7450 Ethernet Service Switch
7750 Service Router
Virtualized Service Router

Releases up to 26.3.R3

Multiservice ISA and ESA Advanced Configuration Guide for MD CLI

3HE 20799 AAAG TQZZA
Edition: 01
July 2026

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

Table of contents

List of tables

List of figures

Preface

About This Guide

Each Advanced Configuration Guide is organized alphabetically and provides feature and configuration explanations, CLI descriptions, and overall solutions. The Advanced Configuration Guide chapters are written for and based on several Releases, up to 26.3.R3. The Applicability section in each chapter specifies on which release the configuration is based.

The Advanced Configuration Guides supplement the user configuration guides listed in the *7450 ESS, 7750 SR, and 7950 XRS Guide to Documentation*.

Audience

This manual is intended for network administrators who are responsible for configuring the routers. It is assumed that the network administrators have a detailed understanding of networking principles and configurations.

Application Assurance — Application Identification and User-Defined Applications

This chapter describes Application Assurance (AA) Application Identification and User-Defined Applications configurations.

Topics in this chapter include:

Applicability

The information and configuration in this chapter were initially based on SR OS Release 11.0.R3, but updates were made based on SR OS Release 19.10.R2. The MD-CLI corresponds to SR OS Release 25.7.R1.

There are no specific prerequisites for this example.

Overview

This chapter is intended for Application Assurance (AA) network architects and engineers. It provides best practice information to customize the AA policy and classify any type traffic to meet the service provider reporting, charging or control requirements.

In addition to the signatures built and supported by Nokia, service providers can create their own application signatures based on various criteria. This customization capability can be used to classify traffic hosted on the provider network (web portal, streaming service) or hosted on the Internet and not yet covered by the default AA signature set.

Basics and terminology

The following main components are used for AA classification:

- Application Filters (app-filters)— define applications based on Layer 3 to Layer 7 criteria. App-filters provide a mapping between one or more protocol signatures or customized traffic patterns into an application of interest.
- Applications — such as BitTorrent®, Netflix®. Traffic is classified into applications using app-filters.
- Application Groups (app-groups)— such as peer-to-peer, multimedia streaming. For the purpose of reporting and control, applications of similar type/function can be grouped together in app-groups.

- Charging Groups — such as zero rating, default. For the purpose of charging or control, applications and app-group can be grouped together in charging groups.

The following table is a high-level example to illustrate how app-filters are used to defined applications and show their logical grouping into app-groups and charging groups.

Figure 1: App-filters – applications – app-groups – charging groups

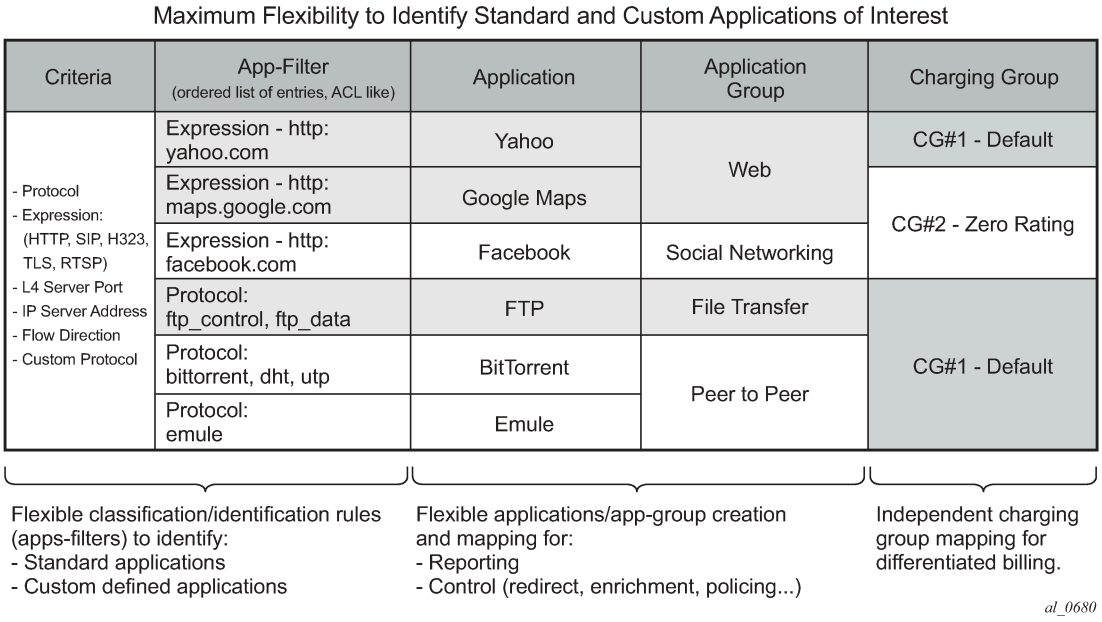

- BitTorrent® and Emule® applications are defined using their protocol signature and grouped in the P2P app-group.
- FTP application is defined using both FTP data and FTP control protocol signatures, the app is mapped in the file transfer app-group.
- Google Maps® and Yahoo® web sites are defined using http expression and grouped together in the Web app-group.

Configuration

Classification criteria (App-filter)

The operator can take full advantage of the flexible AA policy configuration to classify traffic from any application of interest using various criteria ranging from Layer 3 to Layer 7 expressions.

Expression match criteria allows to further refine traffic classification by identifying traffic from HTTP, HTTPS (SSL/TLS), SIP, H323, RTSP, Citrix protocol signatures.

The different app-filter match criteria are listed below:

- L7 expression
  - HTTP: host, URI, user agent, referer

  - SSL/TLS: certificate org name, common name, SNI
  - H323: product ID
  - SIP: URI, user agent, media type
  - RTSP: host, URI, user agent
  - Citrix: application published name
  - RTMP: page-host, page-uri, swf-host, swf-uri
- IP protocol number
- IP server address
- TCP/UDP server port
- Custom protocol
- Protocol signature

The following table shows the supported operators to define expression-based app-filters:

Table 1: Supported operators in expression-based app-filters

^ :	Expression starts with
$:	Expression ends with
*:	Wildcard - anything before or after
\I:	Forces case sensitivity
\d:	Any single decimal digit [0-9]
\.:	Any single character
*:	Asterisk character

The following table shows some examples of expression match combinations:

Table 2: Examples of expression match combinations

^abcd*:	match 'abcd' at beginning, can end with anything
abcd:	match 'abcd' anywhere
*abcd$:	match 'abcd' at the end
^abcd$:	exact expression match 'abcd'
^ab*cd$:	string starts with 'ab', ends with 'cd' (anything else in between)
^ab\dcd$:	string starts with 'ab', followed by a decimal digit, ends with 'cd'

Note:
It is possible to combine different criteria or expressions within the same filter in which case an implicit AND operation between the criteria within the same filter is done by the system.

Application definition example

The following example provides a basic configuration example with the application FTP made of two protocol signatures FTP control and FTP data; the application is mapped into the application group file transfer:

Create the application group.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-group "File Transfer" {
                    }
```

Create the application.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    application "FTP" {
                        app-group "File Transfer"
                    }
```

Create the app-filters.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            application "FTP"
                            protocol {
                                eq "ftp_control"
                            }
                        }
                        entry <1..65535> {
                            admin-state enable
                            application "FTP"
                            protocol {
                                eq "ftp_data"
                            }
                        }
```

Note:
When the application is created, the operator is expected to configure the collection of statistics at the subscriber level for this new application (usually only for business VPNs).

User-defined applications

General recommendations

To classify traffic properly, Nokia recommends to follow the guidelines and best practices defined in this section before creating a new application:

- Analyze the application traffic
  - Identify what type traffic is used (Wireshark®).
  - Use the application the same way the end user would use it, the same application can create various flows.
- Configure the appropriate App-filters
  - Following the analysis of the application done above, create the application.
  - Follow the App-filter best practices chapter to understand in which range to add the filters.
  - More than one App-filter can be required to identify a single application.

AppDB/default AA policy

The default AA policy called AppDB (Application Database) is provided by Nokia and should be used on most deployments. Contact your regional support organization for more details on how to obtain it.

This configuration includes applications and application-groups most providers can use by default and is designed to allow the addition of any custom entries required by service providers to identify additional services/applications.

The following customization options exist:

- Generating a configuration for a particular region (such as APAC)
- Generating a new configuration (or updating a configuration) containing specific applications

Before adding new entries to the template and customizing the configuration, it is recommended to follow the next guidelines on app-filters and ranges. These guidelines are key to allow an easy upgrade path from the policy configuration provided by Nokia.

App-filters

App-filters are an ordered list of entries. It is important to keep the order of this list consistent with the classification objective.

For instance, a common configuration mistake is to configure a filter rule for the HTTP protocol signature before HTTP expression filters. If that were the case, then app-filters using HTTP expressions would not be used as the system would find an acceptable match with the protocol signature before walking the list of expressions configured. This mistake is described in the following example:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
```

```
---snip---
                    app-filter {
                        entry 20100 {
                            admin-state enable
                            description "Default HTTP Protocol"
                            application "HTTP"
                            protocol {
                                eq "http"
                            }
                        }
                        entry 20110 {
                            admin-state enable
                            description "Google"
                            application "Google"
                            expression 1 {
                                type http-host
                                eq "*.google.com$"
                            }
                        }
```

This is an incorrect app-filter order. App-filter entry #20100 always matches before the HTTP expression entry #20110.

It is not necessary to specify a protocol when defining an expression filter, the protocol is implicit based on the type of expression match criteria used (for instance, HTTP, SIP, H323).

App-filter ranges

The App-filter list is an ordered list, it is key to configure each app-filter in the right order and in the correct range.

The operator can customize the policy and create applications and app-filters by using the following ranges shown in Table 3: Customer reserved App-filter ranges (other ranges are used by the default policy):

Table 3: Customer reserved App-filter ranges

Range Name	Description	Start	End
Extended top range	Top range, matches before any other filters	1	1499
High priority	Top range for high priority matches	2000	4999
Expression range A	HTTP Host, Host+URI; optionally with IP/Port match	19000	22999
Expression range B	Other Expression Match; optionally with IP/Port match	33000	34999
Extended protocols	Protocol-signature + Port|IP|Dir. match	40000	58999
Custom protocols	Custom protocol signature match	61000	61999
Trusted/validate ports	1st packet validate, 1st packet trusted match	62000	63999

Ordering basics:

- Layer 7 expression-based filters are located before their parent protocol signature (for example, expression matches on http are located before the HTTP protocol app-filter; the same applies to TLS, SIP, H323, RTSP, Citrix).
- HTTP host and URI are located before the HTTP referer for accounting accuracy (for example, YouTube® from within Facebook® is classified as YouTube®)
- App-filters combining protocol signatures with Layer 4 port, IP protocol, IP address, or flow direction are always located before the protocol signature only filter range.

HTTP

Protocol

HTTP is a client/server protocol using TCP/IP at the transport layer to deliver resources such as HTML files, images, videos and more.

HTTP 1.1 enables HTTP clients to use a persistent connection to a server allowing them to reuse the same TCP session for multiple HTTP transactions. Text, images, video, scripts and other objects can be downloaded individually in different transactions through the same TCP session.

Figure 2: HTTP persistent connection describes a typical persistent HTTP connection between a web client and a server with multiple HTTP transactions within the same TCP session:

Figure 2: HTTP persistent connection

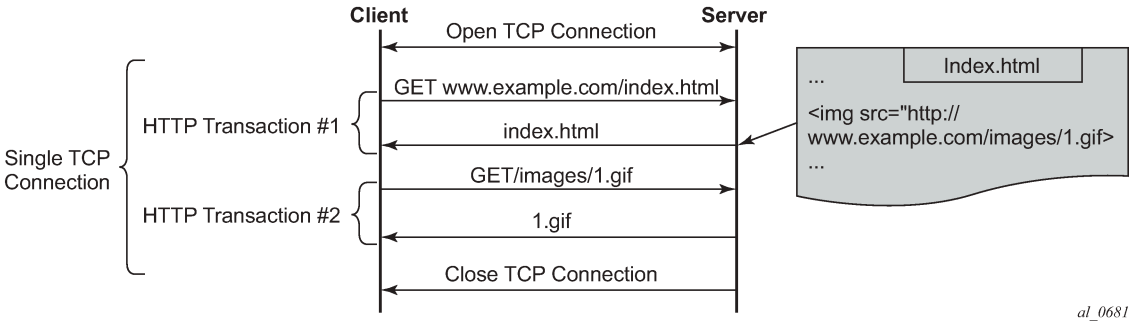

User-defined expression-based HTTP applications use the first HTTP transaction to classify the flow (optionally this behavior can be modified).

HTTP request

The Figure 3: HTTP request shows the content of a typical HTTP request to wikipedia.org which includes the following header fields: HTTP Host, HTTP URI, HTTP User Agent and HTTP referer fields:

Figure 3: HTTP request

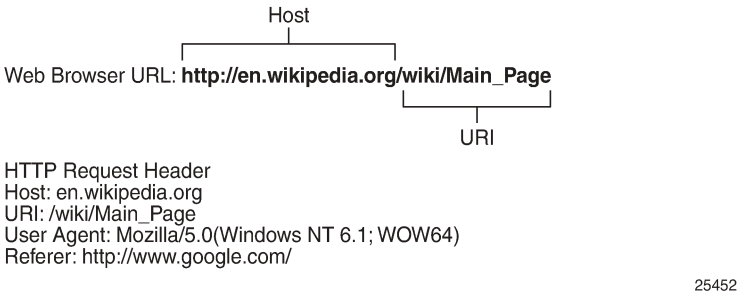

- HTTP Host — Represents the domain name (does not include "http://").
- HTTP URI — The URL trailer after the host domain name (begins with slash "/").
- HTTP Referer — The address of the previous web page from which a link to the currently requested page was followed (in this example, the referer is www.google.com which means the user clicked on a link from a Google search pointing to wikipedia.org).
- HTTP User Agent — This identifies the web browser or application making the HTTP request.

Configuration examples

HTTP host (Wikipedia)

Classifying HTTP traffic from this web site can be done using a single expression tail anchored on the HTTP host:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Wikipedia Web Access"
                            application "Wikipedia"
                            expression 1 {
                                type http-host
                                eq "*.wikipedia.org$"
                            }
                        }
```

This can be confirmed using Wireshark®.

Figure 4: Wireshark® www.wikipedia.org

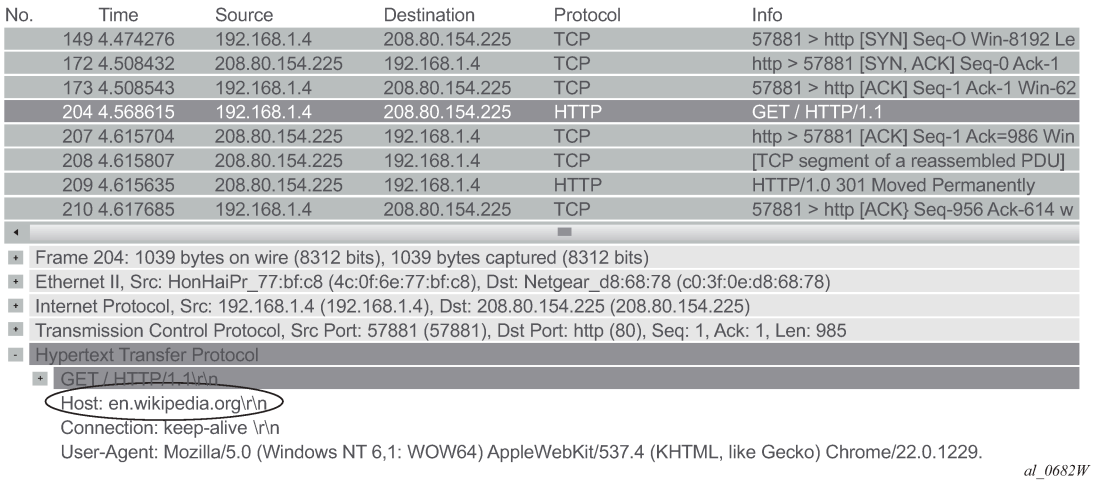

Classification per URI within the same host

Operators may need to apply different charging rules to different content located on the same HTTP domain (different URI, same HOST).

Table 4: Classification rules for the ISP ON-NET content services displays an example of classification rules for the ISP ON-NET content services:

Table 4: Classification rules for the ISP ON-NET content services

URL	Charging rule	AA application
www.ispdomain.com/video	Rule #1 – 0 Rating	ISP-Portal-Video
www.ispdomain.com/images	Rule #2 – Charge X	ISP-Portal-Images
www.ispdomain.com/*	Rule #3 – Charge Y	ISP-Portal-Default

HTTP 1.1 can reuse the same TCP connection for many transactions to the same server. Classifying each HTTP transaction to www.ispdomain.com independently requires a specific AA configuration.

SR OS allows to selectively enable **http-match-all-requests** in app-filters to improve the system performance and limit the HTTP analysis per domain.

The following configuration example allows traffic classification of different URIs of the same domain (www.ispdomain.com) independently, therefore allowing differentiated charging and control:

- **http-match-all-requests** is enabled on all host+uri app-filters to www.ispdomain.com
- default app-filter required to match any traffic to www.ispdomain.com

```
configure {
    application-assurance {
        group 1 {
```

```
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Zero rated content"
                            application "ISP Portal Video"
                            http-match-all-requests true
                            expression 1 {
                                type http-host
                                eq "^www.ispdomain.com$"
                            }
                            expression 2 {
                                type http-uri
                                eq "^/video*"
                            }
                        }
                        entry <1..65535> {
                            admin-state enable
                            description "Image charging"
                            application "ISP Portal Images"
                            http-match-all-requests true
                            expression 1 {
                                type http-host
                                eq "^www.ispdomain.com$"
                            }
                            expression 2 {
                                type http-uri
                                eq "^/images*"
                            }
                        }
                        entry <1..65535> {
                            admin-state enable
                            description "Default charging"
                            application "ISP Portal Default"
                            http-match-all-requests true
                            expression 1 {
                                type http-host
                                eq "^www.ispdomain.com$"
                            }
                        }
                    }
```

SSL/TLS (HTTPS)

Protocol

HTTPS uses SSL/TLS to encrypt traffic between the client and the server. Because this communication is encrypted, it is not possible to identify the HTTP Host or URI. However, AA can still identify the service requested by the subscriber by looking at the TLS certificate information or Server Name Indication exchanged in the clear before the TLS session is established.

Note:
SSL/TLS expression-based app-filters are not limited to HTTPS. HTTPS is not a protocol in itself, but it is HTTP traffic, tunneled encrypted into SSL/TLS on port 443.

SSL/TLS certificates

The following snapshot (Figure 5: Wireshark® HTTPS www.whatsapp.com) from Wireshark shows the SSL/TLS certificate exchanged using the mobile application *Whatsapp*®.

Figure 5: Wireshark® HTTPS www.whatsapp.com

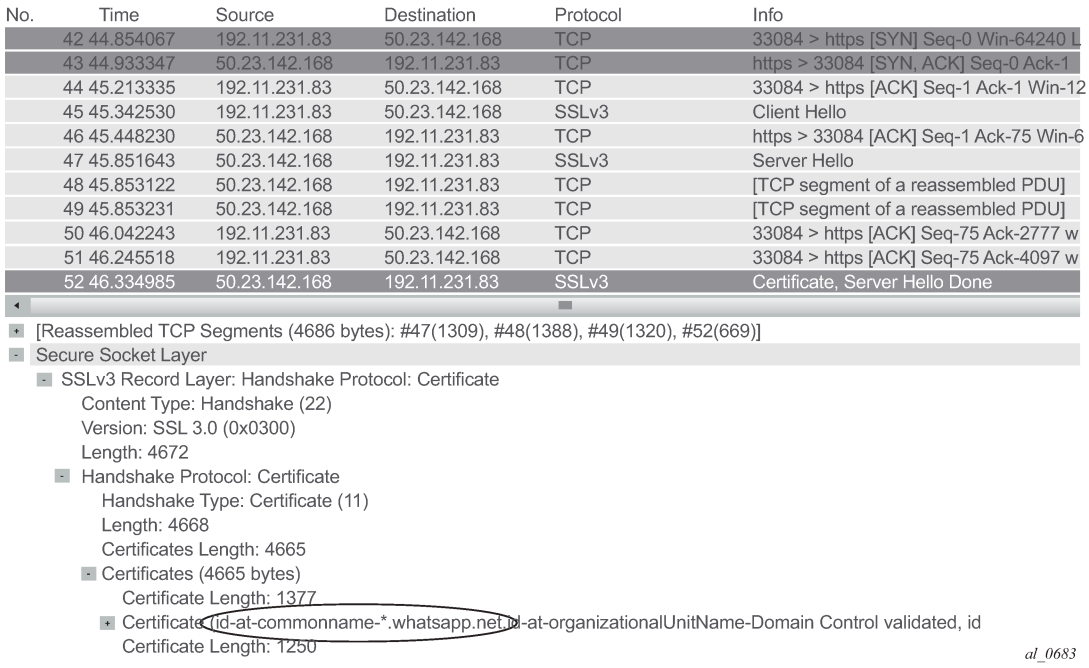

The certificate information can be found in the Server Hello message sent by the server, capturing SSL/TLS (HTTPS) traffic from this application can be done using a single app-filter entry tail anchored on the TLS Common Name Certificate:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Whats App tls and image/voice/video traffic"
                            application "Whats App"
                            expression 1 {
                                type tls-cert-subj-common-name
                                eq "*.whatsapp.net$"
                            }
                        }
```

Server name indication

SSL/TLS traffic can optionally be identified using the Server Name Indication (SNI) which is an extension to the TLS protocol.

The SNI is found in the TLS Client Hello, the http-host expression in the app-filter is reused to classify this traffic:

Figure 6: HTTPS SNI

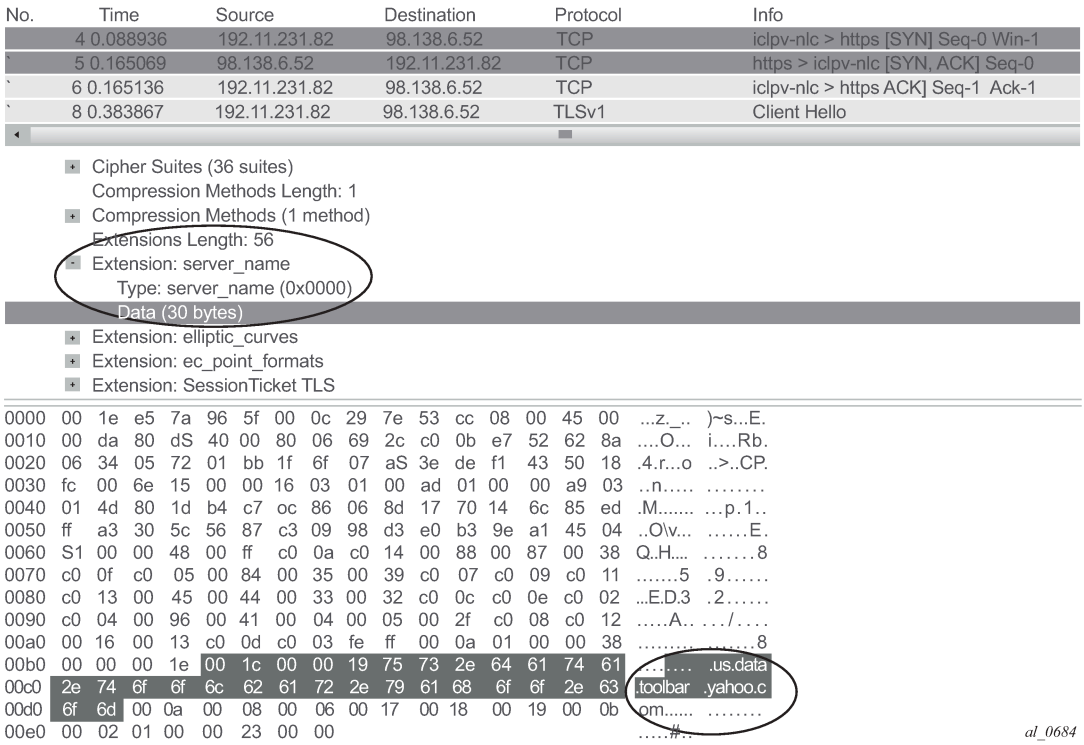


```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Yahoo HTTP or TLS SNI"
                            application "Yahoo"
                            expression 1 {
                                type http-host
                                eq "*.yahoo.com$"
                            }
                        }
                    }
```

SIP

Protocol

SIP is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors SIP control flows and associates RTP/RTCP media flows accordingly in the sip_rtp protocol signature.

