support BGP-LS.

Note: Only the VSR-NRC supports topology discovery for the NRC-P. Do not use any other devices, such as the vCPAA, because they are not supported.

For more information about configuring the VSR-NRC for use with NRC-P, see the *NSP Deployment and Installation Guide*.

5.4.3 IP/MPLS Simulation application

The IP/MPLS Simulation application provides the ability to simulate changes in the IP topology that was discovered by the IP/MPLS Optimization application.

This application is run in a separate VM and imports the IP topology and LSPs from the NRC-P IP/MPLS Optimization application. The specific simulation functions supported on this application are:

1. Modifying link attributes
2. Modifying the status links
3. Creating or deleting LSPs
4. Modifying the profile of imported LSPs
5. Optimizing LSPs via the GCO algorithm

For each change, the Simulate button is activated to determine the visual impact of that change. The general functionality supported in addition to the above functionality are:

1. Import IP/TE topology only
2. Import LSPs
3. Import profiles
4. Delete an imported topology

6 NRC-X

6.1 Overview

6.1.1 Cross-domain coordination

The Network Resource Controller - Cross Domain, or NRC-X, provides cross domain coordination between multiple layers, domains, and IP/optical integration functions. The NRC-X automatically discovers the cross-layer links between the IP routers and the optical switches using LLDP and LLDP snooping. In a brownfield deployment where a customer has pre-configured IP-optical links, the NRC-X will automatically discover the network, run its IP-optical correlation algorithms, and detect all misconfigured IP-optical links.

The NRC-X discovers the entire L0-L3 (IP, optical ODU, and optical OCH) topology. It processes information acquired from the other NSP modules, traverses the IP-optical layers and links, and computes the SRLG and the latency values end-to-end on the optical paths over which the IP interfaces ride. Once those values are computed, they are passed on to the IP layer for further processing. Doing this helps the IP layer to prevent SRLG risks during cross-layer end-to-end IP/MPLS computation. Passing the latency values helps the IP layer to establish latency-aware IP/MPLS LSPs.

The NRC-X is installed on a separate Virtual Machine (VM) and can run independently of other NSP modules. The NRC-X communicates with the NSD and NRC-P modules. REST interfaces are used to retrieve and update topology information.

The NRC-X supports the Cross Domain Coordinator application.

The NRC-P supports the IP/MPLS Optimization application.

For information about accessing NRC-X functionality through this application, see 3.1.4 “Help system” (p. 22).

For information about accessing NRC-X functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

6.1.2 NRC-X deployment

The following figure shows a typical NRC-X deployment on virtual machines.