The operator can use a SIP expression match criteria in app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports SIP expression match criteria on SIP URI, SIP user agent and SIP media type. The following snapshot from Wireshark® shows a SIP control exchange using the voice-video application Vonage® followed by the RTP media audio flow; the expression fields that can be matched using AA app-filters are highlighted:

Figure 7: SIP Wireshark® capture

Configuration example

The following configuration example provides the configuration to classify Vonage® SIP/RTP desktop traffic using SIP URI expression:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Vonage"
                            application "Vonage"
                            expression 1 {
                                type sip-uri
                                eq "*voncp.com*"
                            }
                        }
```

H323

Protocol

Similar to SIP, H323 is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors H323 control flows and associates the RTP media flow accordingly in the h323_rtp protocol signature.

The operator can use an H323 expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports H323 expression match criteria on the H323 product ID. The following snapshot from Wireshark shows an H323 control exchange using the Telepresence application LifeSize® followed by the RTP media audio flow; the expression field that can be matched using AA app-filters is highlighted:

Figure 8: H323 Wireshark® capture

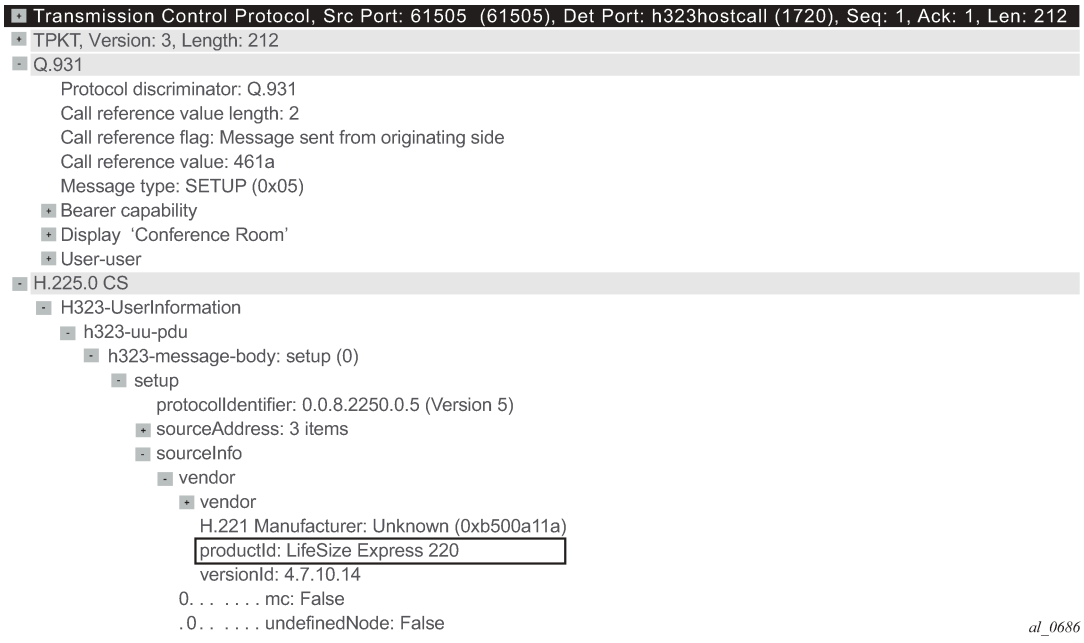

Configuration example

The following configuration example provides the configuration to classify LifeSize® H323/RTP traffic using the H323 product ID expression:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                app-filter {
                    entry <1..65535> {
                        admin-state enable
                        description "LifeSize H323 traffic"
                        application "LifeSize"
                        expression 1 {
                            type h323-product-id
                            eq "^LifeSize*"
                        }
                    }
```

RTSP

Protocol

RTSP is a signaling protocol used for controlling media streaming content such as audio and video over RTP/RDT. AA automatically monitors the RTSP control flows and associates its RTP/RDT media flow with the rtp_rtsp protocol signature.

The operator can use an RTSP expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful to identify specific streaming applications.

AA supports RTSP expression match criteria on the RTSP Host, URI, User Agent. The Figure 9: RTSP setup request shows an RTSP setup request to YouTube® followed by the RTP media audio flow; the expression fields that can be matched in RTSP SETUP request using AA app- filters are highlighted:

Figure 9: RTSP setup request

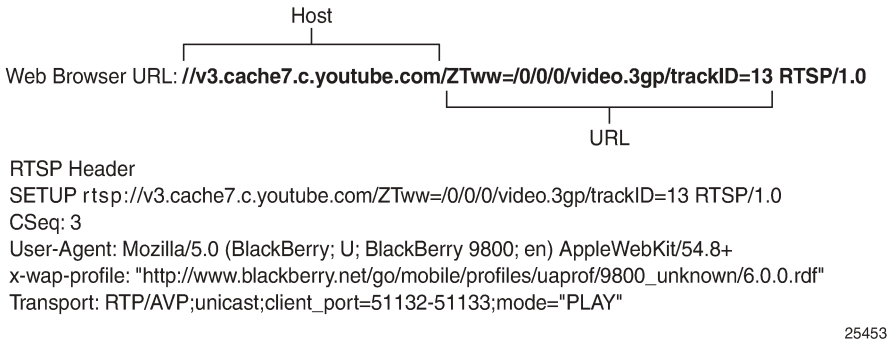

Configuration example

The following configuration example provides the configuration to classify YouTube® RTSP/RTP traffic using RTSP Host expression:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "YouTube RTSP/RTP Video"
                            application "YouTube"
                            expression 1 {
                                type rtsp-host
                                eq "*.youtube.com$"
                            }
                        }
```

Citrix

Protocol

Independent Computing Architecture (ICA) is a Citrix Systems® protocol used in Citrix's WinFrame, Citrix XenApp (formerly called MetaFrame/Presentation Server), and Citrix XenDesktop products.

Citrix makes it possible to run applications remotely on large servers, therefore making better use of server resources while at the same time allowing people using other platforms to use the applications, for example, run Microsoft® Word on a UNIX workstation.

Citrix_ica protocol signature detects any remote application using Citrix (the protocol needs to be unencrypted and configured to non-seamless). The Citrix ICA session is started from a client and can be anything from Remote Desktop, SAP to Microsoft® Word.

The Citrix expression match app-filter is used to classify traffic based on the Citrix-published application. This published application is configured on the server and in the preceding example, it can be for instance RDP, SAP, Word, XLS or Microsoft® Word depending how the server is configured.

Configuration example

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Citrix SAP Application"
                            application "Citrix SAP"
                            expression 1 {
                             type citrix-app
                             eq "^SAP$"
                        }
```

IP address and TCP/UDP port

Traffic from specific servers can be classified using IPv4/v6 server-address app-filter rules. It is used usually to identify traffic from an internal (on-net) server as opposed to an Internet (off-net) server.

The server-address app-filter automatically detects the client from the server by identifying which side opens the connection. It implicitly classifies traffic based on the server IP address or port number. For example, if A initiates a TCP connection to B, then flows from A to B and from B to A can be classified with a match on server-address B. Similarly, a flow initiated from B to A can be classified using a match on server-address A.

Server address

The following configuration example uses a server-address app-filter to classify traffic from server 10.1.1.1 in the application called Application-1:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Server #1 10.0.0.1"
                            application "Application-1"
                            server-address {
                                eq {
                                    ip-prefix 10.0.0.1/32
                                }
                            }
                        }
```

Server address and server port

The following configuration example uses server-address and server-port app-filters to classify traffic from server 10.0.0.2 on port 1234 in the application called Application-2. It is particularly useful when the same server is used to provide different services that need to be classified separately:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Server #2 10.0.0.2 port 1234 only"
                            application "Application-2"
                            server-address {
                                eq {
                                    ip-prefix 10.0.0.2/32
                                }
                            }
                            server-port {
                                eq {
                                    port-number 1234
                                }
                            }
                        }
```

Server port and protocol signature

It is possible to combine a protocol signature with a port number in the same app-filter, this is typically done in business VPNs for specific internal applications not detected using existing AA protocol signatures.

The following configuration example classifies a business VPN application running on TCP port 4000 and not detected by any other signatures. It combines the protocol signature unknown_tcp with the needed port number. This allows keeping the classification untouched for the rest of the protocols/applications and is the recommended approach:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "Business VPN Application X Port 4000"
                            application "Business VPN Application X"
                            protocol {
                                eq "unknown_tcp"
                            }
                            server-port {
                                eq {
                                    port-number 4000
                                }
                            }
                        }
```

Note:
It is important to follow the app-filter range recommendations for a correct classification of traffic using IP address or port number.

Flow setup direction

Traffic can be classified based on flow-setup-direction app-filter. The flow setup direction can be either subscriber-to-network or network-to-subscriber.

Network side and subscriber side is AA terminology related to where AA is enabled:

- In broadband and mobile networks, AA is enabled per subscriber. This means the subscriber side represents the ESM/mobile/transit subscriber while the network side represents Internet or other subscribers.
- In business VPNs, AA is enabled on a VPN SAP/spoke SDP and the subscriber side represents the local VPN site (SAP/spoke/transit).

The following example shows the configuration to classify http traffic hosted by AA subscribers (for example, broadband subscribers running a web server):

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
```

```
---snip---
                app-filter {
                    entry <1..65535> {
                        admin-state enable
                        description "HTTP Server on the subscriber side"
                        application "HTTP server"
                        flow-setup-direction network-to-subscriber
                        protocol {
                            eq "http"
                        }
                    }
```

IP protocol

Traffic can be classified using an IP protocol number for non TCP/UDP traffic.

The following example provides the configuration to classify ICMP IPv4/v6 traffic:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "ICMP v4"
                            application "ICMP"
                            ip-protocol {
                                eq icmp
                            }
                            protocol {
                                eq "non_tcp_udp"
                            }
                        }
                        entry <1..65535> {
                            admin-state enable
                            description " ICMP v6"
                            application "ICMP"
                            ip-protocol {
                                eq ipv6-icmp
                            }
                            protocol {
                                eq "non_tcp_udp"
                            }
                        }
```

Custom protocol

Custom protocols can be used to classify TCP/UDP applications using hexadecimal string matching (up to 16 hex octets) at a configurable payload offset in the data payload. The expression string length and offset must not exceed 128 bytes.

To illustrate this feature the Solaris® application GoGlobal is used. It provides remote access to a server (similar to VNC®). The following snapshot (Figure 7: SIP Wireshark® capture) from Wireshark® shows a TCP SYN/ACK session establishment followed by the first data exchange:

Figure 10: Wireshark® GoGlobal

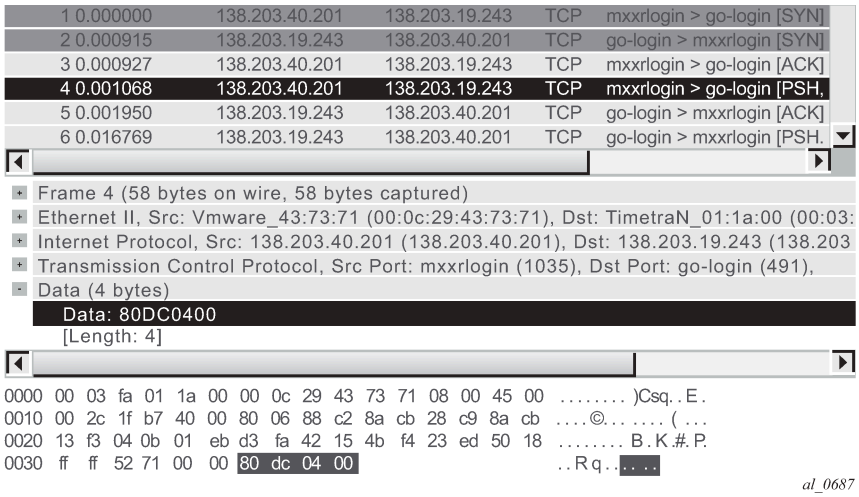

Wireshark® shows that each TCP session payload starts with 80DC0400 (no offset) after the three-way TCP handshake. As a result, the configuration required to classify this traffic is as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
---snip---
                    custom-protocol "custom_01" {
                        admin-state enable
                        description "goglobal tcp"
                        ip-protocol tcp
                        expression 1 {
                            direction client-to-server
                            eq "\x80\xdc\x04\x00"
                        }
                    }
---snip---
                     app-filter {
                        entry <1..65535> {
                            admin-state enable
                            description "GoGlobal "
                            application "GoGlobal"
                            protocol {
                                eq "custom_01"
                            }
                        }
```

Typical configuration mistakes

An operator creating new user-defined applications can make a few typical mistakes which are listed below:

- App-filters in admin-state disable — The default app-filter state is disabled. An **admin-state enable** command must be executed in order for it to be enabled.

- App-filters with no match criteria — This is a more troublesome mistake because it catches all the traffic entering the filter in a particular application.

Troubleshooting application identification

Show commands

Router/partition statistics

Partition level statistics are not updated in real time. Instead, statistics for a particular flow are updated either at flow closure or every five minutes. The five-minute sliding window interval is a common interval for all flows in an ISA MDA. Different ISA MDAs have a different five-minute windows as this interval is set at the MDA boot time.

The following command can be used to view the statistics for all applications configured in the ISA Group 1, Partition 1:

```
show application-assurance group 1:1 application count
```

Alternatively, it is possible to sort the display by octets, packets, flows:

```
show application-assurance group 1:1 application count top [octets | packets | flows] [max-
count <max-count>]
```

The operator can also identify which app-filters are being hit by the AA policy per partition (this command is not available per subscriber), it is particularly useful to identify which filters are used and optionally prune unnecessary app-filters from user-defined applications:

```
show application-assurance group 1:1 policy app-filter
```

Note:
The app-filter policy is usually relatively large, in which case additional 7750 SR CLI functionality can be used to filter out the output and only show the relevant information.

The following example was created for the application FTP:

```
[/]
A:admin@PE-1# show application-assurance group 1:1 policy app-filter
                                | match 'application \"FTP\"' pre-lines 3 post-lines 2
                    exit
                    entry 44030 create (2 flows, 1401 B)
                        protocol eq "ftp_control"
                        application "FTP"
                        no shutdown
                    exit
                    entry 44031 create (2 flows, 1205 B)
                        protocol eq "ftp_data"
                        application "FTP"
                        no shutdown
                    exit
```

Because partition level statistics are not updated in real time, it is recommended for troubleshooting purposes to use subscriber statistics or sub-study statistics.

Subscriber statistics

Subscriber-level statistics can be updated in real time. AA is usually configured by the operator to collect subscriber-level statistics for all application groups in residential and Wifi, while business VPNs typically collect Application group and all applications for each site with AA enabled.

The following commands can be used to view per subscriber statistics for all app-groups or applications configured in ISA group 1, partition 1 for the ESM subscriber "Bob" or business VPN SAP 1/1/10:10:

```
show application-assurance group 1:1 aa-sub esm "bob" app-group count
show application-assurance group 1:1 aa-sub sap 1/1/10:10 application count
```

In case only app-group statistics are collected per subscriber, the aa-sub-study feature can be used to collect per application-level statistics for selected subscribers, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                statistics {
                    aa-sub-study application {
                        aa-sub {
                            esm "bob" { }
                        }
```

When done, the system shows all application level statistics for this subscriber:

```
show application-assurance group 1:1 aa-sub-study esm "bob" application count
```

Similar to partition-level statistics, aa-sub and aa-sub-study statistics can be sorted by octets, packets, flows:

```
show application-assurance group 1:1 aa-sub-study esm "bob" application count top [octets |
 packets | flows] [max-count <max-count>]
```

Note:
When the number of flows per ISA card reaches a threshold then per subscriber statistics are not available in real time anymore and only the snapshot command can be used to display the statistics recorded in the previous five-minute interval window:

show application-assurance group 1:1 aa-sub-study esm "bob" snapshot application count

AppFilterMiss

The default policy configuration provides a failsafe application at the very end of the app-filter list to classify any remaining traffic in the AppFilterMiss application. There should never be any traffic in this application. This failsafe filter is used as a debug to make sure that there are no major issues in the configuration.

Note:
Traffic can typically be classified as AppFilterMiss when not all protocol signatures are mapped to a particular application. This could happen when upgrading to a new ISA software and enabling new protocol signature detection while not ensuring first that the correct application was provisioned. See the Release Note upgrade section for more details on AA signature upgrade.

Tools

Flow-record-search

Traditional show commands may not provide enough information when troubleshooting flow identification and the operator can use the ISA flow-record-search tool to dump the ISA flow table for more information. This feature comes with a large number of filtering options documented in the user guide.

Each flow gives visibility into: Flow ID, Sub-Type, Sub-Name, Initiator, Direction, Source IP, Dest. IP, IP Protocol, Source Port, Dest. Port, FC, DSCP, Classified, Protocol, Application, App- Group, Charging Group, Packets tx, Bytes Tx, Packets-discarded, Bytes-discarded and so on.

See below for the most commonly used commands.

The following command shows all the flows in an ISA card per ISA group:partition (can be a very long output, up to 3M entries):

```
tools dump application-assurance group 1:1 flow-record-search isa 1/2
```

The following command shows all the flows per AA subscriber in a group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob"
```

The following command shows all the active flows per AA subscriber in a group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob" flow- status
 active
```

The flow-record-search command is also available with additional details by adding search-type detail at the end of the command line. Because of the length of the output it is recommended to paste the CLI output content in a notepad file.

HTTP host recorder

AA cflowd allows operators to export the HTTP domain extracted from HTTP flows to the NSP cflowd collector. This allows the operator to understand which HTTP hosts are visible in the network.

However, in case a cflowd collector is not deployed, AA provides the HTTP host recorder tool command to record HTTP hosts seen by AA. See the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for more details.

The following debug command is configured in classic CLI:

```
debug
    application-assurance
        group 1:1
```

```
            http-host-recorder
                filter
                    default-filter-action record
                    record http-host-app-filter-candidates
                exit
                rate 100
                no shutdown
            exit
        exit
    exit
exit
```

```
tools dump application-assurance group 1:1 http-host-recorder top octets
```

Port recorder

This function is particularly useful in business VPN (it can also be used in residential networks). The port-recorder AA tool function is similar to the http-recorder. It allows the operator to record which ports are used on selected applications.

It is most commonly used with the applications Unidentified TCP and Unidentified UDP but it can be configured to record any other applications (the debug configuration is in classic CLI):

```
debug
    application-assurance
        group 1:1
            port-recorder
                application "Unidentified TCP"
                application "Unidentified UDP"
                rate 100
                shutdown
            exit
        exit
    exit
exit
```

```
tools dump application-assurance group 1:1 port-recorder top bytes
```

Conclusion

This example, which is intended for Application Assurance network architects and engineers, provides the information required to modify an existing AA policy following AA best practices and guidelines, and provides the necessary troubleshooting information to better understand application classification using Application Assurance.

Application Assurance — App-Profile, ASO and Control Policies

This chapter provides information about Application Assurance (AA) app-profile, Application Service Options (ASOs) and control policy configurations.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 12.0.R4, but the MD-CLI in the current edition is based on SR OS Release 25.7.R1.

Nokia recommends to use the AppDB prior to configuring traffic control policies. The AppDB is a default configuration file to define all of the applications of interest, including all of the relevant application-groups, applications and app-filters to classify traffic, and can be obtained through Nokia's support organization.

Overview

In addition to providing valuable traffic analysis and statistics information using the 7750 Service Router (SR) or 7450 Ethernet Service Switch (ESS) and Application Assurance (AA), one of the key objectives of the AA solution is to provide the tools to manage subscriber traffic at the application level. Examples of traffic management actions include:

- Throttling low priority bandwidth hungry applications during peak hours.
- Prioritizing and remarking selected applications.
- Implementing a walled-garden environment providing open access to selected free web services only, redirecting all other requests from unregistered subscribers to a registration portal with payment services.
- Enrich HTTP header with subscriber identification parameters to offer subscribers transparent access to premium content.
- In-browser notification which triggers the display of administrative, informational or promotional messages in selected browser-sessions.
- Stateful session filtering with Application Level Gateway (ALG) support to protect subscribers against unsolicited flows.
- Parental control services interworking with an external Internet Content Adaptation Protocol (ICAP) server for rating the requested web sites.

Application traffic control policies can be applied as global policies for all subscribers, or they can be activated for individual subscribers or groups of subscribers.

This chapter describes the basics of activating Application Assurance on a subscriber through the use of App-Profile and demonstrates the use of static or dynamic traffic control policies using Application Service Options (ASOs) and Application QoS Policies (AQP). It also provides detailed information for configuring Bandwidth, Flow-Count and Flow-Rate Policing including Time of Day (ToD) policing. Other policy control actions can be found in the Advanced Configuration Guide or in the 7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide.

Configuration

Activation of AA services

App-profile

Application profiles (app-profile) enable application assurance services for an Enhanced Subscriber Management (ESM), Distributed Subscriber Management (DSM), or transit subscriber, or for a SAP or spoke SDP which are commonly referred to as AA subscribers (**aa-sub**). Each app-profile is unique in the system and defines the services that the AA subscriber receives.

Assigning an app-profile to an ESM subscriber affects every host of that subscriber. Similarly, applying an app-profile to a SAP or spoke SDP affects all traffic within that SAP or spoke SDP.

App-profiles are defined at the AA group partition level (in case of a partitioned ISA-AA group) as in the following configuration example:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-profile "1-1/15M" {
                        description "App-Profile Description"
                        capacity-cost 15
                        divert true
                        characteristic "Parental Control" {
                            value "enabled"
                        }
                    }
```

The app-profile parameters are:

- **divert true** — Diverts all traffic from and to this subscriber to an ISA-AA. Configuring **divert false** effectively disables all AA services for subscribers using this app-profile.

  Default value: **divert false**.

- **characteristic [**<*characteristic-name*> **value** <*value-name*>**]** — one or more optional ASO service characteristics can be used to apply an AA control policy to the subscriber.
- **capacity-cost** <*cost*> — An application profile capacity cost is used to load balance AA subscribers across multiple ISA-AA cards. A common practice is to define a cost proportional to the expected peak

BW for the subscribers using this profile (in kbps or Mbps). The capacity cost is out of the scope of this example. The range is 1 to 65535, default 1.

This app-profile example uses the following naming convention:

*<group-id>-<partition-id>/<BW>*M where

- *<group-id>* — The ISA-AA group ID on which this profile is created.
- *<partition-id>* — The AA partition ID on which this profile is created.
- *<BW>* — Defines the maximum bandwidth used by the subscriber, which is used for aa-subscriber cost load balancing and subscriber rate limiting. The *M* stands for Mbps.

In general the operator can choose to use either ASO characteristics override or multiple app-profiles to apply different AA QoS policies to ESM subscribers or business VPN sites. For flexibility and scale, it is recommended to use ASO overrides whenever possible. This is described in more detail in the following sections.

Note:
Prior to using special characters in a policy object name the operator should verify the list of special characters supported by the 5620 SAM; for instance the 5620 SAM does not support the use of ":" in the app-profile name therefore it should be avoided.

Residential and Wi-Fi services

The app-profile can be assigned or modified for ESM, DSM, or Transit IP subscribers either at subscriber creation time or while the subscriber is in service:

- Subscriber creation — An app-profile can be assigned at subscriber creation time through RADIUS, DHCP Option 82, Local User Database, static configuration or through a default app-profile.
- In service app-profile modification — An app-profile can be dynamically modified in service through a RADIUS Change of Authorization (CoA). An app-profile can also be dynamically modified in service through Gx.

In case no app-profile is returned at subscriber creation by RADIUS, LUDB, or DHCP, or when no static configuration is present, the system can apply a default app-profile if configured within the subscriber group-interface (or MSAP policy) sub-sla-mgmt:

```
            sub-sla-mgmt {
                defaults {
                    app-profile "1-1/15M"
                }
            }
```

Business VPN and other service interfaces

App-profiles are statically assigned to a SAP, spoke SDP, or transit prefix VPN site via the 5620 SAM or CLI.

The following configuration shows how to enable application assurance on a SAP or spoke SDP in a business VPRN service:

```
configure {
    service {
        vprn "business-VPRN-100" {
```

```
            admin-state enable
            description "L3 Service Customer 1"
            service-id 100
            customer "1"
            interface "to-site1" {
                ipv4 {
                    primary {
                        address 192.168.1.1
                        prefix-length 24
                    }
                }
                sap 1/1/10:11 {
                    app-profile "1-1/15M"
                }
            }
            interface "to-site2" {
                ipv4 {
                    primary {
                        address 192.168.2.1
                        prefix-length 24
                    }
                }
                spoke-sdp 12:100 {
                    app-profile "1-1/15M"
                }
            }
```

Defining application service options

ASOs for traffic control - introduction

To determine which application control policies need to be applied to an AA-subscriber, an app-profile with a number of service characteristics (ASOs) is associated with each subscriber. These service characteristics are then used as match criteria in AQP policy rules to determine which rules to apply.

Therefore ASOs are service characteristics assigned to a subscriber and are used to identify the traffic control policy rule (AQP) applicable to a subscriber or a group of subscribers.

Most policy rules are applicable to multiple subscriber profiles; nevertheless it is possible that a specific subscriber requires a dedicated policy.

ASO characteristics and values

For each service option that can be used by one or more subscribers, an ASO characteristic is defined with a number of values that represent all available choices for that service characteristic. The names and values of the ASO characteristics are configurable string values; best practice is to use strings that provide a meaningful description of the service characteristic they represent.

Each ASO characteristic requires a default value and each app-profile inherits the default value of all the ASO characteristics created in a partition unless a characteristic is referenced directly in the app-profile or overwritten as described below.

ASOs are defined at the AA group partition level (in case of a partitioned ISA-AA group). In the following configuration example two different ASO characteristics are defined: "P2P-Sub-DL" and "Parental Control":

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-service-options {
                        characteristic "P2P-Sub-DL" {
                            default-value "unlimited"
                            value "1M" { }
                            value "500k" { }
                            value "unlimited" { }
                        }
                        characteristic "Parental Control" {
                            default-value "disabled"
                            value "disabled" { }
                            value "enabled" { }
                        }
```

The ASO values and default value of a characteristic can be displayed using the following **show** command:

```
[/]
A:admin@BNG# show application-assurance group 1:1 policy app-service-option "P2P-Sub-DL"
===============================================================================
Application-Assurance Application Service Options
===============================================================================
Characteristic "P2P-Sub-DL"
Value                           Default
-------------------------------------------------------------------------------
1M                              No
500k                            No
unlimited                       Yes
===============================================================================
```

When configuring service characteristics for optional service options, it is recommended to configure a default value which does not trigger any AQP policy action (the default value does not match any AQP match criteria) such that the behavior of existing subscribers and app-profiles does not change until the operator specifically configures or signals a non-default characteristic value for the subscriber or the app-profile. In the preceding example, "Parental Control disabled" and "P2P-Sub-DL unlimited" have no corresponding AQP by design; therefore if these particular service options were applied to a subscriber they would not match a QoS policy entry.

Service options for AA subscribers

ASO assignment in app-profile

ASOs can be statically assigned in the app-profile; this type of ASO characteristic assignment is typically reserved to the default service options enabled on a large number of subscribers.

Figure 11: Service tier example using ASO, app-profile and AQP shows an example of AA service definition (ASO and app-profile) for a Gold and Bronze service tier definition where two app-profiles "Gold" and "Bronze" are defined with the following characteristics:

- "Gold" app-profile — No specific policy actions or ASO characteristics are configured statically in the app-profile.
- "Bronze" app-profile — A specific ASO characteristic and value is assigned to the profile to limit peer to peer download traffic to 1 Mbps (this example does not show the app-qos-policy nor policer configuration, this will be described later).

Figure 11: Service tier example using ASO, app-profile and AQP

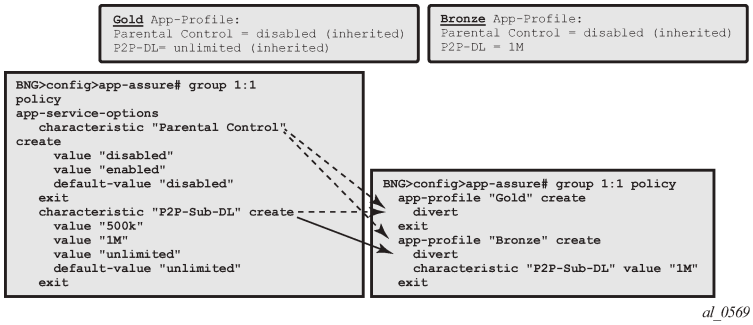

Each app-profile inherits the default values of all the ASO characteristics defined in a AA group-partition; in the preceding example, this is the reason why the app-profile "Gold" inherits "Parental control disabled" and "P2P-Sub-DL unlimited". The app-profile "Bronze" inherits "Parental control disabled" while "P2P-Sub-DL 1M" is assigned to this profile statically.

The operator can identify per app-profile which characteristics values are inherited from their default value and which are statically assigned using the following **show** commands:

```
[/]
A:admin@BNG# show application-assurance group 1:1 policy app-profile "Gold"
                app-profile "Gold" create
                    divert
                    characteristic "P2P-Sub-DL" inherits default-value "unlimited"
                    characteristic "Parental Control" inherits default-value "disabled"
                exit

[/]
A:admin@BNG# show application-assurance group 1:1 policy app-profile "Bronze"
                app-profile "Bronze" create
                    divert
                    characteristic "P2P-Sub-DL" value "1M"
                    characteristic "Parental Control" inherits default-value "disabled"
                exit
```

Note:
Using ASO overrides, described later, it is possible to implement the same choice of AA service options using a single app-profile.

ASO overrides per subscriber via RADIUS or Gx

Prior to SR OS 12.0.R1 the operator can assign (and modify: CoA) the app-profile per ESM or Transit-IP subscribers using the "Alc-App-Prof-Str" [26-6527-45] RADIUS attribute.

SR OS 12.0.R1 added support for ASO characteristic overrides for ESM and Transit-IP subscribers via RADIUS using the attribute "Alc-AA-App-Service-Options" [26-6527-193]. This attribute can be returned

during the subscriber creation process or while the subscriber is in service through RADIUS CoA. Refer to 7450 ESS, 7750 SR, and VSR RADIUS Attributes Reference Guide for more details related to the use of the AA RADIUS attributes.

An example of a RADIUS CoA message returned to the system to modify both the app-profile and one ASO characteristic is as follows:

```
    NAS-Port-Id = "1/1/10:4088"
    Framed-IP-Address = 192.168.211.30
    Alc-App-Prof-Str = "1-1/15M"
    Alc-AA-App-Service-Options = "P2P-Sub-DL=1M"
```

The ASO characteristics and values assigned to a subscriber (statically via app-profile or overridden) can be displayed using the following show command:

```
[/]
A:admin@BNG#  show application-assurance group 1:1 aa-sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/15M
App-Profile divert      : Yes
Capacity cost           : 15
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
---snip---

-------------------------------------------------------------------------------
Traffic                        Octets              Packets                Flows
-------------------------------------------------------------------------------
---snip---

-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                  Value                          Derived from
-------------------------------------------------------------------------------
P2P-Sub-DL                      1M                             dyn-override
Parental Control                disabled                       default
===============================================================================
```

In the preceding show command output, the **derived from** field describes how the characteristics and values are assigned to the subscriber:

- app-profile — The characteristic's value statically configured in the app-profile.
- dyn-override — The characteristic's value received from RADIUS or Gx.
- default — The characteristic's default value inherited (not statically configured in the app-profile nor dynamically modified).

SR OS also supports the signaling the app-profile or ASO characteristics override via Gx, see the Application Assurance – Usage Monitoring and Policy Control via Diameter Gx Protocol chapter for more details.

ASO overrides for business VPN and other services

ASO characteristic override values can be statically assigned to business VPN SAP, spoke SDP, and transit prefix subscribers.

The operator can provision the AA policy override parameters, multiple characteristics overrides per AA-sub can be defined per override policy, as in the following configuration example:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy-override {
                    aa-sub {
                        sap 1/1/10:11 {
                            characteristic "P2P-Sub-DL" {
                                value "1M"
                            }
                            characteristic "Parental Control" {
                                value "enabled"
                            }
                        }
                    }
```

Application control policies

App-QoS policy (AQP)

App-profile / ASO / AQP workflow summary

App-profiles enable application assurance services for an AA-subscriber. Each app-profile is unique in the system and defines the service that the AA subscriber receives.

To determine which control policies need to be applied to an AA-subscriber, a number of service characteristics (ASO) are associated with each AA-subscriber.

As described earlier, these service characteristics can either be configured directly within the app-profile or assigned using overrides and they are then used as match criteria in AQP policy rules to determine which application policy rules to apply.

The app-qos-policy (AQP) is an ordered list of entries defining policy actions for flows diverted to Application Assurance. Each AQP entry is composed of match criteria and actions.

Flows are evaluated against all entries of the AA QoS policy defined in the AA group partition that the subscriber app-profile belongs to (in case of a partitioned AA group).

Figure 12: App-Profile, ASO, AQP workflow summary provides a configuration example summary with app-profile, ASO, AQP, and policers:

Figure 12: App-Profile, ASO, AQP workflow summary

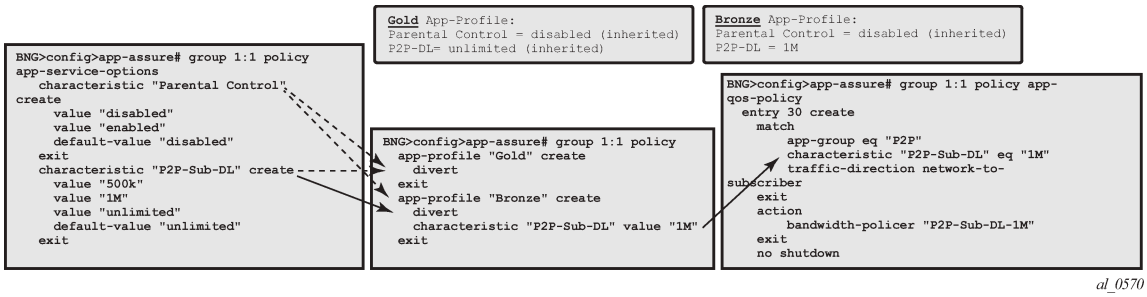

Match and action criteria

AQP match criteria

Multiple match criteria can be specified per AQP entry in which case the action only applies to flows that match all criteria. The most common match criteria are: characteristic, application, app-group, and charging-group.

The following AA match criteria can be used in an AQP:

- **app-group**
  - **{eq | neq }**<*app-group-name*>
- **application**
  - **{eq | neq }**<*app name*>
- **charging-group {eq | neq}**
  - **{eq | neq }**<*charging-group-name*>
- **traffic-direction {subscriber-to-network | network-to-subscriber | both}**
- **characteristic** <*characteristic-name*>: up to 4 characteristics and values per AQP
  - **{eq | neq }**<*value-name*>
- **ip-protocol**
  - **{eq | neq }**<*protocol-id* | *protocol-name*>
- **src-ip**
  - **{eq | neq }**
    - **ip-prefix**<*ipv4-address-and-prefix* | *ipv6-address-and-prefix*>
    - or **ip-prefix-list**<*ip-prefix-list-name*>
- **dst-ip**
  - **{eq | neq }**
    - **ip-prefix**<*ipv4-address-and-prefix* | *ipv6-address-and-prefix*>
    - or **ip-prefix-list**<*ip-prefix-list-name*>
- **src-port**

  - **{eq | neq }**
    - **port-number**<*0..65535*>
    - or **port-list**<*port-list-name*>
    - or **range**
      - **start**<*0..65535*>
      - **end**<*0..65535*>
- **dst-port**
  - **{eq | neq }**
    - **port-number**<*0..65535*>
    - or **port-list**<*port-list-name*>
    - or **range**
      - **start**<*0..65535*>
      - **end**<*0..65535*>
- **dscp**
  - **{eq | neq }**<*dscp-name*>
- **aa-sub <sub-type {eq | neq}>**
  - esm
    - **{eq | neq }**<*esm-subscriber-name*>
  - esm-mac
    - **{eq | neq }**<*string*>
  - sap
    - **{eq | neq }**<*sap-id*>
  - spoke-sdp
    - **{eq | neq }**<*spoke-sdp-id*>
  - transit
    - **{eq | neq }**<*transit-aa-subscriber-name*>

AQP actions

The following AA traffic control policies can be specified in an AQP:

- **drop <boolean>**
- **bandwidth-policer**
  - <*single-bucket* | *dual-bucket* | *anl* | *flow*> <*bw-policer-name*>
- **flow-count-limit-policer**
  - **policer-name**<*policer-name*>
  - **event-log**<*event-log-name*>
- **flow-setup-rate-policer**
  - **policer-name**<*policer-name*>

  - **event-log**<*event-log-name*>
- **remark**
  - **dscp**
    - **in-profile**<*dscp-name*>
    - **out-profile**<*dscp-name*>
  - **fc** <*fc-name*>
  - **priority** <*low* | *high* | *auto*>
- **http-error-redirect** <*redirect-name*>
- **http-redirect** — Redirect traffic to a landing page
  - **flow-type**<*dropped-flows* | *admitted-flows*>
  - **name**<*http-redirect-name*>
- **mirror-source**
  - **mirror-service**<*mirror-service-name*>
  - **all-inclusive**<*boolean*>
- **session-filter** <*session-filter-name*> — Session filter firewall
- **url-filter** : category-based URL filtering using ICAP
  - **characteristic**<*characteristic-name*>
  - **name**<*url-filter-name*>
- **http-notification** <*http-notification-name*>
- Additional drop actions:
  - **error-drop**: configure a drop action for packets cut-through due to IP packet errors (bad IP checksums, tcp/udp port 0, and so on.)
    - **event-log**<*event-log-name*>
  - **overload-drop**: configure a drop action for packets cut-through due to overload
    - **event-log**<*event-log-name*>
  - **fragment-drop**: configure a drop action for IP fragmented packets
    - **drop-scope**<*all* | *out-of-order*>
    - **event-log**<*event-log-name*>

Default versus application-specific AQP policies

Application QoS policy

It usually requires the examination of a few packets to identify the protocol/application of a flow. When AQP entries are defined to match on IP header criteria (IP address, IP prefix list, TCP/UDP port number, IP protocol, DSCP) or application criteria (application, app-group or charging group), the AQP action is only applied to matching application flows after a flow has been classified as an application.

Default QoS policy

If the AQP entry does not include match criteria against application (application, app-group, and charging-group) or IP header information (IP address, IP prefix list, TCP/UDP port number, IP protocol, DSCP) then the AQP policy is applied to all matching flows starting with the first packet of a flow before protocol and application identification is complete. Such AQPs are called default subscriber policies.

For an AQP to be qualified as a default subscriber policy, the match criteria must be limited to any combination of ASO characteristic values, traffic direction, and optional AA subscriber name.

AQP match and actions for the default QoS policy and application QoS policy are summarized in Table 5: Default QoS policy, application QoS policy table :

Table 5: Default QoS policy, application QoS policy table

Policy	AQP match	AQP action
Default QoS	ASO characteristic/values traffic direction aa-sub	Remark FC, DSCP, priority Bandwidth, flow-count, flow-rate policing Session-filter Url-filter Mirror error-drop, overload-drop, fragment-drop Drop
Application QoS	ASO characteristic/values traffic direction aa-sub application app-group charging-group IP address, IP prefix list TCP/UDP port number DSCP IP protocol number	Remark FC, DSCP, priority Bandwidth, flow-count, flow-rate policing HTTP notification HTTP redirect HTTP enrichment mirror Drop

To ensure fair access to the ISA-AA bandwidth and flow resources, it is recommended to configure default AQP policy entries limiting bandwidth and flow resources per AA sub.

Figure 13: Default downstream bandwidth policing shows a default subscriber policy limiting the downstream bandwidth (network-to-subscriber direction) to 25Mbps per subscriber:

Figure 13: Default downstream bandwidth policing

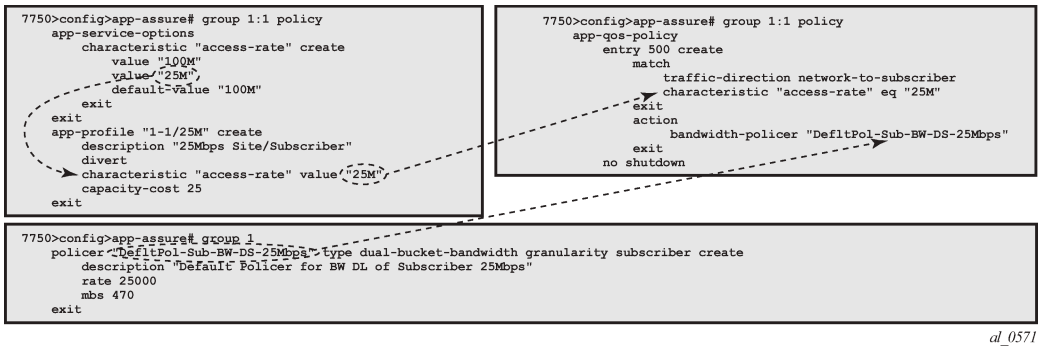

Implicit default subscriber policy

Session-filter, url-filter, overload-drop, fragment-drop, and error-drop can only be used as part of a default subscriber policy; therefore these actions are not compatible with application or IP header match criteria within the same AQP.

AQP entries evaluation

Multiple AQP match entries per flow

A single flow can match multiple AQP entries, in which case multiple actions can be selected based on the order of the AQP entries order (the lowest number entry has the highest priority); the drop action takes precedence over any other AQP entry. The maximum numbers of actions that can be applied on a single flow are:

- 1 drop action
- Any combination of (applied only if no drop action is selected)
  - Up to 1 mirror action
  - Up to 1 FC, 1 priority, and 1 DSCP remark action
  - Up to 4 BW policers (1 single rate AA-sub, 1 dual rate AA-sub, 2 single rate system level)
  - Up to 12 flow policers (3 subscriber flow-count, 3 subscriber flow-rate, 3 system flow-count, 3 system flow-rate)
  - Up to 1 HTTP redirect
  - Up to 1 HTTP error redirect
  - Up to 1 HTTP enrichment
  - Up to 1 URL-filter
  - Up to 1 HTTP-notification
  - Up to 1 session-filter firewall
- 1 error drop

- 1 overload drop
- 1 fragment drop

An AQP entry match that would cause the preceding limits to be exceeded is ignored (no actions from that rule are selected) and the conflict counter for this AQP is incremented.

The operator can display hits and potential conflicts per AQP entry using the following show command:

```
[/]
A:admin@BNG# show application-assurance group 1:1 policy app-qos-policy

===============================================================================
Application QOS Policy Table
===============================================================================
Entry         Admin State                    Flow Hits          Flow Conflicts
-------------------------------------------------------------------------------
30            in-service                             0                      0
-------------------------------------------------------------------------------
No. of AQP entries: 1
===============================================================================
```

AQP evaluation

Flows are evaluated against all entries of the AA QoS policy at different steps during the lifetime of the flow:

- Flow creation — The default subscriber policy AQP entries for matching flows are applied starting with the first packet of a flow so before application identification completes.
- Application identification completion— The application QoS policies are applied when flow identification has been completed.

  Note:
  The default QoS policy entries are applied to the subscriber's flows for packets received before and after application identification is completed.

- Policy change — When a configuration change to the AA policy is committed, all diverted flows for subscribers using this policy partition are evaluated again against all AQP entries. This re-evaluation happens as a paced background task; therefore AQP control changes may not be applied immediately to all existing flows.

Policing

Policers

AA policer templates are configured as part of the AA group configuration by specifying the policer name, type, and granularity. Policers are unidirectional by definition so that separate policers must be defined per flow direction if the traffic needs to be policed in both directions and a separate AQP for each flow direction is therefore required.

The operator can configure the following types of policers:

- Bandwidth policers

  - Single bucket system level
  - Single bucket flow level
  - Single bucket AA subscriber level
  - Dual bucket AA subscriber level
- Flow count limit policer: system or AA subscriber level
- Flow setup rate policer: system or AA subscriber level

Subscriber level policers are instantiated per AA subscriber, meaning:

- The system automatically uses a dedicated policer for every single subscriber, even when multiple subscribers match the same AQP entry.
- The same policer can be referenced in different AQP entries; in this case all subscribers' flows matching any of these AQP entries are policed by the same subscriber policer. Example: if the same subscriber level policer "1Mbps" is referenced in AQP entry 100 matching application BitTorrent® and in AQP entry 110 matching application EDonkey®, then the sum of both the BitTorrent® and EDonkey® traffic cannot exceed 1 Mbps.

System level policers on the other hand are shared by all AA subscribers matching an AQP entry. These policers are typically used in residential and Wi-Fi service deployments to limit the total bandwidth for an application or application group, for all subscribers or for a group of subscribers on the system or partition. An example would be a system level 500 Mbps policer to limit the aggregated downstream bandwidth of peer-to-peer applications for all subscribers with a "Bronze" app-profile to 500 Mbps.

Note:
In case multiple ISA-AA cards are used per system, the overall maximum throughput using a system level policer is equal to the policer rate limit times the number of ISA cards in the system.

Bandwidth policing

Single bucket subscriber/system bandwidth policer

Single bucket policers police the matching traffic against a configured peak information rate (PIR). Traffic exceeding the PIR can be marked as out of profile or dropped.