Figure 6-1 NRC-X deployment on virtual machines

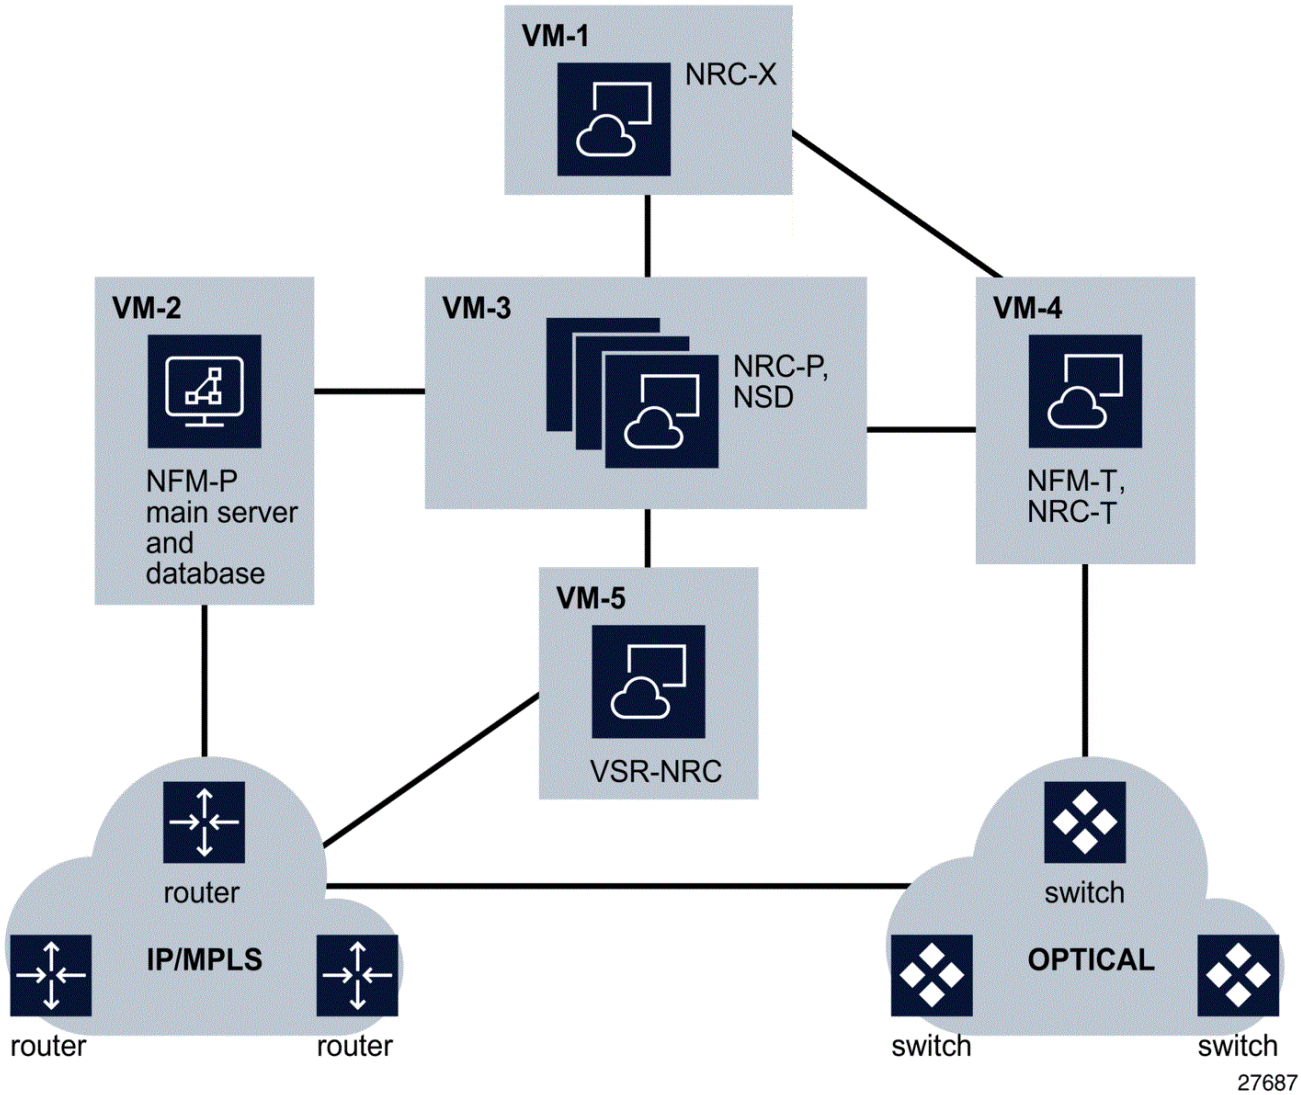

6.1.3 Inter-layer links discovery

The NRC-X automatically discovers inter-layer links, which are physical links between IP and optical endpoints. These links are important for the correlation of IP services and optical transport services. The Cross Domain Coordinator application uses deterministic methods and heuristics to discover links. The application supports deterministic discovery based on LLDP snooping on 1830 PSS nodes. In order for discovery to be successful, the optical transport service needs to be provisioned and the ational state needs to be *UP*.

Within the application, users can commit or delete auto-discovered links, as well as create links manually. Once committed, manually-created links are available to other NSP modules. Physical links can also be added or deleted in other applications, such as the Service Fulfillment application. In that case, the *Sync with Network* operation synchronizes the data between components. Links present in other applications will take precedence in the case of a conflict. It may also be necessary to refresh the views within the Cross Domain Coordinator application in order to capture all changes to the network. The NRC-X will detect manually misconfigured IP-optical links and allow users to resolve conflicts found during link discovery.

6.1.4 Optically disjoint IP routing

The NRC-X allows the NRC-P to establish IP/MPLS LSPs that are disjoint within the optical layer. In order to ensure optical diversity, optical Shared Risk Link Groups (SRLGs) are retrieved from transport controllers, correlated with the corresponding IP links, and exported into the NRC-P. In the case of LAGs, the union of SRLGs is calculated.

Optical topology is retrieved from the NFM-T, the prerequisite being that SRLGs are provisioned by the NFM-T into GMRE. The recommended deployment mode is to define SRLGs for all L0 TE links.

IP-optical Correlation and SRLG upload/withdrawl to the NRC-P is controlled via the Cross Domain Coordinator application. The NRC-X only exports SRGs without modification or detection of overlaps. The values sent to the NRC-P can also be revoked.

6.1.5 IP routing based on optical latency

The NRC-X supports the correlation of static latency parameters from optical links to the IP layer, which can used for latency-based IP routing. This allows for the circumvention of delay measurements within the IP layer.

The NRC-X retrieves optical latency information from NFM-T, correlates it with the corresponding IP links, and updates the IGP information within the NRC-P accordingly. In the case of LAGs, the maximum latency of all LAG members is considered.

7 NSD

7.1 NSD

7.1.1 NSD functionality

The Network Services Director, or NSD, is the network service fulfillment module of the NSP. It automates IP/MPLS and carrier Ethernet service provisioning by mapping abstract service definitions to detailed service templates using operator-defined policies. The NSD also provides provisioning for complex multi-technology services across multi-domain networks.

The NSD maintains abstracted service models that are based on YANG standards, and maps the models to device-specific models that are normalized for multi-vendor provisioning transparency.

The NSD provides network-aware management using a central service connection resource database to track tunnel bandwidth. As the NSD provisions a service, it performs an intelligent database search to choose the optimal path based on the required bandwidth, span, latency, cost, path diversity, and other constraints. Using the resource database and policies, the NSD directs service connection requests to tunnels or paths that have low utilization and thus averts link congestion.

An NSD operator can customize the binding of service connections to tunnels or paths using service-specific policies. If there is no service connection path that meets the specified requirements, the NSD can use a policy to request a new path from the NRC.

The NSD works with the NSP Assurance and Analytics functions for use cases such as IP network-aware provisioning automation with service validation, and bandwidth-on-demand for IP services.

7.1.2 NSD application support

The NSD supports the following applications:

- Service Fulfillment
- Policy Management
- Task Scheduler

For information about accessing NSD functionality through these applications, see 3.1.4 “Help system” (p. 22).

For information about accessing NSD functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

7.1.3 NSD deployment

The following figure shows a typical NSD deployment on virtual machines.

Figure 7-1 NSD deployment on virtual machines

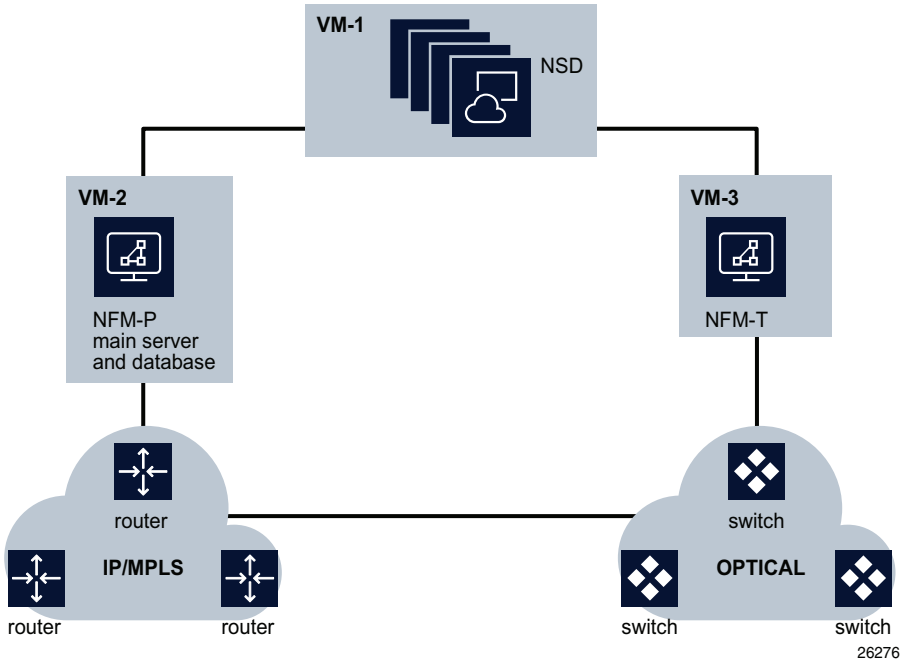

NFM-P

When integrated with the NSD, the NFM-P provides a managed IP/MPLS network model. The NSD leverages this model to perform automated service provisioning and modification on the NFM-P's network.

NFM-T

When integrated with the NSD, the NFM-T provides a managed optical network model. The NSD leverages this model to perform automated service provisioning and modification on the NFM-T's network.