The configuration template for a single rate bandwidth policer is as follows:

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer <policer-name> {
                    description <string>
                    granularity {subscriber|system}
                    mbs <maximum burst size>
                    pir <pir rate>
                    adaptation-rule pir {max|min|closest}
                    time-of-day-override-override <time-of-day-override-id> {
                    }
                    action {priority-mark|permit-deny}
```

where:

- **action** — defines the action that must be taken by the policer for non-conforming traffic.
  - **priority-mark** — non-conforming traffic are marked as out of profile (increasing the chances that non-conforming packets will be discarded in case of congestion on the egress queues).
  - **permit-deny** — non-conforming packets are dropped.
- **rate** — peak information rate in kbps
- **mbs** — maximum burst size in kilobytes
- **adaptation-rule pir <max|min|closest>** — The policers work at discrete operational rates supported by the hardware. The adaptation rule specifies how the actual operational policer rate (supported by the hardware) must be selected as compared to the configured PIR. During operation, both the operational and configured rate can be displayed using the operational **show application-assurance group** *<n>* **policer** *<policer-name>* **detail** command.
- **tod-override** — defines a time of day override policy applicable to a policer; this is described in more detail at the end of the policing section.

The following shows a single bucket subscriber level policer configuration example:

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "P2P-Sub-DL-1M" {
                    description "limit P2P downstream traffic per subs to 1Mbps"
                    granularity subscriber
                    mbs 19          # max. burst size in kilobytes
                    pir 1000        # PIR rate in kbps
                }
```

The following shows a single bucket system level policer configuration example:

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "P2P-Sys-DL-100M" {
                    granularity system
                    mbs 1875
                    pir 100000
                }
```

Dual bucket subscriber bandwidth policer

Dual-bucket policers police the matching traffic against a configured peak information rate (PIR) and committed information rate (CIR). Traffic below CIR is marked in profile; traffic between CIR and PIR is marked as out of profile. Traffic above the PIR is dropped.

Dual-bucket policers can only be used as subscriber policers; system policers cannot be defined as dual-bucket policers.

The configuration is similar to the single-bucket policer, but adds the configuration of a CIR and a committed burst size (CBS), and the action cannot be configured:

```
configure {
    application-assurance {
```

```
        group 1 {
            policer {
                dual-bucket-bandwidth-policer <policer-name> {
                    description <string>
                    mbs <max-burst-size-in-kilobytes>
                    pir <pir-rate-in-kbps>
                    cbs <committed-burst-size-in-kilobytes>
                    cir <cir-rate-in-kbps>
                    adaptation-rule {
                        pir {max|min|closest}
                        cir {max|min|closest}
                    }
```

The following shows a dual-bucket subscriber level policer configuration example:

```
configure {
    application-assurance {
        group 1 {
            policer {
                dual-bucket-bandwidth-policer "P2P-Sub-DL-2M-DB" {
                    mbs 38
                    pir 2000
                    cbs 19
                    cir 1000
                }
```

MBS/CBS calculation for bandwidth policers

The default MBS/CBS value of a bandwidth policer is set to 0, but the operator must modify this value to allow correct interworking with TCP-based applications. The network operator must carefully consider the values to be used in production networks based on the applications in the network and several other factors such as traffic patterns, traffic volume, bursts, and so on. Nokia recommends to configure a minimum MBS for a bandwidth policer, as follows:

MBS = 2 * MTU or 0.00025 * peak rate (whichever is larger)

The formula to calculate the MBS or CBS buffer size, as documented in RFC 6349 is:

Buffer (B) = RTT (s) * rate (bps) / 8

For Internet applications, Nokia recommends to use a common Round Trip Time (RTT) of 150 msec.

An example using a single bucket subscriber level policer rate of 1000 kbps:

MBS (B) = 1,000,000 / 8 * 0.150 = 18750 Bytes or 19 KB

Note: These policer values may need further adjustment depending on the application.

Flow setup rate policer

The default MBS/CBS value of a flow setup rate policer is set to 0, but Nokia recommends to configure a non-zero MBS value for a flow setup rate policer.

Flow setup rate policers police the maximum number of new flows that are accepted per second for matching traffic. The configuration is similar to the single-bucket bandwidth policer, with the rate and MBS now expressed in flows/sec and flows, respectively.

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-setup-rate-policer <policer-name> {
                    description <string>
                    granularity {subscriber|system}
                    peak-flow-setup-rate <setup-rate-in-flows/sec>
                    flow-setup-rate-burst-size <max-burst-size-in-flows>
                    adaptation-rule pir  {max|min|closest}
                    action {permit-deny|priority-mark}
                    time-of-day-override {
                    }
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of flow/seconds allocated per AA subscriber.

Note: In case the policer is used as part of the default AA subscriber policy then the **priority-mark** action has the effect to cut-through non conformant traffic in the ISA instead of drop using **permit-deny**.

Flow count limit policer

Flow count limit policers police the maximum number of concurrent flows for matching traffic:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer <policer-name> {
                    description <string>
                    granularity {subscriber|system}
                    action permit-deny|priority-mark
                    peak-flow-count <max-number-of-flows>
                    time-of-day-override {
                    }
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of concurrent flows allocated per AA subscriber.

Note:
The **priority-mark** has the effect to cut-through non conformant traffic in the ISA instead of drop using **permit-deny**.

Time of day policing

For time-of-day (ToD) policer override, up to eight override rates with time of day specifications can be defined per policer, this time of day override using the system local time.

ToD overrides are supported for all policer types described in the previous section (bandwidth, flow count, and setup rate) and can be configured using either daily or weekly patterns.

The configuration of ToD override on daily or weekly basis is shown in the following template:

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "P2P-Sub-DL-1M-TOD" {
                    granularity subscriber
                    mbs <max-burst-size-in-kilobytes>
                    pir <max rate in kbps>
                    time-of-day-override <override-id> {
                        admin-state enable
                        description <string>
                        mbs <override-mbs-in-kilobytes>
                        pir <override-max-rate-in-kbps>
                        time-range {    # either daily or weekly
                            daily {
                                start <daily start time, e.g. "17:00">
                                end <daily end time, e.g. "22:00">
                                on <day of the week> or on [<day1> ...]
                            }
                            weekly {
                                start {
                                    day <weekly start day, e.g. friday>
                                    time <weekly start time, e.g. "17:45">
                                }
                                end {
                                    day <weekly end day, e.g. saturday>
                                    time <weekly start time, e.g. "00:30">
                                }
                            }
                        }
                    }
```

where:

- **tod-override** <*override-id*> — up to 8 override IDs (with value 1-255) can be configured per policer.
- **time-range** — can be configured to be triggered.
  - On a daily basis at the indicated start/end-time on the specified days.
  - On a weekly basis at the indicated start day and time and end day and time.
  - Times can be indicated as <hh>:<mm> with a 15-minute granularity for the minutes (mm = 0 | 15 | 30 | 45).

A configuration example for a single bucket system level bandwidth policer with the following ToD-override patterns follows:

- Default rate limit: 300Mbps
- Rate limit override to 100Mbps between 5PM and 10PM
- Rate limit override to 200Mbps between 10PM and 12PM

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "P2P-Sys-DL-300M-TOD" {
                    description "Peer to Peer Policer System level Policer"
                    granularity system
                    mbs 5625
```

```
                    pir 300000
                    time-of-day-override 1 {
                        admin-state enable
                        description "Override busy hour #1"
                        mbs 1875
                        pir 100000
                        time-range {
                            daily {
                                start "17:00"
                                end "22:00"
                            }
                        }
                    }
                    time-of-day-override 2 {
                        admin-state enable
                        description "Override busy hour #2"
                        mbs 3750
                        pir 200000
                        time-range {
                            daily {
                                start "22:00"
                                end "24:00"
                                on [friday saturday]
                            }
                        }
                    }
```

The operator can display which policing rate is applied at any moment in time together with all configured override rates using the following command:

```
show application-assurance group <n> policer <policer-name> detail
```

Design and configuration examples

Default AA QoS policy

To ensure fair access for all subscribers to the ISA-AA resources, and avoid that a disproportionate amount of ISA-AA resources are used by one or more subscribers which are misbehaving or receiving large traffic bursts from the Internet, it is recommended to configure the following three types of subscriber-level default AA QoS policies:

- A default bandwidth policer to limit the downstream bandwidth per subscriber (the upstream bandwidth is already limited by ESM/SAP access ingress IOM QoS).
- A default flow count policer to limit the maximum number of active flows per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.
- A default setup rate policer to limit the maximum flow setup rate per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.

The minimum set of app-profiles used in a network is typically determined by the different access bandwidth rates; services characteristics are then used for each profile to apply a default QoS policy to limit bandwidth and flow resources accordingly.

In theory, it is possible to configure a set of default policers for every individual access bandwidth rate that is offered to a subscriber. However, this would result in a large number of policers and corresponding ASO values plus app-profiles that need to be configured. Therefore, a best practice guideline is to

define a small number of bandwidth ranges (not more than five to ten) that cover the full offered access bandwidth spectrum, and define for each bandwidth range a default bandwidth policer plus flow policers with appropriate limits.

As an example, assuming a residential deployment with two bandwidth ranges of up to 25 Mbps and 100 Mbps, the following configuration provides:

- Complete ASO and app-profile configuration.
- Default QoS policy for subscribers in the 25 Mbps range including bandwidth.
- Flow count and flow rate policers are configured by default as permit-deny. Non-conforming traffic is dropped which is common for residential deployments; alternatively the operator can decide to configure these policers as priority-mark to cut-through traffic in the ISA-AA.

In this example the resources are limited per subscriber based on their access rate maximum speed from which flow count and flow rate are derived.

App-profile and ASO

The following configuration provides the app-profile and ASO characteristics used for the default subscriber AQP policy for the 25 Mbps and 100 Mbps access bandwidth range:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-service-options {
                        characteristic "access-rate" {
                            default-value "100M"
                            value "25M" { }
                            value "100M" { }
                        }
                    }
                    app-profile "1-1/25M" {
                        description "25Mbps Site/Subscriber"
                        capacity-cost 25
                        divert true
                        characteristic "access-rate" {
                            value "25M"
                        }
                    }
                    app-profile "1-1/100M" {
                        description "100Mbps Site/Subscriber"
                        capacity-cost 100
                        divert true
                        characteristic "access-rate" {
                            value "100M"
                        }
                    }
```

Default bandwidth policing – 25 Mbps AA-sub

```
configure {
    application-assurance {
        group 1 {
            policer {
```

```
                dual-bucket-bandwidth-policer "DefltPol-Sub-BW-DS-25Mbps" {
                    description "Deflt downstream BW policer for 25Mbps Subs"
                    mbs 470
                    pir 25000
                }
```

The following AQP entry acts as a default AQP policy because it does not include application or IP header match criteria:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "Default DS BW policer for 25Mbps Subs"
                            match {
                                traffic-direction network-to-subscriber
                                characteristic "access-rate" {
                                    eq "25M"
                                }
                            }
                            action {
                                bandwidth-policer {
                                    dual-bucket "DefltPol-Sub-BW-DS-25Mbps"
                                }
                            }
                        }
```

Note:
A similar configuration can be implemented for the 100 Mbps access rate service option.

Default flow-count-limit policing – 25 Mbps AA-sub

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "DefltPol-Sub-FlowCount-US-25Mbps" {
                    granularity subscriber
                    description "Deflt policer to limit active DS flows - 25Mbps Subs"
                    peak-flow-count 10000
                    # action permit-deny        # default
                }
                flow-count-limit-policer "DefltPol-Sub-FlowCount-DS-25Mbps" {
                    granularity subscriber
                    description "Deflt policer to limit active DS flows - 25Mbps Subs"
                    peak-flow-count 10000
                    # action permit-deny        # default
                }
```

The following AQP entry acts as a default AQP policy because it does not include application or IP header match criteria:

```
configure {
    application-assurance {
```

```
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 510 {
                            admin-state enable
                            description "Deflt policer to limit act US flows - 25Mbps Subs"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "access-rate" {
                                    eq "25M"
                                }
                            }
                            action {
                                flow-count-limit {
                                    policer-name "DefltPol-Sub-FlowCount-US-25Mbps"
                                }
                            }
                        }
                        entry 515 {
                            admin-state enable
                            description "Deflt policer to limit act DS flows - 25Mbps Subs"
                            match {
                                traffic-direction network-to-subscriber
                                characteristic "access-rate" {
                                    eq "25M"
                                }
                            }
                            action {
                                flow-count-limit {
                                    policer-name "DefltPol-Sub-FlowCount-DS-25Mbps"
                                }
                            }
                        }
```

Note:
A similar configuration can be implemented for the 100 Mbps access rate service option.

Default flow-setup-rate policing – 25 Mbps AA-sub

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-setup-rate-policer "DefltPol-Sub-FlowRate-DS-25Mbps" {
                    description "Deflt policer to limit DS flow setup rate - 25Mbps Subs"
                    granularity subscriber
                    peak-flow-setup-rate 200
                }
                flow-setup-rate-policer "DefltPol-Sub-FlowRate-US-25Mbps" {
                    description "Deflt policer to limit US flow setup rate - 25Mbps Subs"
                    granularity subscriber
                    peak-flow-setup-rate 200
                }
```

The following AQP entry acts as a default AQP policy because it does not include application or IP header match criteria:

```
configure {
```

```
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 520 {
                            admin-state enable
                            description "Defl t policer to limit US flow setup rate - 25Mbps
Subs"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "access-rate" {
                                    eq "25M"
                                }
                            }
                            action {
                                flow-setup-rate-policer {
                                    policer-name "DefltPol-Sub-FlowRate-US-25Mbps"
                                }
                            }
                        }
                        entry 525 {
                            admin-state enable
                            description "Defl t policer to limit DS flow setup rate - 25Mbps
Subs"
                            match {
                                traffic-direction network-to-subscriber
                                characteristic "access-rate" {
                                    eq "25M"
                                }
                            }
                            action {
                                flow-setup-rate-policer {
                                    policer-name "DefltPol-Sub-FlowRate-DS-25Mbps"
                                }
                            }
                        }
```

Note:
A similar configuration can be implemented for the 100 Mbps access rate service option.

Application BW policing (per subscriber)

The following configuration example provides a per AA subscriber peer-to-peer rate limit of 1 Mbps. It does not include the app-profile configuration because the ASO characteristic and values can be either statically configured within the app-profile or dynamically signaled through RADIUS or Gx using ASO overrides.

AA subscribers with service characteristic "P2P-Sub-DL" value of "1M" will have a bandwidth policer of 1 Mbps applied to peer to peer traffic in the network to subscriber direction:

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "P2P-Sub-DL-1M" {
                    description "Per-subscr BW policer to limit P2P DS traffic to 1Mbps"
                    granularity subscriber
                    mbs 19
                    pir 1000
```

```
                }
```

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-service-options {
                        characteristic "P2P-Sub-DL" {
                            default-value "unlimited"
                            value "1M" { }
                            value "10M" { }
                            value "unlimited" { }
                        }
```

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 30 {
                            admin-state enable
                            description "Per-subscr BW policer to limit P2P DS traffic to
 1Mbps"
                            match {
                                traffic-direction network-to-subscriber
                                app-group {
                                    eq "Peer to Peer"
                                }
                                characteristic "P2P-Sub-DL" {
                                    eq "1M"
                                }
                            }
                            action {
                                abandon-tcp-optimization true
                                drop true
                                bandwidth-policer {
                                    single-bucket "P2P-Sub-DL-1M"
                                }
                                error-drop {
                                }
                                overload-drop {
                                }
                            }
```

Conclusion

This example provides detailed information to properly configure and use app-profiles, ASOs, and AQPs to successfully configure application policy control rules using Application Assurance.

Application Assurance — Asymmetry Removal

This chapter describes Application Assurance asymmetry removal configurations.

Topics in this chapter include:

Applicability

This chapter was originally written for and configured on SR OS Release 11.0.R1. The MD-CLI in the current edition corresponds to SR OS Release 25.7.R2.

This chapter is intended for Application Assurance (AA) network architects and engineers. The prerequisites for this chapter are a base understanding of AA configuration and operation for single homed deployments. This chapter applies to dual-homed SAPs and spoke SDPs configurations, in a business or residential AA context. AARP is not used for ESM AA subscribers.

Overview

This chapter provides best practices recommendations to configure AA asymmetry removal.

Asymmetry means that the two directions of a traffic flow (to-sub and from-sub) take different paths through the network. Asymmetry removal is a means of eliminating traffic asymmetry between a set of dual-homed SAP or spoke SDP endpoints. This can be across endpoints within a single node or across a pair of inter-chassis link connected routers, which is the topology described in this chapter. Asymmetry removal ensures all packets of a dual-homed AA subscriber are diverted to an AA ISA to achieve accurate per subscriber traffic identification and policy enforcement.

Traffic asymmetry is created when there are dual-homed links for a service, and the links are simultaneously carrying traffic. Asymmetry removal for transit subscribers must be implemented in the first routed hop on the network side of the subscriber management point, so there is a deterministic and fixed SAP or spoke SDP representing the downstream subscriber management node. This ensures there are no more than two paths that the flows can take, both covered by the asymmetry removal solution.

Configuration

Application Assurance Redundancy Protocol (AARP) provides the data plane connectivity for dynamically keeping a dual-homed AA subscriber’s traffic on the same ISA-AA for AA processing. An AARP instance is configured between the dual-homed routers to establish connectivity with the same AARP instance number on each node.

When asymmetry exists between dual-chassis redundant systems, Ipipe spoke SDPs are used to interconnect these services between peer nodes over an Inter-Chassis Link (ICL). The following sections describe the configuration and operation of the services for use with the Application Assurance Redundancy Protocol.

AARP service configuration

Figure 14: AA asymmetry removal topology shows the following services configured to establish communications between the AARP instances in each of the paired nodes. The network topology is a VPRN (or IES) service configured in each node, with a dual-homed SAP from each node to a downstream access element such as a BNG.

Figure 14: AA asymmetry removal topology

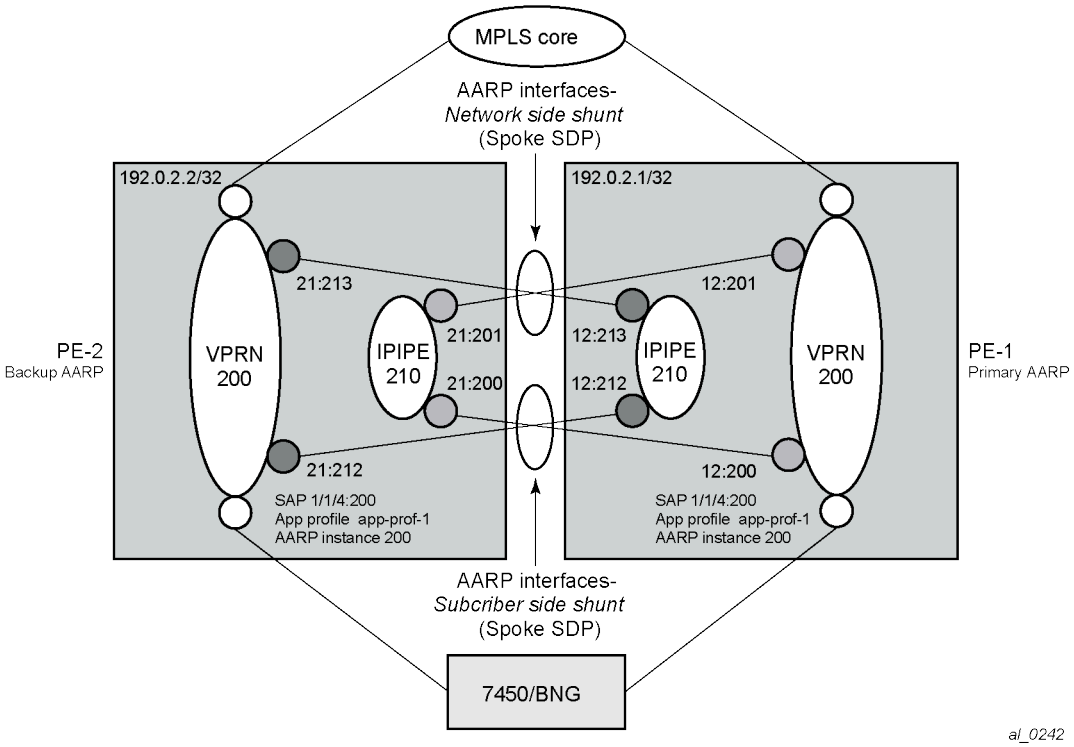

The initial configuration on PE-1 and PE-2 includes an ISA AA group and partition with an identical AA policy "app-prof-1" and divert enabled. Also, the system needs basic routing and LDP configuration for the SDP and the spoke SDPs to be established.

Table 6: AA asymmetry removal topology shows the system IP addresses, VPRN service, SAPs, and app-profile on PE-1 and PE-2.

Table 6: AA asymmetry removal topology

On PE-1	On PE-2
system IP address: 192.0.2.1	system IP address: 192.0.2.2

On PE-1	On PE-2
dual-homed service: VPRN 200	dual-homed service: VPRN 200
dual-homed SAP: 1/1/4:200	dual-homed SAP: 1/1/4:200
app-profile diverting: yes	app-profile diverting: yes

Configuration commands for AARP

To enable AARP, AARP instances and AARP interfaces on both nodes must be configured. The AARP operation has the following dependencies between the nodes:

- Shunt links configured and operationally up, both subscriber side shunt and network side shunt.
- Peer communications established between nodes, AARP instance operational status is up when peers are communicating.
- Dual-homed SAPs and spoke SDPs configured with a unique AARP instance (matched by dual-homed interface).
- App-profile "app-prof-1" configured against SAP or spoke SDP with divert enabled (making the sub an aa-sub). The app-profile is the trigger to divert the traffic in the node with the active AARP instance to one of the ISAs in that node, per normal AA divert behavior.

The following AARP configuration on PE-1 and PE-2 is similar. PE-1 is the primary node because it gets the higher priority.

```
# on primary node PE-1:
configure {
    application-assurance {
        aarp 200 {
            admin-state enable
            description "aarp protecting a dual-homed sap"
            peer 192.0.2.2
            priority 200        # higher priority (PE-1 is primary; PE-2 is backup)
        }
```

```
# on backup node PE-2:
configure {
    application-assurance {
        aarp 200 {
            admin-state enable
            description "aarp protecting a dual-homed sap"
            peer 192.0.2.1
            # priority 100        ## default
        }
```

The following SDPs are configured:

```
# on PE-1:
configure {
    service {
        sdp 12 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
```

```
                ip-address 192.0.2.2
            }
        }
```

```
# on PE-2:
configure {
    service {
        sdp 21 {
            admin-state enable
            delivery-type mpls
            ldp true
            far-end {
                ip-address 192.0.2.1
            }
        }
```

The Ipipe shunt configuration is as follows:

```
# on PE-1:
configure {
    service {
        ipipe "Ipipe-210" {
            admin-state enable
            service-id 210
            customer "1"
            vc-switching true
            service-mtu 1552
            spoke-sdp 12:212 {
                aarp {
                    id 200
                    type subscriber-side-shunt
                }
            }
            spoke-sdp 12:213 {
                aarp {
                    id 200
                    type network-side-shunt
                }
            }
        }
```

```
# on PE-2:
configure {
    service {
        ipipe "Ipipe-210" {
            admin-state enable
            service-id 210
            customer "1"
            vc-switching true
            service-mtu 1552
            spoke-sdp 21:200 {
                aarp {
                    id 200
                    type subscriber-side-shunt
                }
            }
            spoke-sdp 21:201 {
                aarp {
                    id 200
                    type network-side-shunt
                }
            }
```

```
        }
```

The dual-homed and interface shunt configuration in VPRN 200 is as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN-200" {
            admin-state enable
            service-id 200
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:200"
                }
            }
            interface "int-BNG-1" {
                sap 1/1/4:200 {
                    description "AA enabled SAP"
                    app-profile "app-prof-1"
                    aarp {
                        id 200
                        type dual-homed
                    }
                }
            }
            aarp-interface "netside_1" {
                spoke-sdp 12:201 {
                    aarp {
                        id 200
                        type network-side-shunt
                    }
                }
            }
            aarp-interface "subside_1" {
                spoke-sdp 12:200 {
                    aarp {
                        id 200
                        type subscriber-side-shunt
                    }
                }
            }
        }
```

```
# on PE-2:
configure {
    service {
        vprn "VPRN-200" {
            admin-state enable
            description "VPRN 200 Dual Homed Routed Service"
            service-id 200
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:200"
                }
            }
            interface "int-BNG-1" {
                sap 1/1/4:200 {
                    app-profile "app-prof-1"
                    aarp {
```

```
                    id 200
                    type dual-homed
                }
            }
        }
        aarp-interface "netside_1" {
            spoke-sdp 21:213 {
                aarp {
                    id 200
                    type network-side-shunt
                }
            }
        }
        aarp-interface "subside_1" {
            spoke-sdp 21:212 {
                aarp {
                    id 200
                    type subscriber-side-shunt
                }
            }
        }
    }
```

Show commands for AARP

The configuration can be verified on each node. The following output displays the example configuration for PE-1.

Starting with the AARP instance in each node, the following command verifies that the AARP instance operational state is up:

```
[/]
A:admin@PE-1# show application-assurance aarp 200

===============================================================================
AARP Instance 200
===============================================================================
Description    : aarp protecting a dual-homed sap
Admin State    : Up                     Oper State     : Up

Local IP       : 192.0.2.1              Peer IP        : 192.0.2.2
Local State    : master                 Peer State     : backup
Local Priority : 200                    Peer Priority  : 100

Local Flags    : none
Peer Flags     : none
Peer End-Point : none

Master Selection Mode       : minimizeSwitchovers

-------------------------------------------------------------------------------
Service References
-------------------------------------------------------------------------------
Service              Reference            Reference Type
-------------------------------------------------------------------------------
VPRN 200             1/1/4:200            Dual-Homed
Ipipe 210            12:212               Subscriber-Side Pipe Shunt
Ipipe 210            12:213               Network-Side Pipe Shunt
VPRN 200             12:200               Subscriber-Side AARP-Interface Shunt
VPRN 200             12:201               Network-Side AARP-Interface Shunt
-------------------------------------------------------------------------------
```

```
No. of service references: 5
-------------------------------------------------------------------------------
===============================================================================
```

Verifying that the AARP instance is up is an indication that the dual-node communications for AARP is working (instance, shunts, and so on). In addition, in the preceding output, verify on both PE nodes that the intended SAPs are dual-homed for that instance.

Now a detailed review of the configured AARP shunt infrastructure services can be shown to make sure they are all properly configured with the intended AARP parameters (such as AARP ID and Type on the network and subscriber side shunts) as displayed in the following output:

```
[/]
A:admin@PE-1# show service id "Ipipe-210" all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 210                 Vpn Id            : 0
Service Type      : Ipipe
MACSec enabled    : no
Name              : Ipipe-210
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 09/16/2025 12:47:16
Last Mgmt Change  : 09/16/2025 12:47:03
Admin State       : Up                  Oper State        : Up
MTU               : 1552
Vc Switching      : True
SAP Count         : 0                   SDP Bind Count    : 2
CE IPv4 Discovery : n/a
CE IPv6 Discovery : n/a                 Stack Cap Sig     : n/a

-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore


-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:212  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:212                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
VC Type           : Ipipe                   VC Tag            : 0
Admin Path MTU    : 0                       Oper Path MTU     : 1552
Delivery          : MPLS
Far End           : 192.0.2.2               Tunnel Far End    :
Oper Tunnel Far End: 192.0.2.2
LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled

Admin State       : Up                      Oper State        : Up
MinReqd SdpOperMTU : 1552
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 524281                  Egress Label      : 524281
```

```
---snip---

Application Profile: None
Transit Policy     : None
AARP Id            : 200
AARP Type          : subscriber-side-shunt

---snip---

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IPv4 Addr: n/a                 Peer CE IPv4 Addr : 0.0.0.0

-------------------------------------------------------------------------------
 Sdp Id 12:213  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 12:213                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
VC Type            : Ipipe                   VC Tag            : 0
Admin Path MTU     : 0                       Oper Path MTU     : 1552
Delivery           : MPLS
Far End            : 192.0.2.2               Tunnel Far End    :
Oper Tunnel Far End: 192.0.2.2
LSP Types          : LDP
Hash Label         : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                      Oper State        : Up
MinReqd SdpOperMTU : 1552
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 524280                  Egress Label      : 524280

---snip---

Application Profile: None
Transit Policy     : None
AARP Id            : 200
AARP Type          : network-side-shunt

---snip---
```

Next, the configuration of the VPRN service of the dual-homed SAP can be reviewed to ensure it reflects the attached endpoints for the shunt Ipipe spoke SDPs:

```
[/]
A:admin@PE-1# show service id "VPRN-200" all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 200                Vpn Id            : 0
Service Type      : VPRN
MACSec enabled    : no
Name              : VPRN-200
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 09/16/2025 12:47:03
Last Mgmt Change  : 09/16/2025 12:47:03
```

```
Admin State       : Up                  Oper State        : Up

Router Oper State : Up
Route Dist.       : 64496:200           VPRN Type         : regular
Oper Route Dist   : 64496:200
Oper RD Type      : configured
AS Number         : None                Router Id         : 192.0.2.1
ECMP              : Enabled             ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Local Rt Domain-Id: None                D-Path Lng Ignore : Disabled

Auto Bind Tunnel
Allow Flex-Alg-Fb : Disabled
Resolution        : disabled
Weighted ECMP     : Disabled            ECMP Max Routes   : 1
Strict Tnl Tag    : Disabled

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : None
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled
LOG all events    : Disabled

SAP Count         : 1                   SDP Bind Count    : 2

-------------------------------------------------------------------------------
---snip---

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:200  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:200                  Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : n/a                     VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 1552
Delivery          : MPLS
Far End           : 192.0.2.2               Tunnel Far End    :
Oper Tunnel Far End: 192.0.2.2
LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled

Admin State       : Up                      Oper State        : Up

---snip---

Application Profile: None
Transit Policy    : None
AARP Id           : 200
```

```
AARP Type          : subscriber-side-shunt

---snip---

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IPv4 Addr: n/a                  Peer CE IPv4 Addr : 0.0.0.0

-------------------------------------------------------------------------------
 Sdp Id 12:201  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 12:201                   Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                      VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 1552
Delivery           : MPLS
Far End            : 192.0.2.2                Tunnel Far End    :
Oper Tunnel Far End: 192.0.2.2
LSP Types          : LDP
Hash Label         : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                       Oper State        : Up

---snip---

Application Profile: None
Transit Policy     : None
AARP Id            : 200
AARP Type          : network-side-shunt

---snip---
```

Continuing deeper into the same VPRN service **show** output, or using the following **show** command, it can be verified that the dual-homed SAP itself is properly configured and associated with that service and AARP instance:

```
[/]
A:admin@PE-1# show service id "VPRN-200" sap 1/1/4:200 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 200
SAP                : 1/1/4:200                Encap             : q-tag
Description        : AA enabled SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 09/16/2025 12:35:38
Last Mgmt Change   : 09/16/2025 12:47:03
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
```

```
qinq-pbit-marking  : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
Qinq-vlan-                                   Qinq-vlan-
translation        : None                    translation Ids   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                       Tot Static Hosts  : 0
Calling-Station-Id : n/a

Application Profile: app-prof-1
Transit Policy     : None
AARP Id            : 200
AARP Type          : dual-homed

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Bandwidth          : Not-Applicable
Oper DCpu Prot Pol : _default-access-policy
Virtual Port       : (Not Specified)

-------------------------------------------------------------------------------
---snip---
```

Network to subscriber traffic flow

When the AARP is operationally up, AARP tracks which ISA is the primary ISA for each dual-homed AARP instance and uses the inter-chassis services (spoke SDP AARP shunts) to move all traffic for each instance traffic to the node with the primary ISA. Traffic from the backup AARP interface is sent the primary AARP node for AA divert. Afterward the traffic is sent back to the backup AARP interface.

Looking at traffic in the network to subscriber direction (Figure 15: Network to subscriber traffic flow):

- Traffic arriving on PE-1 is diverted to the local primary ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP interface is sent to the primary node for AA processing. The ingress FP forwards packets to network-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the primary ISA where the packets are processed as if this traffic was traveling in the to-sub direction toward the dual-homed endpoint on PE-1, then returned to PE-2.
- Entering PE-2, the traffic from the subscriber side shunt interface is not diverted to ISAs in that node and egresses on the AARP instance SAP.

With this behavior, traffic always returns to the original ingress node before egressing toward the subscriber (network path for the flows are not modified).

Figure 15: Network to subscriber traffic flow

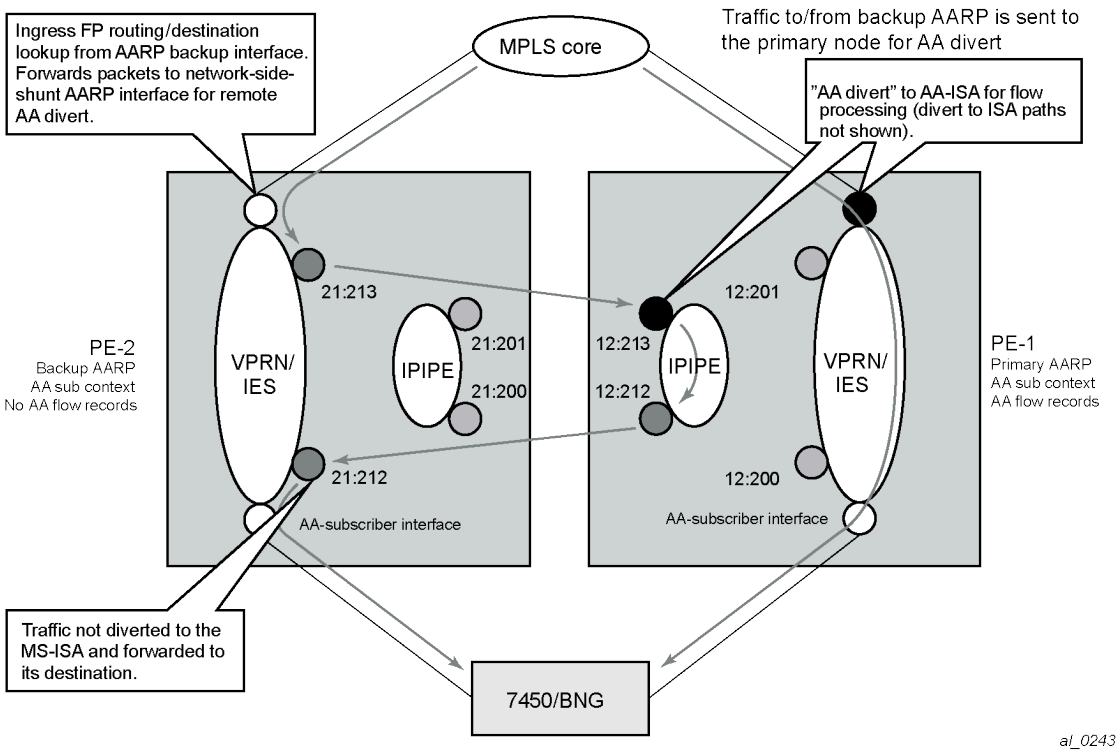

Subscriber to network traffic flow

Looking at traffic in the subscriber to network direction (Figure 16: Subscriber to network traffic flow):

- Traffic arriving on PE-1 is diverted to the local primary ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP ISA is sent to the primary node for AA processing (not diverted to an ISA in PE-2). The ingress FP forwards packets to subscriber-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the primary ISA where the packets are processed as if the traffic was flowing in the from-sub direction on the dual-homed endpoint, then returned to PE-2 over the Ipipe AARP subscriber-side-shunt.
- Entering PE-2, the traffic from the network side shunt interface is forwarded by the VPRN or IES service to its destination.

Figure 16: Subscriber to network traffic flow

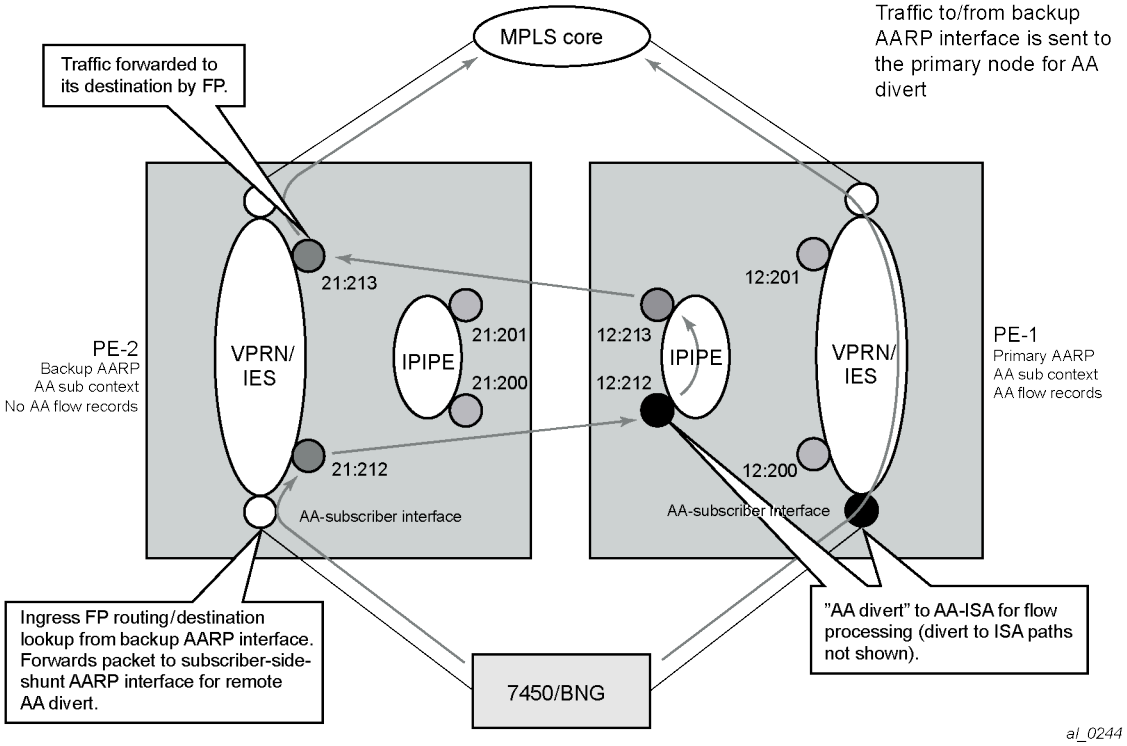

Typical configuration mistakes

Operators configuring AARP can make some typical mistakes listed below that keep the AARP instance in operational state down:

- The spoke SDP AARP shunt instances' IDs must be aligned with the respective spoke SDP on the peer node: if not, it results in a flag indicating that the shunts are down in the **show** output.
- Ipipe service MTU alignment — The Ipipe service MTU values must be the same in both nodes, otherwise it results in the VPRN or IES services in operational status up, but the AARP instance remains down.

Conclusion

This chapter is intended for Application Assurance (AA) network architects and engineers to provide the information required to understand and configure dual-node asymmetry removal following the intended service configuration as used by the AARP implementation.

Application Assurance — GTP Roaming Firewall

This chapter describes Application Assurance GPRS Tunneling Protocol (GTP) roaming firewall.

Topics in this chapter include:

Applicability

This chapter was initially written based on SR OS Release 20.2.R2, but the MD-CLI configuration in the current edition is based on SR OS Release 25.7.R1. The AA FW feature related to the GTP roaming interface in the AA GTP filtering functionality is supported in SR OS Release 20.2.R2 and later.

Overview

Wireless network operators rely on GPRS Tunneling Protocol (GTP) for the delivery of mobile data services across the access network. However, GTP is not designed to be secure, exposing the mobile access network to attacks from both its own subscribers and its partner networks.

The Application Assurance (AA) SR OS 20.2.R2 firewall feature extends AA-Integrated Service Adapter (AA-ISA) application-level analysis to provide an in-line stateful service that integrates into a 7750 SR router. The feature provides protection for mobile operator infrastructure against attacks from compromised mobile gateways: Serving Gateways (SGWs) or Packet data network Gateways (PGWs).

AA stateful packet filtering, combined with AA L7 classification and control, provides operators with advanced, next-generation firewall (FW) functionality. This AA stateful FW feature runs on AA-ISA and, using stateful inspection, not only inspects packets at Layers 3 to 7, but also monitors and keeps track of the connection state, as described in the Application Assurance — Stateful Firewall chapter. Figure 17: AA GTP roaming FW deployment shows an example AA GTP roaming FW deployment.

Figure 17: AA GTP roaming FW deployment

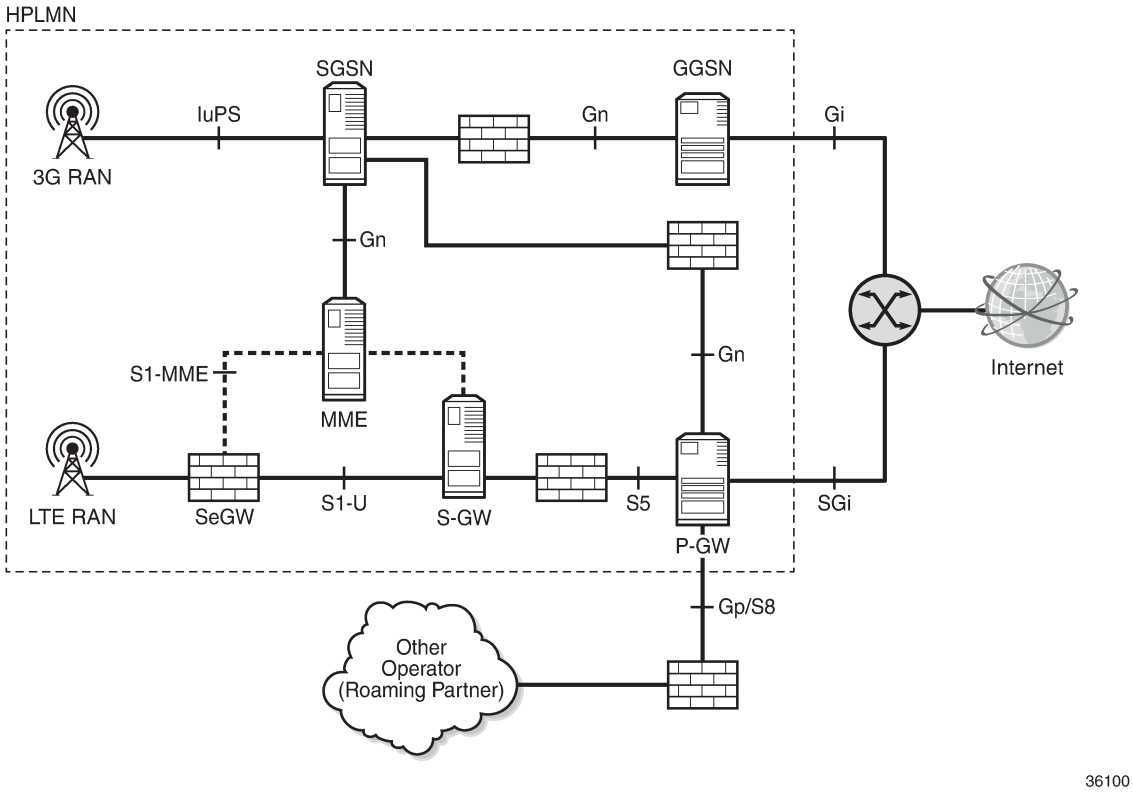

S8/Gp AA FW deployment

AA FW is deployed as a GTP FW on S8/Gp (or S5/Gn) interfaces, either as part of a 7750 SR router in the form of an AA-ISA hardware module or as a separate Virtual SR (VSR) appliance. AA FW provides operators with network security, such as:

- GTP protocol validation, which checks for anomaly attacks that involve malformed, corrupt, or spoofed traffic:
  - header length checks
  - Information Element (IE) length validation
  - invalid reserved field validation
  - reserved IE validation
  - missing mandatory IE validation
  - sequence number validation
  - Tunnel endpoint identifier (TEID) validation - blocks GTP tunnel creations that have not been signaled correctly
- PGW and SGW redirection protection
- GTP-in-GTP check

- Handover control to prevent session hijacking
- User equipment (UE) source IP address anti-spoofing protection
- Protection against unauthorized Public Land Mobile Network (PLMN) and Access Point Name (APN) access:
  - filter messages based on the combination of an International Mobile Subscriber Identity (IMSI) prefix and an APN
- Protection against unsupported GTP message types:
  - filter messages based on message type
  - filter messages based on message length
- Protection against flooding attack:
  - GTP traffic bandwidth policing, which limits the GTP bandwidth from a roaming partner SGW or PGW
  - GTP tunnel limiting, which limits the number of concurrent GTP tunnels and the setup rate of these tunnels from a roaming partner SGW or PGW
- Protection against IP fragmentation-based attacks: drop rules for IP fragmentation of GTP messages

AA FW supports both GTPv1 and GTPv2. It is typically deployed as an L3 VPRN service. SAPs or spoke SDPs are diverted to AA for a GTP FW. L2 VPLS connectivity is supported by AA. AA transit subscribers (identified by SGW IPs) are auto-created under the parent-diverted SAPs or spoke SDPs.

GTP TEID validation

Compromised mobile gateways (GSNs) can send storms of GTP traffic with invalid GTP TEIDs to cause a denial of service (DoS) attack. By inspecting GTP messages for control plane (GTP-C), AA FW supports stateful correlation of upstream and downstream GTP flows (DstIP + TEID) of the same packet data network (PDN) session. AA drops packets with TEIDs that have not been negotiated correctly.

The operator can enable AA to drop GTP traffic with an invalid TEID using:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        gtp-tunnel-database {
                            validate-gtp-tunnels true
```

The GTP-C inspection discards any message type that is invalid for S5/S8 (Gn/p) interfaces and ensures the correct IE states and order.

UE IP address anti-spoofing

Source IP address spoofing is initiated by a malicious UE that hijacks (spoofs) an IP address of another UE and invokes a download from a malicious server on the Internet. After the download begins, the malicious

UE exits the session. The UE under attack (receiving the download traffic) gets charged for traffic it did not solicit.

AA FW associates the GTP-C messages of the UE IP address IE with the GTP for user plane (GTP-U) packets to ensure that the packets carried in the upstream have the correct source IP address (inner IP within the GTP-U tunnel). Because the UE address is negotiated within the packet data protocol (PDP) context creation handshake, any packets originating from the UE that contain a different source address are detected by AA FW and dropped.

To enable UE IP address anti-spoofing protection, the operator needs to enable **validate-source-ip-addr**, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        gtp-tunnel-database {
                            validate-gtp-tunnels true
                            validate-source-ip-addr true
```

GTP anomaly prevention —sequence number check

Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of the protocol specifications. Packets are denied by AA FW if they fail the sanity check. The following are some examples of GTP sanity checks:

- invalid GTP header length
- invalid IE length
- invalid reserved fields
- invalid sequence number
- missing mandatory IEs

Also, AA FW performs sequence number validation whereby it ensures no out-of-sequence GTP packets. By default, sequence number validation is disabled. To enable it, The following CLI command enables sequence number validation:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        gtp-tunnel-database {
                            validate-gtp-tunnels true
                            validate-sequence-number true
```

GTP packets with wrong sequence numbers are dropped when sequence number validation is enabled.

GTP message-type filtering

AA FW performs GTP message validation, in which packets with invalid message types (that are not applicable to the roaming interfaces) are denied by the AA GTP-C inspection command:

Table 7: Denied GTP message types for roaming interface

	GTP-U port	GTP-C port
Denied GTPv1 message types	None	GTPU_PDU GTPV1_END_MARKER GTPV1_MSG_ERR_IND GTPV1-ALL-MBMS message-types GTPV1-ALL-Location management message-types
Denied GTPv2 message types	N/A	GTP_PKT_ERROR_INDICATION GTP_PKT_DNLK_DATA_FAIL_INDICATION GTP_PKT_STOP_PAGING_INDICATION GTP_PKT_CRE_INDR_TNL_REQ GTP_PKT_CRE_INDR_TNL_RSP GTP_PKT_DEL_INDR_TNL_REQ GTP_PKT_DEL_INDR_TNL_RSP GTP_PKT_RELEASE_BEARERS_REQ GTP_PKT_RELEASE_BEARERS_RSP GTP_PKT_DNLK_DATA GTP_PKT_DNLK_DATA_ACK GTP_PKT_MOD_ACCESS_BEARERS_REQ GTP_PKT_MOD_ACCESS_BEARERS_RSP

Also, AA FW allows the operator to further restrict allowed message types (shown in the following table) by configuring GTP message type filter entries to deny (or allow) the following message types:

Table 8: Allowed GTP message types (Cat-1)

	GTP-U port	GTP-C port
Allowed GTPv1 message types	GTPV1_MSG_ECHO_REQ GTPV1_MSG_ECHO_RESP GTPV1_SUPP_EXT_HDR_NOTIF GTPV1_MSG_ERR_IND GTPV1_END_MARKER	GTPV1_MSG_ECHO_REQ GTPV1_MSG_ECHO_RESP GTPV1_SUPP_EXT_HDR_NOTIF GTPV1_MSG_VER_NOT_SUPP_IND GTPV1_MSG_PDP_CREATE_REQ GTPV1_MSG_PDP_CREATE_RESP

	GTP-U port	GTP-C port
	GTPU_PDU	GTPV1_MSG_PDP_UPD_REQ GTPV1_MSG_PDP_UPD_RESP GTPV1_MSG_PDP_DEL_REQ GTPV1_MSG_PDP_DEL_RESP GTPV1_MSG_NET_INIT_REQ GTPV1_MSG_NET_INIT_RESP GTPV1_MSG_MSINFO_REQ GTPV1_MSG_MSINFO_RESP
Allowed GTPv2 message types	N/A	GTP_PKT_ECHO_REQ GTP_PKT_ECHO_RSP GTP_PKT_VERSION_NOT_SUPPORTED GTP_PKT_CRE_SES_REQ GTP_PKT_CRE_SES_RSP GTP_PKT_MOD_BEARER_REQ GTP_PKT_MOD_BEARER_RSP GTP_PKT_DEL_SES_REQ GTP_PKT_DEL_SES_RSP GTP_PKT_CHG_NOT_REQ GTP_PKT_CHG_NOT_RSP GTP_PKT_MOD_BEARER_CMD GTP_PKT_MOD_BEARER_FAIL_INDICATION GTP_PKT_DEL_BEARER_CMD GTP_PKT_DEL_BEARER_FAIL_INDICATION GTP_PKT_BEARER_RESOURCE_CMD GTP_PKT_BEARER_RESOURCE_FAIL_INDICATION GTP_PKT_SUSPEND_NOTIFICATION GTP_PKT_SUSPEND_ACK GTP_PKT_RESUME_NOTIFICATION GTP_PKT_RESUME_ACK GTP_PKT_CRE_BEARER_REQ GTP_PKT_CRE_BEARER_RSP GTP_PKT_UPD_BEARER_REQ GTP_PKT_UPD_BEARER_RSP GTP_PKT_DEL_BEARER_REQ GTP_PKT_DEL_BEARER_RSP

	GTP-U port	GTP-C port
		GTP_PKT_TRACE_SESSION_ACTIVATION GTP_PKT_TRACE_SESSION_DEACTIVATION GTP_PKT_UPDATE_PDN_CONNECTION_SET_REQ GTP_PKT_UPDATE_PDN_CONNECTION_SET_RSP GTP_PKT_DELETE_PDN_CONNECTION_SET_REQ GTP_PKT_DELETE_PDN_CONNECTION_SET_RSP

By default, the GTP message filter allows all GTP messages.