7.1.4 NSD service access QoS

This feature includes an implementation of a normalized model for access QoS: a generic QoS policy that can be used for 7450 ESS, 7750 SR, 7210 SAS, and third party routers. The enhancement in QoS also facilitates Bandwidth on Demand functionality.

This feature includes the following QoS setup and usage procedures:

1. The operator/admin defines QoS catalog including, for example, Gold, Silver, and Bronze categories. This user also uses the NFM-P GUI to define the QoS Generic Policies. The policy model is generic, therefore, it can be applied to 7450 ESS, 7750 SR, 7210 SAS, and third party routers.

2. While provisioning an endpoint, either via ReST NBI or Service Fulfillment application, the user (a tenant user or operator) can select one of the predefined QoS categories (gold, silver, or bronze).
3. If Bandwidth on Demand is used (meaning the bandwidth constraints are modified), then the user can only select another policy.

   The behavior is as follows for each of the supported node categories:
   - 7450 ESS and 7750 SR: the changes only affect the SAP that is being changed
   - 7210 SAS: queue changes are not addressed
   - Third party routers: changes are defined by the corresponding driver

7.1.5 Brownfield LSP and SDP tunnels

The NSD is capable of discovering LSP and SDP tunnels created previously within the NFM-P, including multi-vendor LSP and SDP tunnels, with the following exceptions:

- A single SDP tunnel that uses multiple LSPs
- Multiple SDP tunnels that use the same LSP

Service tunnels (SDP)

The NSD will discover, and will allow users to create services with bandwidth constraints on service tunnels created previously within the NFM-P. The NSD will operate with initial allocated bandwidth on these tunnels and will keep track of used bandwidth for all the services created by the NSD. It is assumed that the NSD is the only entity creating services on these tunnels. The NSD can be used to delete or resize the allocated bandwidth, as well as to modify the LSPs associated with service tunnels previously created within the NFM-P.

RSVP-TE LSPs

The NSD will discover, and will allow users to create service tunnels on RSVP-TE LSPs created previously within the NFM-P. The NSD will operate with initial allocated bandwidth on these LSPs and will keep track of used bandwidth for all the service tunnels created on these LSPs by the NSD. It is assumed that NSD is the only entity creating service tunnels on these LSPs. The NSD cannot be used to delete, resize the allocated bandwidth, or modify LSPs previously created within the NFM-P.

Bandwidth update on existing LSP

When the reserved bandwidth of a previously-discovered LSP is modified, the NSD receives an event, and will update the both initial and available bandwidth on the LSP and SDP tunnel. This case should apply to all LSP and SDP tunnels managed by the NSD, regardless of their origin, with the following exceptions:

- A single SDP tunnel that uses multiple LSPs
- Multiple SDP tunnels that use the same LSP

The following bandwidth utilization considerations apply:

- When an LSP is used by an SDP tunnel, but is not yet bound to any service, the initial and available bandwidth of the LSP is updated. However, since the LSP is used by an SDP tunnel,

the SDP tunnel takes the entire bandwidth. As there are no services using the SDP tunnel, the available bandwidth must be equal to current bandwidth.

- When an LSP is used by an SDP tunnel and there are services bound to the SDP tunnel, the SDP tunnel will take all the LSP current bandwidth. The LSP available bandwidth must be 0, and depending on the services bound to the SDP tunnel, the available bandwidth of the SDP tunnel is adjusted to reflect the current bandwidth, minus the total bandwidth of all services running on that SDP tunnel.

7.2 MDM

7.2.1 MDM description

NSP Release 18.6 delivers a model-driven mediation (MDM) framework to support the management of new Nokia and multi-vendor network elements. The MDM is a component within the NSP architecture that provides mediation between model-driven NSP applications and Nokia or third-party network devices. The MDM provides an adaptation layer that uses adaptors to convert NSP application requests to device specific directives using standard protocols such as NETCONF, SNMP and CLI over SSH or Telnet. MDM adaptor bundles must be installed to allow NEs to be discovered by the MDM and managed through the MDM-aware applications. For details, see the *NSP Deployment and Installation Guide*.

The MDM servers can be optionally deployed in an NSP complex with NSD and NRC. The MDM server is installed as a component of the NSD_NRC. To include MDM in the NSP server installation, you must add an [mdm] entry to the hosts file of the NSD and NRC modules.

7.2.2 MDM and NSP applications

The NSP model-driven applications can be used in the same way with any NE, regardless of the NE mediation source. In this release, the NSP supports the following MDM-aware applications:

- Service Fulfillment
- Network Supervision
- Modeled Device Configurator
- Device Administrator
- Fault Management
- Telemetry

Note: Restrictions may apply to the MDM-aware applications. See the *NSP Release Notice*.

You can use the Service Fulfillment application to create a service between NEs from different vendors and managed using different mediation protocols.

For information about accessing MDM functionality through these applications, see 3.1.4 “Help system” (p. 22).

For information about accessing MDM functionality through the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

7.2.3 LLDP link discovery

In the MDM framework, the NSD supports the LLDP-based discovery of physical links in the network topology. The discovery mechanism uses the MDM to read the LLDP neighbor information. The NSD adds the discovered links to the common store so that all applicable applications can access the link information. This functionality is supported only for the physical links with the “Nearest Bridge” transmission scope.

Part III: Services and Templates

Overview

Purpose

This volume describes the services, templates, policies, and tasks that can be created using the NSD and NRC applications.

Contents

8 Services

Service description

8.1 Introduction

8.1.1 Overview

This section describes each of the service types that can be provisioned from the Service Fulfillment application. Every type of service that can be provisioned can later be edited. The editing process generally consists of modifying the same parameters that were initially populated during service provisioning, with the exception of the service name.

For information about provisioning services using the Service Fulfillment application, see 3.1.4 “Help system” (p. 22).

For information about provisioning services using the NSD and NRC modules' REST APIs, see 3.1.3 “REST APIs” (p. 21).

8.1.2 Saving service configuration without deployment

The Service Fulfillment application allows you to save a configured service without deploying it to the network. The system stores all user-configured parameters in the database for the service. However, a saved service does not reserve any bandwidth and network resource, and the system does not check the availability of resources for a saved service. The system checks for resources when the saved service is deployed to the network, and validates the resources or returns the appropriate errors to the user.

You can save only fully configured services. The Service Fulfillment application treats any saved service as a service that a user can deploy immediately. The system does not support saving services that are partially configured.

The Service Fulfillment application allows you to modify a saved service. Then you can either deploy the service with the modified configuration or save the modified service again. A modified saved service does not reserve any network resources, which are allocated and validated only when a service is deployed.