The following command is used to configure GTPv2 message filtering:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        message-type-gtp-v2 {
                            default-action {permit|deny}
                            entry <id:516..770 | expression> {
                                action {permit|deny}
                                value <gtpv2-message-value>
```

The following command is used to configure GTPv1 message filtering:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        message-type {
                            default-action {permit|deny}
                            entry <id:1..255 | expression> {
                                action {permit|deny}
                                value <gtpv1-message-value>
```

Note:
If the operator configures a message type that is invalid for the roaming interface to be denied, it is dropped and counted under that filter entry (and not tagged as dropped because of the "wrong-interface" in the event log). However, configuring the message-type filter to "permit" a message type that is invalid for the roaming interface does not take effect, because the packet with the specified message type is dropped by the GTP-C protocol inspection process.

Unauthorized APN attack – APN filtering

APN filtering checks GTP-C messages to determine if a roaming subscriber is allowed to access a specified external network (access point name – APN). The "create session request" and "create pdp request" GTP message types contain an APN IE in the header of a GTP packet. An APN IE consists of an external network ID (for example, nokia.com) and, optionally, a unique ID that identifies the operator PLMN.

APN filtering prevents malicious UEs from initiating a "create PDP/session request" flood attack toward the PGW/GGSN for invalid or disallowed APNs. The operator can configure an AA GTP filter to perform APN filtering to restrict roaming subscribers access to specific external networks.

An APN filter, an IMSI prefix, and an SGSN address pool can be used together to filter GTP packets, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-roaming-FW" {
                        imsi-apn-filter {
                            default-action {permit|deny}      ## default permit
                            entry <entry-id:1031..2030 | expression> {
                                src-gsn <ip-prefix | ip-prefix-list <ip-prefix-list> >
                                apn <apn-string>
                                imsi-mcc-mnc-prefix < [1 to 8 digits] | "ANY_IMSI" >
                                action {permit|deny}         ## default permit
```

By default, AA FW allows all APNs.

Unauthorized PLMN access – IMSI prefix filtering

The PLMN of a subscriber home network is identified by combining the Mobile Country Code (MCC) and Mobile Network Code (MNC). MCC-MNC is also known as the International Mobile Station Identity (IMSI) prefix. The IMSI prefix acts as a PLMN identifier.

GTP IMSI prefix filters can be configured to deny GTP incoming traffic from invalid roaming partners. Conversely, GTP IMSI prefix filters can allow only incoming traffic from those network operators that have signed roaming agreements. Any GTP packets with IMSI prefixes not matching the configured prefixes are dropped.

As shown in the Unauthorized APN attack – APN filtering section, an IMSI filter entry can also be optionally combined with an SGSN/SGW IP address (or IP address prefix list) to further restrict allowed IMSI prefix traffic to specific SGSN/SGW nodes.

Unauthorized network access

An attacker, using an unauthorized GSN, can cause a DoS attack using spoofed PDP context delete messages (DoS attack) or spoofed update PDP context requests to hijack existing sessions. Such attacks can also spoof a create PDP context request to gain unlawful Internet access. Session hijacking can come

from either the SGW/SGSN or the PGW/GGSN. An unauthorized GSN can hijack GTP tunnels or cause a DoS attack by intercepting another GSN and redirecting traffic to it.

Operators can use AA FW to configure pools of trusted GSN IP addresses (using AA IP-prefix-list) to stop spoofed requests from untrusted GSNs. AA IP prefix lists can be configured to model GSN groups, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list <ip-prefix-list-name> {
                    prefix <ip-prefix/ip-prefix-length> {
                        name <prefix-name>
```

These lists are then referenced in session filters, such that only sessions that match the lists can be permitted, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter <session-filter-name> {
                    default-action {
                        action deny
                    }
                    entry 1 {
                        match {
                            src-ip {
                                ip-prefix-list <ip-prefix-list-name>
                            }
                        }
                        action {
                            permit
                        }
```

Configuration

AA GTP filtering functionality includes support for the AA FW feature related to the GTP roaming interface. The AA GTP filters are optional Application QoS Policy actions (AQP actions). AQPs have partition-level scope, which allows different FW policies to be implemented by using AA partition concepts within the same AA-ISA.

The configuration topology in Figure 18: Configuration topology shows how the VSR equipped with AA FW functionality provides protection for the S8 interfaces.

Figure 18: Configuration topology

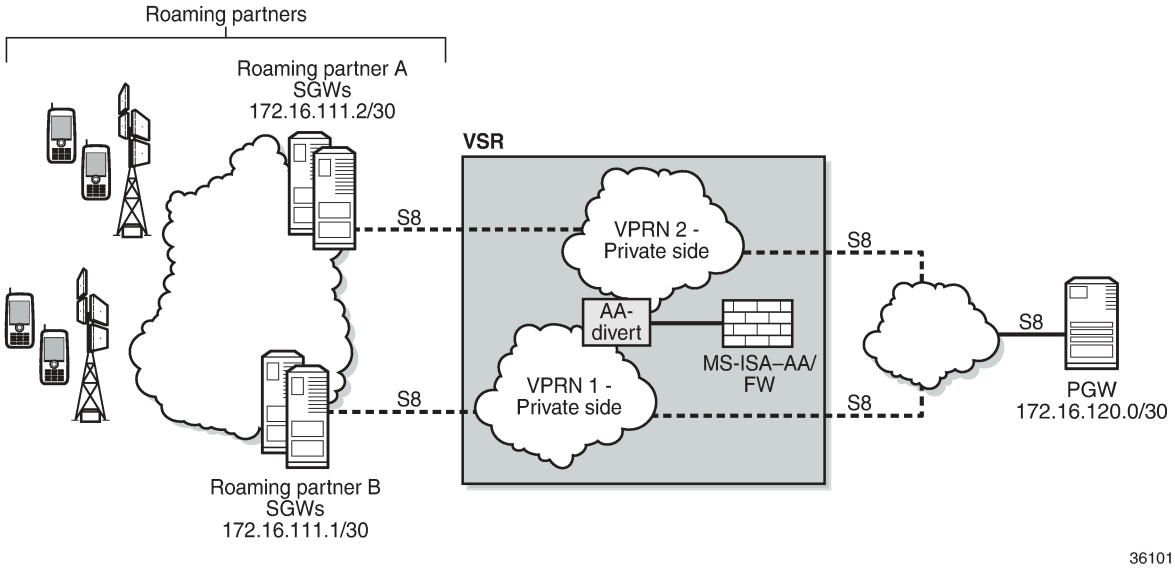

Pre-setup requirements

Configuration of a VSR router is required if a 7750 SR is not already used in the access network on the S8 interfaces. If a 7750 SR is already deployed, AA-ISA must be configured.

See the Application Assurance — Stateful Firewall chapter for basic knowledge about AA-FW functionality.

Platform-dependent configuration

VSR

For GTP FW deployment in VSR, the following configuration supports load balancing of traffic across multiple CPU cores:

```
configure {
    isa {
        application-assurance-group 1 {
            vm-traffic-distribute-by-teid true
```

Note: The **vm-traffic-distribution-by-teid** command is only supported on VSR, not on hardware (ISA2 or ESA). To distribute work across cores, HW uses a different mechanism on its own.

7750 hardware

GTP FW deployment in 7750 SR hardware is supported by ISA2:

```
configure {
    isa {
        application-assurance-group 1 {
            minimum-isa-generation 2
```

Allocation of memory for stateful GTP processing

To support stateful GTP processing (for example, TEID, sequence number, and UE IP validation FW operations), the operator must configure the system to allocate sufficient memory resources, as follows:

```
configure {
    isa {
        application-assurance-group 1 {
            shared-resources {
                gtp-tunnel-database 100
```

Configuration to divert SAPs/VPRN traffic into AA-ISA

In this configuration example, one VPRN is used per wireless roaming partner network. In the example, two roaming partner networks are used for illustration. In real networks, this number is much bigger.

However, before configuring SAPs for diversion, the operator can optionally define some Application Service Option (ASO) characteristics to provide different FW policies for different roaming partners, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-service-options {
                        characteristic "FW-Protection" {
                            default-value "OFF"
                            value "OFF"
                            value "ON"
                        }
                        characteristic "strict-FW-Protection" {
                            default-value "OFF"
                            value "OFF"
                            value "ON"
                        }
                    }
```

For more information about ASO configuration, see the Application Assurance — App-Profile, ASO and Control Policies chapter.

After configuring any ASO characteristics, define an application profile and transit IP policy; for example:

```
configure {
    application-assurance {
```

```
        group 1 {
            partition 1 {
                policy {
                    app-profile "default" {
                        description "App profile that applies to the whole SAP"
                        divert true
                        characteristic "FW-Protection" {
                            value "ON"
                        }
                    }
                    app-profile "strict-FW" {
                        description "App profile that applies strict FW rules to the SAP"
                        divert true
                        characteristic "FW-Protection" {
                            value "ON"
                        }
                        characteristic "strict-FW-Protection" {
                            value "ON"
                        }
                    }
                }
                transit-ip-policy 1 {
                    default-app-profile "strict-FW"
                    detect-seen-ip true
                    transit-auto-create {
                        admin-state enable
                    }
                }
```

Traffic of the following two VPRNs needs to be diverted into AA-ISA to provide firewall protection. The preceding application profiles and transit IP policy are applied to the SAPs of the two VPRNs:

```
configure {
    service {
        customer "1" {
            description "Default customer"
            customer-id 1
        }
        customer "2" {
            customer-id 2
        }
        vprn "VPRN 100" {
            admin-state enable
            description "L3 service roaming partner 2"
            service-id 100
            customer "2"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.1:100"
                }
            }
            interface "to-NetA" {
                 ipv4 {
                    primary {
                        address 192.168.100.1
                        prefix-length 24
                    }
                }
                sap 1/1/9:100 {
                    app-profile "default"
                }
            }
```

```
        }
        vprn "VPRN 200" {
            admin-state enable
            description "L3 service roaming partner 1"
            service-id 200
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "192.0.2.1:200"
                }
            }
            interface "to-site1" {
                ipv4 {
                    primary {
                        address 10.1.201.1
                        prefix-length 24
                    }
                    neighbor-discovery {
                        static-neighbor 10.1.201.2 {
                            mac-address 00:00:5e:00:53:21
                        }
                    }
                }
                sap 1/1/9:201 {
                    app-profile "strict-FW"
                    transit-policy {
                        ip 1
                    }
                }
            }
            static-routes {
                route 10.10.68.1/32 route-type unicast {
                    next-hop "10.1.201.2" {
                        admin-state enable
                    }
                }
            }
        }
```

This configuration achieves the following.

1. Roaming traffic is diverted to AA-ISA for FW protection.
2. Customer 1 traffic has a "strict" FW rule attribute, while customer 2 traffic is subject to basic FW rules.
3. Within AA-ISA, the customer 1 diverted SAP is treated as a parent SAP. Instead of treating the whole SAP as a single subscriber, subscribers are auto-created within this SAP, based on the IP address of the SGWs/SSGNs.

Note:
If the operator does not require per SGW/SSGN control (such as limiting the total bandwidth of a SGW to prevent DoS attack), the transit IP policy from the SAP configuration can be removed. This causes AA to treat the whole SAP as a single subscriber, as in the case of the customer 2 SAP.

Configuration FW events log

The following configures a log that captures events related to various AA firewall actions:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                event-log "FW_events_log" {
                    admin-state enable
                    buffer-type circular
                    max-entries 100000
                }
```

Note:
Alternatively, because of the limited size of the log and the large amount of traffic AA can handle, Nokia recommends that the operator use the syslog mechanism instead of local logging, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                event-log "FW_events_log" {
                    syslog {
                    }
```

The event log can be referenced in various FW actions that are configured later in this chapter.

```
[/]
A:admin@PE-1# tools dump application-assurance group 1:1 event-log "FW_events_log" isa 1/2

=====================================================
Application-Assurance event-log "FW_events_log"
Current Time:          "08/15/2025 07:38:36" (UTC)
  group[:partition]:   1:1
  isa:                 1/2
  admin state:         no shutdown
  buffer-type:         circular
  max-entries:         100000
=====================================================
Event-source
Action        SubType     SubName                            Direction Src-ip
Dst-ip                          Ip-protocol Src-port Dst-port Timestamp
"gtp filter gtp-filter-partner1 reason: filtered-gtp-message-type, teid: 0x0001d100,
 MT: 36, version: 2"                                                          deny
   transit     "1_10.10.68.1/32"                 from-sub  10.10.68.1
10.10.68.3                      udp         2123     2123     "08/15/2025 18:54:50"
"gtp filter gtp-filter-partner1 reason: filtered-gtp-message-type, teid: 0x00019100,
 MT: 37, version: 2"                                                          deny
   transit     "1_10.10.68.1/32"                 to-sub    10.10.68.3
10.10.68.1                      udp         2123     2123     "08/15/2025 18:54:55"

Total Records:   2
=====================================================
```

The following command clears all the entries within the specified log:

```
clear application-assurance group 1:1 event-log "FW_events_log"
```

Configuration to limit total traffic from SGWs

Nokia recommends that a total limit be placed on how much bandwidth and how many flows an SGW/SGSN can generate toward the network. The exact limit values are a function of the number of end devices that are served by the roaming partner SGW/SGSN and capacity limits of the Home PLMN (HPLMN) PGW/GGSN, plus some additional margin.

In the following example, it is assumed that traffic from each roaming SGW does not exceed 1200 concurrent flows per second (serving about 200 roaming UEs) and 50 Mb/s. These need to be replaced in actual deployments with appropriate values that reflect the specific network deployment.

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "limit_roamingSGW_Flows" {
                    granularity subscriber
                    limit-gtp-flows true
                    peak-flow-count 1200
                }
                single-bucket-bandwidth-policer "limit_roamingSGW_bw" {
                    granularity subscriber
                    mbs 500
                    pir 50000
                }
```

The configured policers are applied as actions from within the default subs-policy AQP entry:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "limit per SGW flow and b/w- partner 1"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "strict-FW-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "limit_roamingSGW_bw"
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_roamingSGW_Flows"
                                }
                            }
                        }
                    }
```

For GTP traffic flow count policing, it is important that "aqp-initial-lockup" is enabled:

```
configure {
    application-assurance
        group 1 {
            partition 1 {
                aqp-initial-lookup true
```

Note:
All the preceding actions apply to the traffic direction "subscriber-to-network". These actions do not apply to traffic in the other direction (downlink), because the purpose of the AA FW is to protect the network resources from upstream traffic from compromised roaming partner SGWs.

Note:
No policers are placed for the traffic of customer 2, because its profile does not have "strict policing" enabled. A policer can be configured to limit the total bandwidth and flows from all SGWs served to the customer 1 SAP, as follows:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "limit_roamingSGWs_total_Flows" {
                    granularity subscriber
                    limit-gtp-flows true
                    peak-flow-count 12000
                }
                single-bucket-bandwidth-policer "limit_roamingSGWs_total_bw" {
                    granularity subscriber
                    mbs 5000
                    pir 500000
                }
```

The configured policers are applied as actions from within the default subs-policy AQP entry, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 501 {
                            admin-state enable
                            description "limit total SGW flow and b/w- partner 2 "
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "strict-FW-Protection" {
                                    eq "OFF"
                                }
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "limit_roamingSGWs_total_bw"
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_roamingSGWs_total_Flows"
                                }
                            }
                        }
```

GTP filtering – disallow traffic from unauthorized SGWs

To use GTP filtering to disallow traffic from unauthorized SGWs, perform the following steps:

1. Create AA IP lists
2. Use AA IP lists in session filters
3. Reference session filters within AQP entries

Create AA IP lists

Create AA IP lists, by creating an AA IP prefix list that contains SGW IP addresses or range of addresses for each customer, as follows:

Roaming partner 1

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list "Roaming1_ALL_SGWs" {
                    description "SGWs subnet for roaming partner 1"
                    prefix 172.16.100.0/24 {
                    }
                }
```

Roaming partner 2

```
                ip-prefix-list "Roaming2_ALL_SGWs" {
                    description "SGWs subnet for roaming partner 2"
                    prefix 172.16.110.100/30 {
                    }
                }
```

Use AA IP lists in session filters

The AA IP prefix lists can be referenced and used in AA FW rules using session filters to be referenced in AQPs, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter "restricted_access_partner1" {
                    description "SGWs_allowed_partner1"
                    default-action {
                        action deny
                    }
                    entry 10 {
                        description "allow GTP-U from authorized subnets"
                        match {
                            ip-protocol udp
                            dst-port {
                                eq 2152
                            }
```

```
                        src-ip {
                            ip-prefix-list "Roaming1_ALL_SGWs"
                        }
                    }
                    action {
                        permit
                    }
                }
                entry 11 {
                    description "allow GTP-C from authorized subnets"
                    match {
                        ip-protocol udp
                        dst-port {
                            eq 2123
                        }
                        src-ip {
                            ip-prefix-list "Roaming1_ALL_SGWs"
                        }
                    }
                    action {
                        permit
                    }
                }
                entry 20 {
                    description "allow DNS"
                    match {
                        ip-protocol tcp-udp
                        dst-port {
                            eq 53
                        }
                        src-ip {
                            ip-prefix-list "Roaming1_ALL_SGWs"
                        }
                    }
                    action {
                        permit
                    }
                }
            }
            session-filter "restricted_access_partner2" {
                description "SGWs_allowed_partner2"
                default-action {
                    action deny
                    event-log "FW_events_log"
                }
                entry 10 {
                    description "allow GTP-U from authorized subnets"
                    match {
                        ip-protocol udp
                        dst-port {
                            eq 2152
                        }
                        src-ip {
                            ip-prefix-list "Roaming2_ALL_SGWs"
                        }
                    }
                    action {
                        permit
                    }
                }
                entry 11 {
                    description "allow GTP-C from authorized subnets"
                    match {
                        ip-protocol udp
```

```
                            dst-port {
                                eq 2123
                            }
                            src-ip {
                                ip-prefix-list "Roaming2_ALL_SGWs"
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 20 {
                        description "allow DNS"
                        match {
                            ip-protocol tcp-udp
                            dst-port {
                                eq 53
                            }
                            src-ip {
                                ip-prefix-list "Roaming2_ALL_SGWs"
                            }
                        }
                        action {
                            permit
                        }
                    }
                }
```

Note:
Optionally, you can combine the session filter entries for the two roaming partners into a single session filter (for scale reasons). AA supports a total of 300 session filters. If there are less than 300 roaming partners, you can use a session filter per partner for customization purposes (related to, for example, IP subnets). If the number of partners is greater than the maximum number of session filters, you need to aggregate entries into a fewer number of session filters. Be aware of overlapping IP addresses from different VPRNs in different roaming partner networks.

Reference session filters within AQP entries

The configured session filters need to be referenced within AQP entries, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 510 {
                            admin-state enable
                            description "apply FW rules for roaming partner 1"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "strict-FW-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                session-filter "restricted_access_partner1"
                            }
                        }
```

```
                        entry 511 {
                            admin-state enable
                            description "apply FW rules for roaming partner 2"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "strict-FW-Protection" {
                                    eq "OFF"
                                }
                            }
                            action {
                                session-filter "restricted_access_partner2"
                            }
                        }
```

Restrict downstream traffic (optional)

Operators can optionally restrict downstream traffic to specific destinations and protocols, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter "restricted_downstream_partner1" {
                    description "allow only traffic to only signed up partners"
                    default-action {
                        action deny
                    }
                    entry 10 {
                        description "allow GTP-U from authorized subnets"
                        match {
                            ip-protocol udp
                            dst-ip {
                                ip-prefix-list "Roaming1_ALL_SGWs"
                            }
                            dst-port {
                                eq 2152
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 11 {
                        description "allow GTP-C to authorized subnets"
                        match {
                            ip-protocol udp
                            dst-ip {
                                ip-prefix-list "Roaming1_ALL_SGWs"
                            }
                            dst-port {
                                eq 2123
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 20 {
                        description "allow DNS"
                        match {
                            ip-protocol tcp-udp
```

```
                            dst-ip {
                                ip-prefix-list "Roaming1_ALL_SGWs"
                            }
                        }
                        action {
                            permit
                        }
                    }
                }
                session-filter "restricted_downstream_partner2" {
                    description "SGWs_allowed_partner2"
                    default-action {
                        action deny
                        event-log "FW_events_log"
                    }
                    entry 10 {
                        description "allow GTP-U to authorized subnets"
                        match {
                            ip-protocol udp
                            dst-ip {
                                ip-prefix-list "Roaming2_ALL_SGWs"
                            }
                            dst-port {
                                eq 2152
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 11 {
                        description "allow GTP-C to authorized subnets"
                        match {
                            ip-protocol udp
                            dst-ip {
                                ip-prefix-list "Roaming2_ALL_SGWs"
                            }
                            dst-port {
                                eq 2123
                            }
                        }
                        action {
                            permit
                        }
                    }
                }
```

Note:
The preceding configuration provides the most flexibility and allows IP addresses to overlap between different partner networks. However, it comes at the cost of creating separate session filters for each partner. If IP addresses do not overlap, a single session filter is sufficient.

The session filters need to be referenced from AQP, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 514 {
```

```
                        admin-state enable
                        description "apply FW rules for roaming partner 1"
                        match {
                            traffic-direction network-to-subscriber
                            characteristic "strict-FW-Protection" {
                                eq "ON"
                            }
                        }
                        action {
                            session-filter "restricted_downstream_partner1"
                        }
                    }
                    entry 515 {
                        admin-state enable
                        description "apply FW rules for roaming partner 2"
                        match {
                            traffic-direction network-to-subscriber
                            characteristic "strict-FW-Protection" {
                                eq "OFF"
                            }
                        }
                        action {
                            session-filter "restricted_downstream_partner2"
                        }
                    }
```

Configuration to protect against malformed packets

It is always recommended in FW deployments that overload-drop, error-drop, and fragment-drop are enabled within the default sub-policy, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                    entry 50 {
                        admin-state enable
                        description "drop error and fragmented packets"
                        action {
                            error-drop {
                                event-log "FW_events_log"
                            }
                            fragment-drop {
                                drop-scope all
                                event-log "FW_events_log"
                            }
                            overload-drop {
                                event-log "FW_events_log"
                            }
                        }
```

Note:

- The overload-drop action ensures that AA-ISA, if it gets overloaded, drops the excess traffic instead of cutting it through without applying FW rules.
- The error-drop action ensures that AA-ISA drops malformed IP packets.