8.2 Object life cycle

8.2.1 Planning phase

Object Life Cycle (OLC) is used to manage state transitions of objects inside the NSD as they go from the planning phase to the deployment phase. The planning phase includes four states:

- Planned
- Routing
- Routing Failed
- Routed and Save

8.2.2 Deployment phase

The deployment phase includes five states:

- Waiting for Deployment
- Deploying
- Partially Deployed
- Deployment Failed
- Deployed

Note: Services that were not created by the NSD cannot be modified by the NSD. The state of a service discovered or re-synchronized from the NFM-P is Deployed.

8.3 Service CAC

8.3.1 Bandwidth CAC and validation

The NSD can perform bandwidth CAC and validation on access ports. Every port available for use in E-Line, C-Line, E-LAN, and L3 VPN services will have their available ingress bandwidth and available egress bandwidth displayed as read-only properties in the Service Fulfillment application and the NSP's REST APIs. When any of these ports are discovered, available bandwidth is initialized to port speed. In some cases, such as the 60-port 10/100 card when the port is operationally down, the port speed is zero. On fixed port speed cards, the port speed is populated, allowing services with bandwidth to be configured even when the port is down.

Note: Service CAC is not available on the variable-speed SFP-based cards.

Any changes to port speed will be reflected in the displayed available ingress bandwidth and available egress bandwidth. This may result in these fields displaying a negative value. No alarms or notifications will occur but a WARN level log will be generated.

Service CAC is disabled by default. For information about enabling service CAC, see the *Service Fulfillment* application.

Note: Service CAC is supported on services originating from the NFM-P that have had their 'NSD-managed' flag enabled.

Note: Service CAC is not supported on multi-vendor services for which there is no access port bandwidth tracking.

8.3.2 Bandwidth calculation and booking

A formula is used to calculate both the ingress and egress aggregate bandwidth of all endpoints used by E-Line, C-Line, E-LAN, and L3 VPN services. The formula yields the sum of the CIR values, which is based on each of the configured queues and the scheduler policy of the QoS. This same value is used for E-Line service tunnel bandwidth calculation. No overbooking is applied to the formula. When the NSD creates a service on one of these endpoints, the validation code will make sure that the sum of the formula is less than, or equal to, the current available bandwidth on the port, otherwise the service will not be created and an error is returned.

The bandwidth is only booked after the traversal operation is run to match with the current behavior of the core bandwidth. It is possible that between the validation check and the traversal operation, the port bandwidth was consumed by another service. In this case, the OLC state is changed to Routing Failed, and the user is told that either the access port ingress or egress bandwidth was exceeded. Modifying the CIR will reinitiate the traversal operation. Similar operations occur when adding endpoints to an existing service and modifying endpoints. In the latter case, it is the bandwidth delta which is applied to the available ingress or egress bandwidth. Upon deletion of an endpoint or service, the available ingress or egress bandwidth is increased by the bandwidth of the endpoints.

Service CAC is available on both access and hybrid ports. If there are network interfaces on hybrid ports, these are not tracked as part of the available ingress or egress bandwidth. When an upgrade is performed, the available ingress and egress bandwidths will be calculated based on all existing services within the NSD. This may result in negative values. When in an overbooked state, any

request that will not cause a change to bandwidth reservation, or that will cause a shrink in bandwidth reservation, will be permitted.

8.4 E-Line services

8.4.1 E-Line service description

An E-Line service connects two customer Ethernet ports over a WAN. The NSD supports the creation of E-Line services over IP networks. When an E-Line service is deployed, the selection of the endpoints utilizes automatically the requisite technology tunnels. For example, when the tunneling technology is MPLS, a service tunnel with a single LSP satisfying the service-specified constraints and objectives is automatically selected. The service is then bound to that LSP via the service tunnel. The NSD tracks the LSP available bandwidth and adjusts it automatically to accommodate the E-Line service, which reserves bandwidth on the LSP.

If an existing E-Line service is modified (for example, to increase bandwidth), the service tunnel is resized to accommodate it, if permitted by policy. If the service tunnel resizing fails, the service tunnel may be rerouted onto links that cannot accommodate the resized service tunnel. If the reroute fails, then a new service tunnel is created. It is possible for E-Line services to use service tunnels that were not created using the NSD.

Note: Policies for service-to-tunnel binding dictate the rules associated with the service binding. If no service tunnel meets all the constraints, and this is a new E-Line service, a new service tunnel is created.

Other parameters of the E-Line service are obtained from the specific templates referenced in the abstract API definition. The service definition in the abstract API, the detailed configuration in the service templates, and other network and tunnel parameters form the complete service definition, which is represented in the normalized model for E-Line. Specific configurations based on the devices are then constructed and deployed using the NFM-P.

Note: You can provision SAP-to-SAP E-Line services if you select different ports for each endpoint.

E-Line multi-vendor support

The NSD supports the following multi-vendor endpoint combinations for E-Line services:

- Cisco-Nokia
- Juniper-Nokia
- Cisco-Juniper
- Cisco-Cisco
- Juniper-Juniper

The following considerations apply to the E-Line multi-vendor support:

- Cisco LSP names must be in the format of Tunnel*number*, where *number*is an integer between 0 and 65535.
- Cisco LSP-Path Bindings contain a property called Path Option. This property is be set to 1 for primary and 2 for secondary.
- Cisco and Juniper endpoints support only secondary paths, and do not support standby paths. When Cisco or Juniper endpoints are used and the Tunnel Creation Template has the Protection Type set to *Standby*, the NSD creates secondary paths instead.

When creating an E-Line service on multi-vendor NEs, the NSD attempts to find a tunnel based on the criteria specified in the Tunnel Selection Profile (TSP). If no tunnel exists, and the TSP specifies that new tunnels must be created, then the NSD creates MPLS RSVP-TE tunnels, including the Dynamic LSP and LSP-Path Bindings.

Brownfield E-Line services

The E-Line services created within the NFM-P (brownfield E-Line services) can be managed by the NSD. In order for the NSD to discover these services, their “NSD-managed” flag must be enabled within the NFM-P. Once discovered by the NSD, these services will function the same as E-Line services created within the NSD itself, provided that they meet the NSD requirements. Any change made to these services within NFM-P after discovery will be propagated to the NSD, provided the change impacts the topology of the service.

Note: E-Line services created within the NFM-P have an “Auto-delete” flag. When enabled, services without service sites are automatically deleted. This flag should not be enabled on services being managed by the NSD, as the “NSD-managed” flag is disabled upon service deletion, and remains so even if the service is recreated and resynchronized into the NSD.

8.4.2 Multi-domain E-Line services

The NSD supports multi-domain E-Lines that span any mix of MPLS and non-MPLS domains. The service tunnels must be already created in the MPLS domains. The non-MPLS domains can consist of only peer-to-peer Ethernet links.