- The fragment-drop all action allows the operator to drop all fragmented traffic, drop out-of-order fragments only, or allow fragments through. Because many network DoS attacks use IP fragmentation to initiate attacks, allowing fragments through is not recommended for firewall deployments. As a minimum, if fragmentation is used, the operator is recommended to configure AA to drop out-of-order fragmented packets.
- The preceding actions are applied to all traffic. Therefore, there are no AQP match conditions configured.

Plausibility of GTP messages and GTP message validation

To protect the network from malformed GTP packets and associated attacks as described in the overview section, a GTP filter needs to be created and referenced from an AQP entry.

1. Configure the GTP filter object to:
   a. Enable GTP-C inspection so that the FW:
      i. Ensures the correct IE states and order
      ii. Discards any GTP packet that contains an invalid message type for S5/S8 (Gn/Gp) interface
   b. Enable sequence number checking for GTP-C traffic (for partner 1 traffic)
   c. Enable the GTP filter to check and drop errored GTP packets (anomalies)
   d. Enable GTP message length checking (to minimize exposure to code injection attacks). The maximum is set here (for example, 1250 bytes). The value is operator dependent, and should be replaced with the figure used by the operator.
   e. Drop GTP-in-GTP encapsulated packets

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-filter-partner1" {
                        description "gtp-filter for partner 1"
                        gtp-in-gtp deny
                        max-payload-length 1250
                        gtp-tunnel-database {
                            validate-gtp-tunnels true
                            validate-sequence-number true
                        }
                        log {
                            action deny
                            event-log "FW_events_log"
                        }
                    }
                    gtp-filter "gtp-filter-partner2" {
                        description "gtp-filter for partner 2"
                        gtp-in-gtp deny
                        max-payload-length 1250
                        log {
                            action deny
                            event-log "FW_events_log"
                        }
                    }
```

2. The configured GTP filters need to be referenced from AQP entries, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 512 {
                            admin-state enable
                            description "apply SGW GTP filter rules"
                            match {
                                characteristic "strict-FW-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                gtp-filter "gtp-filter-partner1"
                            }
                        }
                        entry 513 {
                            admin-state enable
                            description "apply SGW GTP filter rules"
                            match {
                                characteristic "strict-FW-Protection" {
                                    eq "OFF"
                                }
                            }
                            action {
                                gtp-filter "gtp-filter-partner2"
                            }
                        }
```

Filtering of GTP message types

In this configuration, traffic from partner 2 is considered "safe/trusted". Therefore, unlike traffic from partner 1, no additional GTP message-type filtering is applied to it, beyond the GTP Cat-1 (see Table 8: Allowed GTP message types (Cat-1)) message filtering applied as a result of enabling GTP-C inspection.

For roaming partner 1 traffic, the GTP filter is configured to block some Cat-1 optional message types (GTPv1 and GTPv2):

- GTPv2: Trace session activation/deactivation (this is optional for S8)
- GTPv1: Allows only the message types used by GTP-U and blocks GTPv1 message types used by GTP-C

Note:
By configuring GTP-C inspection, only Cat-1 message types (see Table 8: Allowed GTP message types (Cat-1)) are allowed and all others are denied. Therefore, there is little to no need for additional GTP message filtering configuration.

The GTP filter is configured as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
```

```
                    admin-state enable
                    gtpc-inspection true
                    log {
                        action deny
                        event-log "FW_events_log"
                    }
                    gtp-filter "gtp-filter-partner1" {
                        description "gtp-filter for partner 1"
                        gtp-in-gtp deny
                        max-payload-length 1250
                        gtp-tunnel-database {
                            validate-gtp-tunnels true
                            validate-sequence-number true
                        }
                        log {
                            action deny
                            event-log "FW_events_log"
                        }
                        message-type {
                            default-action deny
                            entry 1 {
                                action permit
                                value echo-request
                            }
                            entry 2 {
                                action permit
                                value echo-response
                            }
                            entry 3 {
                                action permit
                                value error-indication
                            }
                            entry 4 {
                                action permit
                                value supported-extension-headers-notification
                            }
                            entry 5 {
                                action permit
                                value end-marker
                            }
                            entry 6 {
                                action permit
                                value g-pdu
                            }
                        }
                        message-type-gtp-v2 {
                            default-action permit
                            entry 524 {
                                action deny
                                value trace-session-activation
                            }
                            entry 525 {
                                action deny
                                value trace-session-deactivation
                            }
                        }
                        imsi-apn-filter {
                            default-action deny
                            entry 1031 {
                                action permit
                                apn "ANY_APN"
                                imsi-mcc-mnc-prefix "161379"
                            }
                        }
```

```
                    }
```

TEID validation

For roaming partner 1, to protect the network resources from spoofed TEIDs, the FW is recommended to verify that the TEIDs used in the GTP-U traffic are valid (that is, correctly negotiated via GTP-C), as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-filter-partner1" {
                        gtp-tunnel-database
                            validate-gtp-tunnels true
```

Because roaming partner 2 network is trusted, no TEID validation is needed.

UE IP address anti-spoofing

It is a good practice to protect the network against UEs spoofing a different IP address, as follows:

```
configure exclusive
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-filter-partner1" {
                        gtp-tunnel-database
                            validate-gtp-tunnels true    # required for IP SA validation
                            validate-source-ip-addr true
```

This example applied to partner 1 traffic. Validation of source IP requires the use of a GTP tunnel database.

APN and IMSI filtering

In this example, only Home-Routed (HR) traffic from partner 1 is allowed, regardless of the APN. The rest is denied. This is achieved by configuring an IMSI prefix (="1613797") that corresponds to the Home network.

For roaming partner 2, MVNO traffic is allowed as well as HR traffic. This MVNO traffic (specific IMSI prefix = "1613400" in this example) is only allowed to attach to the mvnoguest.com APN, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
```

```
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-filter-partner1" {
                        imsi-apn-filter {
                            default-action deny
                            entry 1031 {
                                action permit
                                apn "ANY_APN"
                                imsi-mcc-mnc-prefix "161379"
                            }
                        }
                    }
                    gtp-filter "gtp-filter-partner2" {
                        imsi-apn-filter {
                            default-action deny
                            entry 1040 {
                                action permit
                                apn "mvnnoguest.com$"
                                imsi-mcc-mnc-prefix "161340"
                            }
                            entry 1041 {
                                action permit
                                apn "ANY_APN"
                                imsi-mcc-mnc-prefix "161379"
                            }
                        }
                    }
```

Limiting concurrent session creations

To further lower the risk of DoS attacks using massive amounts of session/PDN create messages, it is recommended that the operator configure the maximum concurrent number of endpoints (TEIDs) that an SGW can create.

In this example, the limit that is configured in the GTP filter corresponds to the maximum concurrent TEIDs that can be created by any SGW IP address, as follows:

```
configure exclusive
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtpc-inspection true
                    gtp-filter "gtp-filter-partner1" {
                        gtp-tunnel-database {
                            default-tunnel-endpoint-limit 400
```

Configuring FW statistics

To gain visibility into the traffic passing through the FW and the FW actions taken, it is highly recommended to enable "deny-admit" statistics, as follows:

```
configure {
    log {
        accounting-policy 5 {
```

```
            admin-state enable
            description "LogFileforAAFirewallAccounting"
            collection-interval 10
            record aa-admit-deny
            destination {
                file "logfile-5"
            }
        }
        file "logfile-5" {
            compact-flash-location {
                primary cf3
            }
        }
```

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                statistics {
                    aa-admit-deny {
                        accounting-policy 5
                        collect-stats true
                        gtp-filter-stats true
                        policer-stats true
                        policer-stats-resources true
                        session-filter-stats true
                    }
                    aa-protocol {
                        admin-state disable     # default
                    }
```

Configuring threshold crossing alerts

As well as admit-deny statistics, the operator can optionally enable the FW to generate Threshold Crossing Alerts (TCAs) against the collected statistics.

Note:
NSP also supports TCAs. The operator has a choice to enable TCAs on both the FW and NSP. The advantage of TCAs generated directly from the FW is that they tend to be more real-time relative to NSP TCAs. However, NSP supports a larger TCA scale than the FW.

The operator needs to set the low- and high-water marks according to the conditions of their networks. The following values are for illustration purposes only, for example, for roaming partner 1:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                threshold-crossing-alert {
                    criteria gtp-sanity-drop direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria message-type-default-action
 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
```

```
                    gtp-filter "gtp-filter-partner1" criteria header-sanity direction from-sub
{
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria imsi-apn-filter-default-action
direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria validate-gtp-tunnels direction
from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria validate-sequence-number
direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria validate-src-ip-addr direction
from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria missing-mandatory-ie direction
from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria gtp-tunnel-database-full
direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner1" criteria gtp-tunnel-endpoint-limit
direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter-entry "gtp-filter-partner1" entry-id 524 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter-entry "gtp-filter-partner1" entry-id 525 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
```

GTP tunnel validation must be enabled and the default tunnel endpoint limit must be set. The following error is raised when GTP tunnel validation is not enabled when the TCA watermarks are configured:

```
*[ex:/configure application-assurance group 1 partition 1 threshold-crossing-alert]
A:admin@PE-1# commit
MINOR: MGMT_CORE #5001: configure application-assurance group 1 partition 1 threshold-crossing-
alert gtp-filter "gtp-filter-partner2" criteria gtp-tunnel-database-full direction from-sub -
 gtp-filter does not have the matching config for this tca criteria - configure application-
assurance group 1 partition 1 gtp gtp-filter "gtp-filter-partner2" gtp-tunnel-database
 validate-gtp-tunnels
MINOR: MGMT_CORE #5001: configure application-assurance group 1 partition 1 threshold-crossing-
alert gtp-filter "gtp-filter-partner2" criteria gtp-tunnel-endpoint-limit direction from-sub
 - gtp-filter does not have the matching config for this tca criteria - configure application-
```

```
assurance group 1 partition 1 gtp gtp-filter "gtp-filter-partner2" gtp-tunnel-database default-
tunnel-endpoint-limit
```

The following TCA settings can be configured for roaming partner 2:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                threshold-crossing-alert {
                    gtp-filter "gtp-filter-partner2" criteria imsi-apn-filter-default-action
 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner2" criteria missing-mandatory-ie direction
 from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner2" criteria gtp-tunnel-database-full
 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
                    gtp-filter "gtp-filter-partner2" criteria gtp-tunnel-endpoint-limit
 direction from-sub {
                        high-watermark 100
                        low-watermark 60
                    }
```

Configuring GTP and GTP-C applications

By configuring AA app-filters to define GTP-U and GTP-C applications, the operator can gain further visibility into the volume of traffic of these applications, as follows:

```
configure exclusive
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    application "GTP_c" {
                    }
                    application "GTP_other" {
                    }
                    application "GTP_u" {
                    }
                    app-filter {
                        entry 40000 {
                            admin-state enable
                            application "GTP_u"
                            protocol {
                                eq "gtp"
                            }
                            server-port {
                                eq {
                                    port-number 2152
                                }
                            }
                        }
```

```
                        entry 40010 {
                            admin-state enable
                            application "GTP_c"
                            protocol {
                                eq "gtp"
                            }
                            server-port {
                                eq {
                                    port-number 2123
                                }
                            }
                        }
                        entry 40020 {
                            admin-state enable
                            application "GTP_other"
                            protocol {
                                eq "gtp"
                            }
                        }
                    }
```

The export and display of statistics related to these applications use standard AA per application per partition or per application per subscriber XML records and **show** routines.

Relevant debug routines

CLI show routines

```
[/]
A:admin@PE-1# show application-assurance group 1:1 session-filter "restricted_access_partner1"

===============================================================================
AA Session Filter Instance "restricted_access_partner1"
===============================================================================
Description    : SGWs_allowed_partner1
Default Action : deny
    Event Log  : (Not Specified)
AQP Entries    :
          510
-------------------------------------------------------------------------------
Filter Match Criteria
-------------------------------------------------------------------------------
Entry          : 10
Description    : allow GTP-U from authorized subnets
IP Protocol    : udp
Source IP List : Roaming1_ALL_SGWs
Dest Port      : eq 2152
Action         : permit
    Event Log  : (Not Specified)
Hits           : 1 flows
-------------------------------------------------------------------------------
Entry          : 11
Description    : allow GTP-C from authorized subnets
IP Protocol    : udp
Source IP List : Roaming1_ALL_SGWs
Dest Port      : eq 2123
Action         : permit
    Event Log  : (Not Specified)
Hits           : 2 flows
```

```
-------------------------------------------------------------------------------
Entry          : 20
Description    : allow DNS
IP Protocol    : *
Source IP List : Roaming1_ALL_SGWs
Dest Port      : eq 53
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
No. of entries   : 3
===============================================================================
```

```
[/]
A:admin@PE-1# show application-assurance group 1:1 gtp gtp-filter gtp-filter-name "gtp-filter-
partner1"

===============================================================================
Application Assurance Group 1:1 GTP Filter "gtp-filter-partner1"
===============================================================================
Description                 : gtp-filter for partner 1
Maximum payload length      : 1250
Event log                   : FW_events_log
Event log action            : deny
Default action              : deny
Default GTPv2 action        : permit
Default IMSI-APN action     : deny
GTP in GTP action           : deny
Validate GTP tunnels        : enabled
Validate sequence number    : enabled
Validate source IP address  : enabled
GTP tunnel endpoint limit   : 400
Configured messages         : 6
Configured GTPv2 messages   : 2
Configured IMSI-APN filters : 1
Packets arrived             : 18
Packets denied
  Payload length            : 0
  Message type              : 0
  GTPv2 message type        : 2
  IMSI-APN filter           : 0
  Mandatory header          : 0
  Extension header          : 0
  Information element       : 0
  Invalid TEID              : 0
  Invalid sequence number   : 0
  Invalid source IP address : 0
  Missing mandatory IE      : 0
  GTP in GTP                : 0
  No tunnel resource        : 0
  Tunnel endpoint limit     : 0
Packets permitted           : 16
===============================================================================
```

```
[/]
A:admin@PE-1# show application-assurance group 1:1 gtp

===============================================================================
Application Assurance Group 1:1 GTP
===============================================================================
Admin status      : Up
Event log         : FW_events_log
```

```
Event log action  : deny
Mode              : filtering
GTP-C inspection  : Enabled

-------------------------------------------------------------------------------
GTP Statistics                              sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Incoming packets                                 9                     9
Packets denied
  UDP packet length                              0                     0
  GTP message length                             0                     0
  GTP version                                    0                     0
-------------------------------------------------------------------------------
Packets permitted                                9                     9
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
GTP Policing Statistics                     sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                  9                     9
Packets denied
  gtp-traffic flow-count policer                 0                     0
  Other                                          0                     0
-------------------------------------------------------------------------------
Packets permitted                                9                     9
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
GTP Filter Statistics                       sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                  9                     9
Packets denied                                   1                     1
Packets permitted
  gtp-filter                                     8                     8
  no gtp-filter                                  0                     0
-------------------------------------------------------------------------------
Total GTP packets permitted                      8                     8
===============================================================================
```

```
[/]
A:admin@PE-1# show application-assurance group 1 aa-sub-list

===============================================================================
Application-Assurance Subscriber List for Group 1
===============================================================================
type      aa-sub                    ISA       App-Profile              divert
                                    assigned
-------------------------------------------------------------------------------
group 1:1
-------------------------------------------------------------------------------
sap       1/1/9:100                 1/2       default                  Yes
sap       1/1/9:201                 1/2       strict-FW                Yes
transit   1_10.10.68.1/32           1/2       strict-FW                Yes
-------------------------------------------------------------------------------
Number of aa-subs found in group 1:1              : 3
Total number of aa-subs found                     : 3
===============================================================================
```

CLI tools dump routines

```
[/]
A:admin@PE-1# tools dump application-assurance group 1:1 flow-record-search isa 1/2

=====================================================
Application-Assurance flow record search
Search Start Time:     "08/15/2025 19:04:37" (UTC)
 Search Criteria:
  group[:partition]:   1:1
  isa:                 1/2
  protocol name:       none specified
  application name:    none specified
  app-group name:      none specified
  flow-status:         none specified
  start-flowId:        none specified
  classified:          none specified
  server-ip:           none specified
  server-port:         none specified
  client-ip:           none specified
  bytes-tx:            none specified
  flow-duration:       none specified
  max-count:           none specified
  flow-modified:       none specified
  search-type:         default
=====================================================
FlowId     Init  Src-ip           Dst-ip           Ip-prot     Src-prt Dst-prt
 Protocol               Application
   Pkts-tx    Bytes-tx             Pkts-disc  Bytes-disc
      Time-ofp(UTC)         Time-olp(UTC)
2          no    10.10.68.3       10.10.68.1       udp         2123    2123
 "gtp"                  "GTP_c"
    1          46                   0          0
      "08/15/2025 18:54:50" "08/15/2025 18:54:50"
3          yes   10.10.68.1       10.10.68.3       udp         2123    2123
 "gtp"                  "GTP_c"
    2          359                  0          0
      "08/15/2025 18:54:50" "08/15/2025 18:54:50"
4          yes   10.10.68.3       10.10.68.1       udp         2123    2123
 "gtp"                  "GTP_c"
    2          326                  1          51
      "08/15/2025 18:54:50" "08/15/2025 18:54:55"
7          yes   10.10.68.1       10.10.68.3       udp         63760   2152
 "gtp"                  "GTP_u"
    5          500                  0          0
      "08/15/2025 18:54:50" "08/15/2025 18:54:50"
8          yes   10.10.68.3       10.10.68.1       udp         64784   2152
 "gtp"                  "GTP_u"
    5          500                  0          0
      "08/15/2025 18:54:50" "08/15/2025 18:54:50"
11         yes   10.10.68.1       10.10.68.3       udp         2123    2123
 "gtp"                  "GTP_c"
    1          136                  1          91
      "08/15/2025 18:54:50" "08/15/2025 18:54:50"
SEARCH COMPLETED.
Search End Time:     "08/15/2025 19:04:37" (UTC)
Total Records:       6
=====================================================
```

```
[/]
A:admin@PE-1# tools dump application-assurance group 1:1 admit-deny-stats
```

```
===============================================================================================================
===============================================
Application-Assurance Group 1:1 Admit-Deny Statistics
===============================================================================================================
===============================================
-------------------------------------------------------------------------------------------------------
-----------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
Packet Validation Statistics
(Packets)              (Packets)          (Packets)              (Packets)
-------------------------------------------------------------------------------------------------------
-----------------------------------------------
Error
0                      0                    0                    0
Fragments: Out-Of-Order
0                      0                    0                    0
Fragments: All
0                      0                    0                    0
Overload
N/A                    0                  N/A                    0
GTP Sanity
9                      0                    9                    0
-------------------------------------------------------------------------------------------------------
-----------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
GTP Filter Statistics
(Packets)              (Packets)          (Packets)              (Packets)
-------------------------------------------------------------------------------------------------------
-----------------------------------------------
GTP Filter: gtp-filter-partner1
 Entry: 1    echo-request
0                      0                    0                    0
 Entry: 2    echo-response
0                      0                    0                    0
 Entry: 3    error-indication
0                      0                    0                    0
 Entry: 4    supported-extension-headers-notification
0                      0                    0                    0
 Entry: 5    end-marker
0                      0                    0                    0
 Entry: 6    g-pdu
5                      0                    5                    0
 Message Type Default Action
0                      0                    0                    0
 GTPv2 Entry: 524  trace-session-activation
0                      1                    0                    0
 GTPv2 Entry: 525  trace-session-deactivation
0                      0                    0                    1
 GTPv2 Message Type Default Action
3                      0                    3                    0
 IMSI-APN Entry: 1031
N/A                    0                  N/A                    0
 IMSI-APN Filter Default Action
N/A                    0                  N/A                    0
 Max Payload Length
N/A                    0                  N/A                    0
 Message Type Header Sanity
N/A                    0                  N/A                    0
 Invalid TEID
N/A                    0                  N/A                    0
 Invalid Sequence Number
```

```
N/A                  0                   N/A                   0
 Invalid Source IP Address
N/A                  0                   N/A                   0
 Missing Mandatory IEs
N/A                  0                   N/A                   0
 GTP in GTP Action
N/A                  0                   N/A                   0
 GTP Tunnel DB Resource
N/A                  0                   N/A                   0
 Tunnel Endpoint Limit
N/A                  0                   N/A                   0
GTP Filter: gtp-filter-partner2
 Message Type Default Action
0                    0                     0                   0
 GTPv2 Message Type Default Action
0                    0                     0                   0
 IMSI-APN Entry: 1040
N/A                  0                   N/A                   0
 IMSI-APN Entry: 1041
N/A                  0                   N/A                   0
 IMSI-APN Filter Default Action
N/A                  0                   N/A                   0
 Max Payload Length
N/A                  0                   N/A                   0
 Message Type Header Sanity
N/A                  0                   N/A                   0
 Invalid TEID
N/A                  0                   N/A                   0
 Invalid Sequence Number
N/A                  0                   N/A                   0
 Invalid Source IP Address
N/A                  0                   N/A                   0
 Missing Mandatory IEs
N/A                  0                   N/A                   0
 GTP in GTP Action
N/A                  0                   N/A                   0
 GTP Tunnel DB Resource
N/A                  0                   N/A                   0
 Tunnel Endpoint Limit
N/A                  0                   N/A                   0
-------------------------------------------------------------------------------------------------
-------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
Session Filter Statistics
(Sessions)            (Packets)           (Sessions)            (Packets)
-------------------------------------------------------------------------------------------------
-------------------------------------------
Session Filter: restricted_access_partner1
 Entry: 10
1                    0                     0                   0
 Entry: 11
2                    0                     0                   0
 Entry: 20
0                    0                     0                   0
 Default Action
0                    0                     0                   0
Session Filter: restricted_access_partner2
 Entry: 10
0                    0                     0                   0
 Entry: 11
0                    0                     0                   0
 Entry: 20
0                    0                     0                   0
```

```
 Default Action
0                   0                   0                   0
Session Filter: restricted_downstream_partner1
 Entry: 10
0                   0                   1                   0
 Entry: 11
0                   0                   1                   0
 Entry: 20
0                   0                   0                   0
 Default Action
0                   0                   0                   0
Session Filter: restricted_downstream_partner2
 Entry: 10
0                   0                   0                   0
 Entry: 11
0                   0                   0                   0
 Default Action
0                   0                   0                   0
-------------------------------------------------------------------------------------------------
-------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
Flow Policer Statistics
(Flows)             (Flows)             (Flows)             (Flows)
-------------------------------------------------------------------------------------------------
-------------------------------------------
Subscriber Flow Count Policers
   limit_roamingSGW_Flows
0                   0                   0                   0
   limit_roamingSGWs_total_Flows
0                   0                   0                   0
-------------------------------------------------------------------------------------------------
-------------------------------------------
```

Conclusion

A 3GPP roaming interface using GTP presents a security risk to mobile access networks. The AA GTP stateful firewall protects the network infrastructure from untrusted roaming partner networks.

Application Assurance — Security Gateway Stateful Firewall

This chapter provides information about Application Assurance Security gateway stateful firewall.

Topics in this chapter include:

Applicability

This chapter is applicable to all chassis supporting Application Assurance (AA).

The chapter was initially written based on SR OS Release 13.0.R2, but the MD-CLI in the current edition is based on SR OS Release 25.3.R2.

Overview

The SR OS AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW). The feature provides protection for mobile infrastructure: Mobility Management Entities (MMEs), Serving Gateways (SGWs), and Network Management Systems (NMSs), against attacks from compromised base stations, evolved NodeBs (eNBs), or Femto Access Points (FAPs). AA stateful packet filtering, combined with AA layer 7 classification and control, provides advanced, next-generation firewall functionality. Using stateful packet filtering, the AA FW not only inspects packets at layers 3 to 7, but also monitors the connection state.