In addition to the SAP-to-SAP and SAP-to-SDP service sites, the multi-domain E-Line service also supports SDP-to-SDP connections through the use of pseudowire switching. However, the NEs eligible for SDP-to-SDP pseudowire switching must be pre-configured with a pw-switching flag that is enabled on the NE.

The NSP calculates an optimal end-to-end path that traverses existing service tunnels, including VLAN handoff. Only strictly-routed RSVP-based service tunnels have calculations for the number of hops and accumulated IGP metric and latency. The VLAN handoffs have hard-coded hops, IGP metric and latency to 1. Other service tunnels have very large numbers for hops, IGP metric and latency and are usually less preferred. The non-RSVP service tunnels have zero bandwidth.

Multi-domain E-Line types

The multi-domain E-line service provisioning allows you to create the following types of E-Lines:

- vc-switched—in addition to the two terminating sites, an E-Line can include one or more switching sites
- composite—a composite E-Line consists of multiple component services connected through VLAN handoff to provide end-to-end connectivity

  A composite E-Line can include a vc-switched e-line.

To support the multi-domain functionality, the E-Line service template provides the VC Type parameter, which allows you to specify the type of pseudowire for the E-Line service.

8.4.3 EVPN-based E-Line service

The NSD supports the creation of EVPN-based E-Line services over tunnel types that are supported in a BGP-EVPN MPLS context. The EVPN-based E-Line service is not established over pseudowire. You can configure EVPN-based E-Line services on all the Nokia NEs that support EVPN. The configuration of EVPN-based E-Line services on multi-vendor NEs is not supported.

Configuration

To configure an EVPN-based E-Line service, you need to start the E-Line service creation in the Service Fulfillment application, as usual, and select the Enable EVPN Tunnel Selection check box in the Additional Properties form. After enabling the EVPN service, you are able to select a tunnel type from the following options: LDP, RSVP-TE, SR-ISIS, SR-OSPF, SR-TE, and BGP. There is also the ANY option, which indicates to the nodes that any supported tunnel type in the EVPN context can be selected following the order of preference.

Considerations

The following considerations apply to the EVPN-based E-Line service configuration in the NSD:

- The NSD supports only the configuration of greenfield EVPN-based E-Line service. The modification of existing EVPN-based E-Line services that were created in the NFM-P is not supported.
- The NSD assumes that the network is correctly configured to support the selected tunnel type. The service can fail if the network is not correctly configured. For example, if the network does not have SR-TE LSPs configured, then an EVPN-based E-Line service configured with the SR-TE tunnel type is operationally down.
- The tunnel type parameter is modifiable, as required. However, the NSD does not support switching from the EVPN-based E-Line (the Enable EVPN Tunnel Selection check box is selected) to a pseudowire-based E-Line (the Enable EVPN Tunnel Selection check box is not selected).
- Each service is associated with a unique EVPN instance (EVI) number that the NSD generates automatically and then sends to the NE to auto-derive the unique RD/RT for the NE. The NSD synchronizes the EVIs defined in the network to ensure the EVI uniqueness.
- The EVPN-based E-Line service uses an Ethernet Tag (eth-tag) that is pre-configured by the NSD and not visible in the GUI. The NE uses the Ethernet Tag to identify its remote BGP peer and establish the MP-BGP connection.

Service template

To ensure consistency when configuring multiple similar services, you can create EVPN-based E-Line service templates that you can then apply to your service. Just select the appropriate Tunnel Type for EVPN-based E-Line in the template properties, as required.

8.4.4 E-Line service with MC-LAG termination and pseudowire redundancy

In this release, the NSD supports the creation of E-Line services with MC-LAG termination and pseudowire redundancy at demo quality. The NSD support for MC-LAG and pseudowire redundancy requires the preconfiguration of MC-LAG on the NFM-P, including the MC Peer Group and the MC Lag Group. The NSD discovers the preconfigured MC LAG Group as a service object

in the physical layer. This service object can be selected as an endpoint for service creation. The MC-LAG object represents both the active and the standby NE members of the LAG.

An E-Line service with MC-LAG termination and pseudowire redundancy has zero bandwidth in the core.

Service creation overview

When an MC-LAG is selected as an endpoint for E-Line service creation, the NSD automatically considers it as two separate endpoint requests: one request for each LAG member. The NSD decomposes the request into two endpoint requests: one for the active NE LAG and one for the standby NE-LAG. All the configuration in the original request, including the encapsulation values, is copied to two new separate requests. As a result, the NSD assumes that both MC-LAG members are configured with the same encapsulation mode and type.

The NSD supports the creation of E-Line service with MC-LAG termination and pseudowire redundancy between the following endpoint types:

- from regular port to MC-LAG
- from SC-LAG to MC-LAG
- from MC-LAG to MC-LAG

Note: A regular endpoint consists of a single NE and a single port. When a regular endpoint is selected, the NSD creates a service site on that NE using the specified port as the termination point.

Endpoint configuration

For the SC-LAG, two VLL endpoints must be configured on the NFM-P:

- The first VLL endpoint must have an SC-LAG SAP.
- The second VLL endpoint must have a primary and a secondary spoke SDP bindings to the two SC-LAG members.

For the MC-LAG, two VLL endpoints must be configured on the NFM-P:

- The first VLL endpoint must have an MC-LAG SAP and a spoke SDP binding destined for its MC-LAG peer, with ICB enabled.
- The second VLL endpoint must have one or two spoke SDP bindings to the far end.

8.5 C-Line services

8.5.1 C-Line service description

C-Line services connect two SAPs that can be defined on SONET/SDH, DS3/E3,T1/E1 ports or TDM channels. The NSD supports the creation of C-Line services over IP networks. When a C-Line service is deployed, the selection of the endpoints automatically utilizes the requisite technology (MPLS or L0 WDM) tunnels.

It is possible for C-Line services to use service tunnels that were not created using the NSD.

Note: Policies for service-to-tunnel binding dictate the rules associated with the service binding. If no service tunnel meets all the constraints, and this is a new C-Line service, a new service tunnel is created.

Other parameters of the C-Line service are obtained from the specific templates referenced in the abstract API definition. The service definition in the abstract API, the detailed configuration in the service templates, and other network and tunnel parameters form the complete service definition, which is represented in the normalized model for C-Line. Specific configurations based on the devices are then constructed and deployed using the NFM-P.

Note: The SAP-to-SAP C-Line services can be provisioned if different ports are used for each endpoint.

Brownfield C-Line services

C-Line services created within the NFM-P can be managed by the NSD. In order for the NSD to discover these services, their "NSD-managed" flag must be enabled within the NFM-P. Once discovered by the NSD, these services function the same way as C-Line services created within the NSD, provided that they meet the NSD requirements. Any change made to these services within NFM-P after discovery is propagated to the NSD if the change impacts the topology of the service.

Note: The C-Line services created within the NFM-P have an “Auto-delete” flag. When enabled, services without service sites are automatically deleted. This flag must not be enabled on services managed by the NSD, as the “NSD-managed” flag is disabled upon service deletion, and remains so even if the service is recreated and resynchronized into the NSD.

C-Line NE support

For C-Line creation, the NSD supports the 7x50 and 7705 NE types. Third-party vendor NEs are not supported.

VC types

The C-Line service creation requires you to specify a type of VC (pseudowire). The options are:

- SAToP T1 (unstructured DS1)
- SAToP E1 (unstructured E1)
- CESoPSN (structured)
- CESoPSN CAS (structured with CAS)

8.6 E-LAN services

8.6.1 E-LAN service description

E-LAN services are configured with the same parameters that are used for E-Line service creation. Objectives/constraints are enforced for the LSPs. The default endpoint QoS template is applied to all endpoints. Zero bandwidth is reserved in the core. E-LAN services can use service tunnels that were not created using the Service Fulfillment application.

E-LAN multi-vendor support

The NSD does support the creation of most of the E-LAN services on Cisco nodes. When this is done, the NSD configures a property called bridgeDomainId during the site creation.

The NSD does not support the creation of E-LAN services on Juniper nodes.

8.6.2 EVPN-based E-LAN service

The NSD supports the creation of EVPN-based E-LAN services over tunnel types that are supported in a BGP-EVPN MPLS context. The EVPN-based E-LAN service is not established over pseudowire. You can configure EVPN-based E-LAN services on all the Nokia NEs that support EVPN. The configuration of EVPN-based E-LAN services on multi-vendor NEs is not supported.

The EVPN-based E-LAN service supports the same topology types as those supported by the pseudowire E-LAN service: Hub and Spoke and Full Mesh.

Configuration

To configure an EVPN-based E-LAN service, you need to start the E-LAN service creation in the Service Fulfillment application, as usual, and select the Enable EVPN Tunnel Selection check box in the Additional Properties form. After enabling the EVPN service, you are able to select a tunnel type from the following options: LDP, RSVP-TE, SR-ISIS, SR-OSPF, SR-TE, and BGP. There is also the ANY option, which indicates to the NSD that any supported tunnel type in the EVPN context can be selected following the order of preference.

Considerations

The following considerations apply to the EVPN-based E-LAN service configuration in the NSD:

- The NSD supports only the configuration of greenfield EVPN-based E-LAN service. The modification of existing EVPN-based E-LAN services that were created in the NFM-P is not supported. The opposite is also true: EVPN-based E-LAN services that were created in the NSD cannot be modified in the NFM-P.
- The NSD assumes that the network is correctly configured to support the selected tunnel type. The service can fail if the network is not correctly configured. For example, if the network does not have SR-TE LSPs configured, then an EVPN-based E-LAN service configured with the SR-TE tunnel type is operationally down.
- The tunnel type parameter is modifiable, as required. However, the NSD does not support switching from the EVPN-based E-LAN (the Enable EVPN Tunnel Selection check box is selected) to a pseudowire-based E-LAN (the Enable EVPN Tunnel Selection check box is not selected).

- Each service is associated with a unique EVPN instance (EVI) number that the NSD generates automatically. The NSD synchronizes the EVIs defined in the network to ensure the EVI uniqueness.

  In the Full Mesh topology, the NSD sends the EVI to the NE to auto-derive the RD/RT for the NE. The RT import and export label is the same for all NEs in the Full Mesh topology.

  In the Hub and Spoke topology, the NSD sends the EVI to the NE to auto-derive only the RD for the NE. The NSD generates the RT to ensure that the RT import and export labels on the hub endpoint are inverted with respect to the RT import and export labels on the spoke endpoints. That is, the hub RT import label matches the spoke RT export label, and the hub RT export label matches the spoke RT import label. The unique RT labels are stored in the cache, and the NSD continuously resyncs the existing RT labels from the NFM-P to ensure uniqueness.

Service template

To ensure consistency when configuring multiple similar services, you can create EVPN-based E-LAN service templates that you can then apply to your service. Just select the appropriate Tunnel Type for EVPN-based E-LAN in the template properties, as required.

8.7 IES services

8.7.1 IES service description

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

8.8 L3 VPN services

8.8.1 L3 VPN service description

The NSD supports the creation of L3 VPN services. L3 VPN services utilize layer 3 VRF (VPN/ virtual routing and forwarding) to routing tables for each customer utilizing the service. The customer peers with the service provider router and the two exchange routes, which are placed into a routing table specific to the customer. Multiprotocol BGP (MP-BGP) is required to utilize the service.

The RD and RT is auto-generated as per policy direction and the topology type selected. Other parameters specified in the referenced template complete the service definition. Other parameters of the L3 VPN service are obtained from the specific templates referenced in the abstract API definition. The service definition in the abstract API, the detailed configuration in the service templates, and other network and tunnel parameters form the complete service definition, which is represented in the normalized model for L3 VPN. Specific configurations based on the devices are then constructed and deployed using NFM-P. L3 VPN services can use service tunnels that were not created using the Servic Fulfillment application.

Note: Before provisioning L3 VPN services using the NSD, the user must have MP-BGP configured and working between the PE nodes to support IP VPN. The Peer CE nodes also need to be well configured. Only one AS is supported per provider.

8.8.2 Multi-domain L3 VPN service provisioning from L2 endpoints

Multi-domain L3 VPN services from L2 metro areas are supported. These services are created between PE routers on metro areas. However, because some PE routers are not L3 capable, the NSD performs the path search across the network, from L2 metro areas to L3 core, and finds the best exiting routers from metro to core. Then, the NSD provisions L2 E-Line services on all metro areas and L3 VPN services in the core. Finally, the services are stitched together by the NSD using VLAN hand-off.

The intra-domain tunnels must be created in advance, and all metro domains are interconnected via Ethernet links (VLAN handoff) to the core. Since none of the routers on L2 metro domains are L3 VPN capable, the NSD uses this property to run the path search algorithm. This property can be set using the NSP's REST APIs.

The NSD uses L2 and L3 service templates to define the common attributes for the auto-created services. Profiles are used for QoS and the auto-assignment of L3 RD/RT. The NSD also uses the tunnel selection profile to include and exclude specific tunnels during path search. The path search objectives (such as minimizing hop or cost) and other values specific to the VPN (such as the IP addresses of the L3 access points) are defined either from the Service Fulfillment application or the NSP REST APIs. The NSD uses the QoS CIR values to book the bandwidth on tunnels.

8.8.3 L3 VPN services on Wavence SM NEs

The NSD NRC supports the creation of L3 VPN services on Wavence SM NEs that support L3 VPN through SNMP. To create such a service in the Service Fulfillment application, choose L3 VPN as the service type and then select endpoints that are already configured on Wavence SM NEs.

Note: This feature does not support the creation of L3 VPN services on a mix of Wavence SM and other NE endpoints.

This release supports only Greenfield scenarios (services provisioned in the Service Fulfillment application).

When you create an L3 VPN service on Wavence SM NEs in the Service Fulfillment application, take into account the following considerations:

- Static LSP creation must be completed prior to the creation of the L3 VPN service.
- Both Full Mesh and Hub and Spoke topology types are supported. When Hub and Spoke is selected, the NSD will automatically assign ingress and egress labels on the spoke SDP bindings.
- In the service's advanced properties, set the Auto Bind parameter to None.
- On each service endpoint, configure the Outer Tag, Primary IP Address, and GQP parameters. If required, also configure the Static Route, Next Hop and Preference parameters associated with the Primary IP Address.

When the Full Mesh topology type is selected, the Service Fulfillment application will automatically configure static routes on all endpoints to achieve the required mesh topology.

When the Hub and Spoke topology type is selected, the Service Fulfillment application will automatically configure the black hole static routes on all spoke sites to prevent traffic between them. Default route static routes will also be automatically configured on Spoke endpoints with the next hop being the Hub endpoint.

Note: Only one hub endpoint can be configured, regardless of topology type. This hub endpoint must be configured in the SR domain.

8.8.4 Composite L3 VPN service across multiple Wavence SM domains and one 7x50 SR domain

The NSD supports the creation of a composite L3 VPN service across multiple domains of Wavence SM NEs without MP-IBGP and one domain of 7x50 SR NEs with MP-IBGP.

A physical link (VLAN uplink) is required between access ports on the adjacent Wavence SM and 7x50 SR that connect the two domains.

Prerequisites

To be able to create a composite L3 VPN service, you must perform pre-configuration tasks on all the Wavence SM and 7x50 SR NEs that are part of the service, and on the NFM-P.

Configuration in the Wavence SM NE domain:

- Provision static LSPs between all the Wavence SM NEs

Configuration in the 7x50 SR domain

- Create an L3 VPRN service with service tunnels between all the 7x50 SR NEs. IGP, MPLS (RSVP-TE or LDP) and MP-IBGP must be configured on the NEs.

Configuration on the NFM-P

- Create a physical link between network ports on the Wavence SM NEs.
- Create a physical link between network ports on the adjacent Wavence SM NE and 7x50 SR NE that connect the two domains.

  If the physical links are not present, then the L2 service fails to deploy.

Service creation considerations

To configure a composite L3 VPN service successfully, apply the following guidelines during the service creation:

- Set the Tunnel Type to None. The composite service does not support a specific tunnel type across the two domains.

During the service creation, the NSD performs the following tasks automatically:

- Provisions both local and remote static routes to configure the service level (VRF) route table on each Wavence SM NE.
- Provisions static routes on each end of the physical link (VLAN uplink) to enable the inter-domain IP routing.
- Creates two L3 internal service endpoints, one on each side of the physical link.

8.8.5 L3 VPN multi-vendor support

For L3 VPN services, the NSD supports the RSVP-TE option, since multi-vendor nodes do not support SDP tunnels. As a result, if an L3 VPN service is created on a multi-vendor node, the NSD algorithm tries to find or to create RSVP-TE tunnels and always sets the auto-bind property to RSVP-TE on the multi-vendor nodes.

8.9 Other services

8.9.1 E-Access services

The NSD supports the creation of E-Line services between an NE that is managed by an NSD instance and an NE that is located in a different domain and managed by a different NSD instance or by a third-party system. This applies to cases in which your NSD does not manage the entire network. This type of service is called an E-Access service. The NSD supports the creation of an E-Access service only by way of the NSD REST API. To create an E-Access service, use the POST /api/v4/services/eaccess operation, Create an IP E-Access service.

For the E-Access service creation to work, you must use an SDP tunnel, a path and an LSP that were already configured on each NE between the two NEs. You also need to configure a Tunnel Selection policy to apply to the E-Access service. Select the Use existing tunnels option as the service provisioning rule.

A few pointers to help you configure an E-Access service:

- The NSD books bandwidth only at the Access interface level, not on the LSP in the core.
- The service uses the VC-ID label specified on the Adjacency on the spoke-sdp binding.
- If the Service ID is not already in use, the NSD assigns the VC-ID label as the Service ID. If that Service ID is already in use, the NSD auto-assigns the next available ID.
- If you specify a VC-Type label, then it must match the VC-Type defined at the far end. The default value is Ethernet_Tagged_Mode, which maps to VLAN on the NE.

After creating the E-Access service, you can view the service and its properties in the Service Fulfillment application.

- The Service layer shows just the NE, along with the service endpoint, that is in the network managed by your NSD.
- The Service Tunnel layer shows the Service Tunnel connection. This is the existing SDP that you created for the service.
- The MPLS layer shows the LSP. This is the existing LSP that you created for the service.
- The service does not show an IGP layer because your NSD does not manage the far-end NE and therefore is not aware of the IGP path.
- The service does not show a physical layer because your NSD does not manage the far-end NE.

8.9.2 Services on SDPs with multiple loopback IP addresses

The NSD supports the configuration of services on SDP tunnels using a loopback IP address as either the source or destination IP address when routing services. A potential benefit of having services on SDP tunnels using a loopback IP address is the ability to configure routing on tunnels established on different paths between two NEs. However, you can configure such services only on

brownfield SDP tunnels that were created in the NFM-P or on NEs. The NSD does not support the creation of new service tunnels using loopback IP addresses.

Before you start configuring a service in the NSD, you must create SDP tunnels with loopback IP addresses in the NFM-P or on the NE. The following list captures the high-level configuration tasks required for each NE.

- Configure the loopback interfaces on routers.
- Configure peers on the targeted LDPs. Use the loopback interface name as the local-lsr-id option and enable tunneling to enable LDP over the tunnels.
- Configure the SDP tunnels using the loopback interface IP addresses for the service far end. Optionally, you can apply a steering parameter to the tunnel to help the selection of the correct SDP tunnel when creating the service.

The service tunnels that you created can be viewed in the Service Fulfillment application, on the INVENTORY tab. The service tunnel Destination IP is the IP address of the loopback interface and the service Transport type is MPLS.

If you applied the optional steering parameter to the tunnel, then you can also create a tunnel selection policy for the steering parameter in the Policy Management application.

Service provisioning

8.10 To enable NSD management on services created in the NFM-P

8.10.1 NSD Managed parameter

The NSD can manage a service that was created in the NFM-P if the following conditions are met:

- The NSD supports the service type.
- The service NSD Managed parameter is enabled in the NFM-P.

A service created in the NFM-P and discovered in the NSD functions the same as a service created in the NSD. Any change that impacts the topology of the service made in the NFM-P after discovery is propagated to the NSD.

This procedure describes how to enable the NSD Managed parameter for a service in the NFM-P.

8.10.2 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Click Search and choose a service.

The following table maps the service names defined in the NFM-P to the corresponding NSD service names.

NFM-P service	NSD service
CPIPE	C-LINE
EPIPE	E-LINE
VPLS	E-LAN
VPRN	L3 VPN

3

Enable the check box in the NSD Managed column for the selected service.

4

Save your change and close the form.

End of steps

9 Templates and policies

9.1 Introduction

9.1.1 Scope

This chapter describes the templates and policies that you can provision using the Policy Management application. You can edit every template or policy that you provisioned. The editing process generally consists of modifying the same parameters that were populated during provisioning, with the exception of the template/policy name.

For information about provisioning services using the Policy Management application, see 3.1.4 “Help system” (p. 22).

For information about provisioning templates and policies using the NSD and NRC modules' REST APIs, see the *NSP Developer portal*.

9.2 Service templates

9.2.1 E-Line and C-Line service templates

The NSD and NRC modules support the creation of E-Line and C-Line service templates. The configuration of these templates can be applied to the E-Line or C-Line service creation form in the Service Fulfillment application, thereby simplifying service provisioning. If an E-Line or C-Line service uses a template that specifies the same parameters as those specified via the NSP's REST APIs or the E-Line or C-Line service creation form in the Service Fulfillment application, then the template parameters in are overridden.

9.2.2 E-LAN service templates

The NSD and NRC modules support the creation of E-LAN service templates. The configuration of these templates can be applied to the Service Fulfillment application's E-LAN service creation form, thereby simplifying service provisioning. If an E-LAN service uses a template that specifies the same parameters as those specified via the NSP's REST APIs or the Service Fulfillment application's E-LAN service creations form, the template's parameters are overridden.

9.2.3 L3 VPN service templates

The NSD and NRC modules support the creation of L3 VPN Service templates. The configuration of these templates can be applied to the Service Fulfillment application's L3 VPN service creation form, thereby simplifying service provisioning. If an L3 VPN service uses a template that specifies the same parameters as those specified via the NSP's REST APIs or the Service Fulfillment application's L3 VPN service creation form, the template's parameters are overridden.

9.2.4 Endpoint QoS templates

The NSD and NRC modules support the creation of Endpoint QoS templates.

9.2.5 Generic QoS Policies

The NSD and NRC modules support the creation of Generic QoS policies.

9.3 Service policies

9.3.1 RD/RT Range policy

The NSD and NRC modules support the modification of the global RD/RT Range policy. The RD/RT Range policy is a single default policy that applies to all L3 VPN services. In the future, multiple RD/RT Range policies may be supported.

9.3.2 Tunnel Profile policy

The NSD and NRC modules support the creation of Tunnel Profile policies to apply selection and creation rules to service tunnels. You can use tunnel profile policies when creating services in the Service Fulfillment application. Under the service Additional Properties, choose a tunnel profile that is appropriate for your service, as required.

The Tunnel Profile policy includes Tunnel Selection Rules, which allow you to enable one or multiple rules that are applied in sequence to the service.

The Tunnel Creation Rules control the tunnel maintenance attributes, such as the consumption of the tunnel by all services, the automatic deletion of the tunnel when there are no more services attached to it, the modification of tunnel parameters by services, and the tunnel protection.

9.3.3 Steering Parameters policy

The NSD and NRC modules support the creation of Steering Parameter policies.

9.4 IP/MPLS policies

9.4.1 System IP MPLS Configuration

The System IP MPLS Configuration policy allows you to configure the maintenance mode of an IP link. You can choose one of the following maintenance modes:

- Manual
  You must manually resignal LSPs, trigger GCO or wait for network changes to move resources from the affected link and back onto the best path.
- Automatic
  The NRC-P automatically moves LSPs from the affected resources, usually using resignalling.

9.4.2 Router ID Mapping policies

The NRC-P is able to discover and display multiple IGP instances (OSPF and ISIS), which are each discovered as a unique domain. These domains are interconnected on the same routers, which themselves have multiple instances defined. If the Router IDs for these instances are the same, they will be displayed as a single router on the Service Fulfillment application's multi-domain topology maps. If the Router IDs are different, a Router ID Mapping template must be provisioned in order for the instances to be displayed as a single router on the Service Fulfillment application's multi-domain topology maps.

9.4.3 Path Profile policy

The NSD and NRC modules support the configuration of path profile templates, which are associated to path requests by PCCs. A default path profile template can also be configured, if required. By default, path profile templates will optimize on metric. Additional behavior can be specified, such as bidirectionality for forward and reverse paths between a pair of sources or destinations, and path disjointness between two paths specifying the same profile. A path request can also contain multiple profiles. Path profile templates can also be specified on the PCC.

When a PCE request contains objects specifying constraints and objectives in addition to the path profile template, the following behavior is observed:

- If a PCC request has an associated path profile template, and also has the specific constraints (B = 1) in the METRIC object (such as bandwidth, IGP metric, and TE metric values), then the path computation will use the PCC-specified values, overriding the constraint values specified in the path profile templates.
- If a PCC request has an associated path profile template and no bounds set on the values in the METRIC object, then the default values specified in the path profile template will be used in the path calculation.

Path profile templates can be applied to both SR TE LSPs, and RSVP TE LSPs. For RSVP TE LSPs, the specification of the path profile template applies to all paths for that LSP.

9.5 Mediation Profiles

9.5.1 Mediation Profile

The NSD supports the creation of mediation profiles, which you can then apply to a service template designed to configure a service in a specific way. In this release, you can apply the mediation profiles to L3 VPN, E-LAN and E-Line service templates.. The mediation profiles reference predefined MFM adaptor templates that you must create on the NFM-P, and then deploy to the NSD. The templates are based on json augmentation files that describe additional properties, which are supported at service, site and endpoint levels.

When creating a mediation profile, you need to configure the following parameters:

- Mediation Profile

  The AMI type to which you apply the template. The NSD currently supports the NfmpService and NsdServiceNBI mediation profile types. The NfmpService profile (not an AMI) is designed for NEs that are managed by the NFM-P. The NsdServiceNBI profile is designed for NEs that are managed by the MDM.

- Mediation Version

  The AMI version—it applies only to the templates/augmentations designed for MDM-managed NEs. If you upgrade the AMIs but need to maintain backwards compatibility, you can restrict the templates only to a specific mediation version.

- Global Template Name

  The template/script that implements the desired functionality.