AA FW deployed within a SeGW in ultra-broadband access networks (3G, 4G, or Femto) provides back-end core network security protection, as per Figure 19: LTE SeGW firewall deployment. AA FW offers protection for the following 3rd Generation Partnership Project (3GPP) defined interfaces:

- S1-MME
- S1-U
- Operations, Administration, and Maintenance (OAM)

Figure 19: LTE SeGW firewall deployment

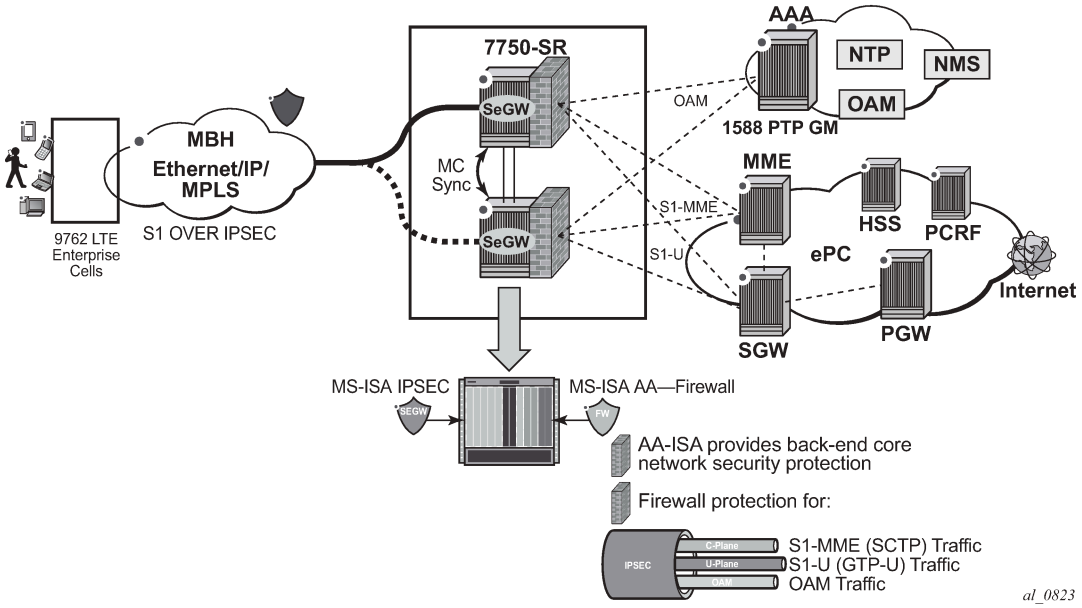

Similarly, the SeGW architecture for Femto deployment is based on two SR OS systems terminating the mobile backhaul side (front-end and connecting to, for example, a base station router gateway and other network elements of the packet core (back-end), as per Figure 20: SeGW in small cells architecture:

Figure 20: SeGW in small cells architecture

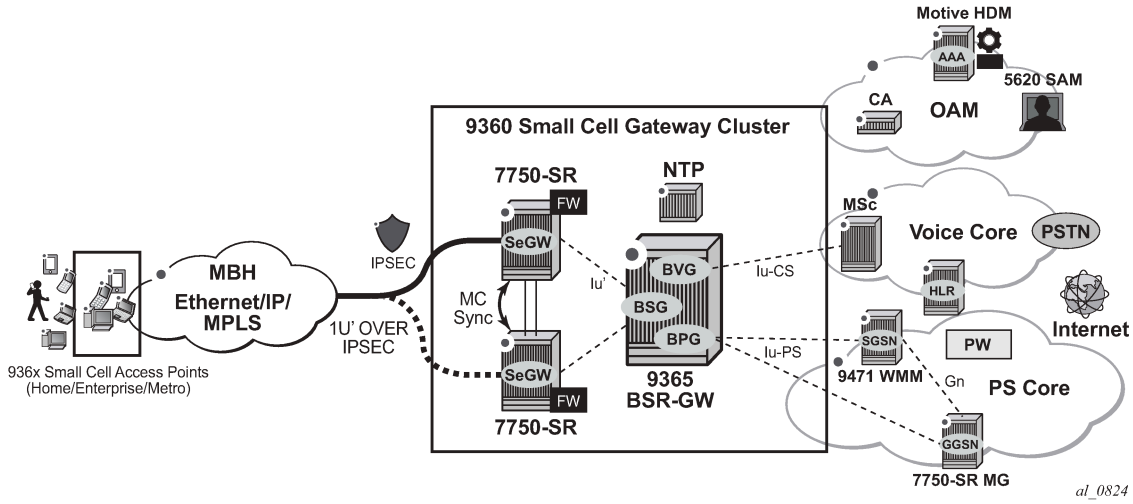

The two SeGWs run in stateful redundant mode: upon partial or total failure of the active SeGW for a set of IPSec tunnels, the other SeGW takes over without terminating the IPSec tunnels, providing hitless failover.

In addition to MS-ISA hardware dedicated to the IPSec function, each SeGW supports one or more additional MS-ISAs running AA to provide firewall capabilities. The firewall rules protect the BSR as well as the BSR-GW and packet core network elements (NEs) from malicious attacks or unauthorized traffic.

The objective of this chapter is to describe the required configuration within AA-ISA to enable AA FW and protection for S1-MME, S1-U, and OAM traffic. Basic knowledge of AA-ISA diversion configuration is assumed.

S1-MME traffic protection

The purpose of AA FW in this deployment is to protect the MME infrastructure against an attack from a compromised eNB or FAP. Network flooding attacks, malformed packets, and port scans are examples of denial of service (DoS) attacks that can be carried out using a compromised eNB or FAP.

AA FW provides inspection of the Stream Control Transmission Protocol (SCTP) used to communicate to the MME. Such inspection includes checking for SCTP payload protocol IDs (PPIDs), source and destination ports, SCTP chunk validation, and malformed SCTP packets, such as checksum validation. In addition, the operator can configure DoS flooding rules, such as policers to limit the bandwidth and flow counts of SCTP traffic.

S1-U traffic protection

The purpose of AA FW in this deployment is to protect the SGW infrastructure against an attack from a compromised eNB or FAP. AA FW supports protection against:

- malformed GPRS Tunneling Protocol User plane (GTP-U) packet attacks

  Checking packet sanity, which include GTP-U mandatory, optional, and extension header checks, as well as checks for invalid reserved information elements (IE) and missing IEs.
- unsupported GTP messages

  Filtering messages based on message type and message length.
- flooding attacks

  Shaping GTP traffic bandwidth, which limits the GTP-U bandwidth that a FAP can send to the core (SGW).

  Limiting GTP tunnels, which limits the number of concurrent GTP tunnels and setup rate of these tunnels from a FAP to the core network.

  To prevent the shared resources of bandwidth and the SGW processor from being consumed by an attacker, Nokia recommends the GTP flow rate limiting configuration.
- IP fragmentation-based attacks

  Applying various drop rules for IP fragmentation of GTP messages.

OAM traffic protection

The purpose of AA FW protection in this deployment is to protect against any abuse of OAM network resources, such as NMS.

Network flooding attacks, malformed packets, and port scans are examples of such attacks that can be carried out using a compromised eNB or FAP.

See the configuration described in the Application Assurance — Stateful Firewall chapter for this context of OAM protection in SeGW.

Configuration

AA-ISA Application QoS Policies (AQPs) can contain AQP actions that provide SCTP and GTP filtering functionality. As with all AQPs, these actions have partition-level scope, which allows different FW policies to be implemented by using AA partition concepts within the same AA-ISA.

The configuration topology in Figure 21: Configuration topology shows the SeGW FW functionality of S1-U and S1-MME interfaces. Geo-redundancy, which is a very common deployment option, is not described in this chapter because it is described in the Multi-Chassis IPSec Redundancy chapter.

Figure 21: Configuration topology

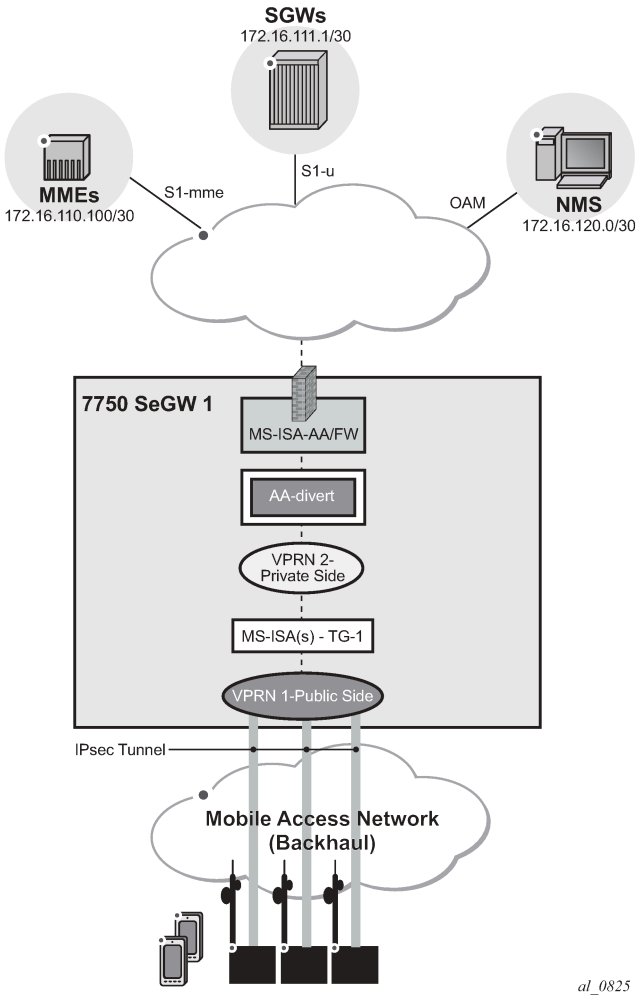

Initial setup with multi-chassis IPSec redundancy

Tunnel ISAs are configured with optional multi-chassis redundancy. See the Multi-Chassis IPSec Redundancy chapter for more information.

The following configuration steps are described in the remainder of this section:

- 1 Divert AA traffic and apply basic firewall rules
- 2 Configure AA-ISA to provide firewall protection to protect MMEs (S1-AP traffic)
- 3 Configure AA-ISA to protect SGW (GTP-U traffic)
- 4 Configure AA-ISA to protect NMS (OAM traffic)

1 Divert AA traffic and apply basic firewall rules

In this section, the following steps are described:

- 1.1 Divert private VPRN traffic into AA-ISA with AA multi-chassis redundancy
- 1.2 Protect against malformed packets
- 1.3 Limit total traffic from any eNB

1.1 Divert private VPRN traffic into AA-ISA with AA multi-chassis redundancy

This step is required for any of the configurations in steps 2, 3, or 4. There is no dependency between steps 2, 3, or 4.

In this example, one private VPRN is used for all traffic to and from eNBs. In some small cell deployments, eNB traffic is split into three different VPRNs: one for control (S1-MME), one for management (OAM), and one for bearer traffic (S1-U GTP-U). In that case, each of these VPRNs needs to be diverted into AA-ISA to provide firewall protection.

First, define an application profile and transit IP policy, such as:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-profile "default" {
                        description "App profile that applies to the whole SAP"
                        divert true
                    }
                }
                transit-ip-policy 1 {
                    description "Per eNB-IP Sub policy"
                    detect-seen-ip true
                    transit-auto-create {
                        admin-state enable
                    }
                }
            }
```

Then, apply the application profile and the transit IP policy to the SAP on the private side of the IPSec tunnel ISA:

```
configure {
    service {
        vprn "VPRN-2" {
            interface "int-IPsec-Private-1" {
                tunnel true
```

```
                sap tunnel-1.private:1 {
                    app-profile "default"
                    transit-policy {
                        ip 1
                    }
                }
```

This configuration achieves:

- Traffic to or from the IPSec tunnel ISA private SAP is diverted to AA-ISA for the purpose of FW protection
- Within AA-ISA, the diverted SAP is treated as a parent SAP. That is, instead of treating the whole SAP as a single subscriber, subscribers are auto-created within this SAP based on the IP address of the eNBs

1.2 Protect against malformed packets

In firewall deployments, it is recommended that **overload-drop**, **error-drop**, and **fragment-drop** are enabled within the default sub-policy, as shown in the following example:

- **overload-drop** ensures that AA-ISA, when overloaded, drops the excess traffic instead of allowing it through, without applying firewall rules.
- **error-drop** ensures that AA-ISA drops malformed IP packets.
- **fragment-drop**: because many network DoS attacks use IP fragmentation to initiate attacks, the operator has the option to drop all fragmented traffic, drop out-of-order fragments only, or allow fragments through. Allowing fragments through is not recommended for firewall deployments.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "apply SeGW session filter rules"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                error-drop {
                                }
                                fragment-drop {
                                    drop-scope all
                                }
                                overload-drop {
                                }
                            }
                        }
                    }
```

1.3 Limit total traffic from any eNB

Nokia recommends that a total limit be placed on how much bandwidth and how many flows an eNB or FAP can generate toward the network, regardless of the type of traffic.

The limit values are a function of the number of end devices that are served by the eNB or FAP, plus some additional margin:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "limit_eNBs_total_Flows" {
                    granularity subscriber
                    peak-flow-count 1000
                }
                single-bucket-bandwidth-policer "limit_eNBs_total_bw" {
                    granularity subscriber
                    mbs 500
                    pir 500
                }
```

Note:
If the traffic from eNB or FAP is separated into different private SAPs, based on traffic type (S1-AP, S1-U, or OAM), as with some deployment topologies, then the policing limit value is dependent on the SAP traffic type as well as the number of end devices. See policing limit settings in steps 2 and 3.

Apply the configured policers as actions from within the default sub-policy AQP entry:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "apply SeGW session filter rules"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "limit_eNBs_total_bw"
                                }
                                error-drop {
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_eNBs_total_Flows"
                                }
                                fragment-drop {
                                    drop-scope all
                                }
                                overload-drop {
                                }
                            }
                        }
                    }
```

Note:
All of the preceding listed actions use the traffic direction of subscriber-to-network. That is, they are not applied to traffic in the other direction (downstream) because the purpose of the firewall is to protect the network resources from upstream traffic coming from compromised eNBs or FAPs.

2 Configure AA-ISA to provide firewall protection to protect MMEs (S1-AP traffic)

The following steps are described in this section:

2.1 Create IP AA lists

First, create an AA IP prefix list that contains eNB IP addresses or range of addresses:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list "ALL_eNBs" {
                    description "eNodeB subnet"
                    prefix 172.16.100.0/24 {
                    }
                }
```

Optionally, create an AA IP list that contains MME IP addresses (in case there are more than one):

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list "MMEs" {
                    description "MMEs subnet"
                    prefix 172.16.110.100/30 {
                    }
                }
```

After the preceding lists are created, they can be referenced and used in AA FW rules using session filters and AQPs.

2.2 Allow only SCTP traffic toward MMEs — no port scanning

A basic setup creates session-filter rules that only allow SCTP traffic between eNBs and MMEs.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
```

```
                session-filter "SeGW_FW" {
                    default-action {
                        action deny
                    }
                    entry 1 {
                        description "allow SCTP to MMEs"
                        match {
                            ip-protocol sctp
                            dst-ip {
                                ip-prefix-list "MMEs"
                            }
                            dst-port {
                                eq 6005
                            }
                            src-ip {
                                ip-prefix-list "ALL_eNBs"
                            }
                        }
                        action {
                            permit
                        }
                    }
```

Note:
In the preceding configuration, SCTP traffic on MMEs is assumed to be running on port 6005.

The newly created session filter needs to be referenced from a default sub-policy AQP action, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "apply SeGW session filter rules"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                session-filter "SeGW_FW"
                                bandwidth-policer {
                                    single-bucket "limit_eNBs_total_bw"
                                }
                                error-drop {
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_eNBs_total_Flows"
                                }
                                fragment-drop {
                                    drop-scope all
                                }
                                overload-drop {
                                }
                            }
                        }
```

Using traffic direction **subscriber-to-network** in the preceding AQP entry achieves two objectives:

- Protecting MMEs by allowing only SCTP traffic to be initiated from eNB subnets toward MMEs. Port scanning toward MME is blocked.

- Allowing MMEs to have full access to eNBs.

Note:
It is important that an AQP, containing a session filter action, does not contain any matching condition other than ASOs, traffic direction, or subscriber ID. Subscriber ID is not applicable in this deployment use case.

2.3 DoS protection — limit the number of SCTP flows from eNBs

In this step, the operator configures a flow count policer to limit the number of SCTP flows that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to set up many SCTP flows.

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "sctp_flow_count" {
                    granularity subscriber
                    peak-flow-count 2
                }
```

In the preceding configuration, an eNB or FAP can have up to two flows at a time. In practice, there should only be one SCTP session, one flow in each direction, per eNB-MME pair. The preceding example uses two flows to leave a margin in case a second, backup, MME needs to communicate with the eNB, while still providing enough protection.

Add the defined policer as a **flow-count-limit-policer** AQP action, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 100 {
                            admin-state enable
                            description "limit SCTP traffic"
                            match {
                                traffic-direction subscriber-to-network
                                ip-protocol {
                                    eq sctp
                                }
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "sctp_flow_count"
                                }
                            }
                        }
```

Configure an AA FW events log

It is sometimes advisable to configure a log that captures events related to various AA FW actions. Because of the limited size of the log and the large amount of traffic AA can handle, consider the usefulness of the information in the log when:

- debugging a configuration
- testing a configuration in a staged environment
- capturing infrequent actions

The following configures a log called "FW_drops_log":

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                event-log "FW_drops_log" {
                    admin-state enable
                    buffer-type circular
                    max-entries 100000
                }
```

The following adds the configured log to the **default-action** of the session filter:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter "SeGW_FW" {
                    default-action {
                        action deny
                        event-log "FW_drops_log"    # added
                    }
                    entry 1 {
                        description "allow SCTP to MMEs"
                        match {
                            ip-protocol sctp
                            dst-ip {
                                ip-prefix-list "MMEs"
                            }
                            dst-port {
                                eq 6005
                            }
                            src-ip {
                                ip-prefix-list "ALL_eNBs"
                            }
                        }
                        action {
                            permit
                        }
                    }
```

The following **tools** command shows the log:

```
[/]
A:admin@SeGW-2# tools dump application-assurance group 1:1 event-log "FW_drops_log" isa 1/2

=================================================
Application-Assurance event-log "FW_drops_log"
```

```
Current Time:          "06/17/2025 13:48:34" (UTC)
  group[:partition]:   1:1
  isa:                 1/2
  admin state:         no shutdown
  buffer-type:         circular
  max-entries:         100000
=====================================================
Event-source             Action        SubType     SubName

    Direction Src-ip  Dst-ip  Ip-protocol Src-port Dst-port Timestamp

Total Records:   0
=====================================================
```

The following command clears all the entries within the specified log:

```
clear application-assurance group 1:1 event-log "FW_drops_log"
```

2.4 DoS protection — limit the SCTP bandwidth from eNB

Similar to the previous step, the operator configures a flow bandwidth policer to limit the amount of SCTP traffic that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to flood the MMEs.

```
configure {
    application-assurance {
        group 1 {
            policer {
                single-bucket-bandwidth-policer "sctp_bw_limit" {
                    granularity subscriber
                    mbs 10
                    pir 30
                }
```

In the preceding example, a single leaky-bucket policer is configured with a rate set to 30 kb/s and maximum burst size of 10 kbytes. This provides enough bandwidth to ensure normal operations, while still providing a ceiling limit of how much traffic any eNB can send toward the MMEs.

The value for this policer is a function of the amount of user equipment (UEs) served by the eNB/FAP. For example, in a small cell deployment, with 32 active users per FAP, the S1-MME bandwidth is estimated to be:

- Uplink — toward MME : 2.7 kb/s
- Downlink — from MME toward FAP : 28 kb/s

The following command adds the defined policer as a subscriber policy:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 100 {
                            admin-state enable
                            description "limit SCTP traffic"
                            match {
                                traffic-direction subscriber-to-network
```

```
                            ip-protocol {
                                eq sctp
                            }
                        }
                        action {
                            bandwidth-policer {
                                single-bucket "sctp_bw_limit"  # added
                            }
                            flow-count-limit-policer {
                                policer-name "sctp_flow_count"
                            }
                        }
                    }
```

Configure additional limits for all traffic to MMEs

To further protect the MMEs from a distributed attack, whereby a number of eNBs or FAPs are compromised, an AA FW can be configured to limit total traffic, not just from a single eNB as described in previous sections, but from all eNBs toward the MMEs.

It is recommended to configure the following three protection limits:

- total bandwidth of SCTP toward MMEs
- total number of flows toward MMEs
- flow setup rate toward the MMEs

The configuration is as follows:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-setup-rate-policer "limit_total_sctp_flows_rate" {
                    granularity system
                    flow-setup-rate-burst-size 100
                    peak-flow-setup-rate 100
                }
                flow-count-limit-policer "limit_total_sctp_flows" {
                    granularity system
                    peak-flow-count 400
                }
                single-bucket-bandwidth-policer "limit_total_sctp_bw" {
                    granularity system
                    mbs 100
                    pir 1200
                }
```

Note:

- The policers are of type **system** and not **subscriber** to be applied to all eNBs or FAPs, as is the case when auto-transit subscribers are created (see 1 Divert AA traffic and apply basic firewall rules).
- The actual limits of these policers are a function of the total number of eNBs served by the SeGW. In the preceding configuration, it is assumed that there are 400 eNBs. Therefore, the total limit is 400 flows of SCTP traffic.
- A flow setup rate limit of 100 is set to protect MMEs from a storm of new SCTP flows.

The policers are then referenced from within the appropriate AQP entry that matches the MMEs traffic and SCTP:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 110 {
                            admin-state enable
                            description "limit system traffic towards MMEs"
                            match {
                                traffic-direction subscriber-to-network
                                dst-ip {
                                    eq {
                                        ip-prefix-list "MMEs"
                                    }
                                }
                                src-ip {
                                    eq {
                                        ip-prefix-list "ALL_eNBs"
                                    }
                                }
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "limit_total_sctp_bw"
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_total_sctp_flows"
                                }
                                flow-setup-rate-policer {
                                    policer-name "limit_total_sctp_flows_rate"
                                }
                            }
                        }
```

Note:
It is possible, but redundant, to add the **ip-protocol eq sctp** command as a match condition, because the configured session filter already ensures that only SCTP traffic can flow between eNBs and MMEs.

2.5 Allow only specified SCTP PPIDs toward the MMEs

In this step, the operator blocks all except the specified SCTP messages that contain configured PPIDs, using an AA SCTP filter configuration:

```
*[ex:/configure application-assurance group 1 partition 1
                                                sctp-filter "SCTP-PPID-Filter"]
A:admin@SeGW-2# ?

 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 description           - Text description
 event-log             - Event log for packets dropped by the SCTP filter
 ppid                  + Enter the ppid context
 ppid-range            + Enter the ppid-range context
```

The filter specifies either a range of PPIDs or individual PPIDs.

```
*[ex:/configure application-assurance group 1 partition 1 sctp-filter "SCTP-PPID-Filter" ppid
 entry 1]
A:admin@SeGW-2# ?

 Immutable fields      - value

 action                ^ SCTP filter PPID entry action
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 value                 ^ PPID entry value
```

The PPIDs can be specified either by their values or by names. Names are specified in RFC 4960. See Table 9: SCTP PPIDs .

Table 9: SCTP PPIDs

Value	SCTP PPID	Value	SCTP PPID
0	Reserved by SCTP	31	Service Area Broadcast Protocol (SABP)
1	IUA	32	Fractal Generator Protocol (FGP)
2	M2UA	33	Ping Pong Protocol (PPP)
3	M3UA	34	CalcApp Protocol (CALCAPP)
4	SUA	35	Scripting Service Protocol (SSP)
5	M2PA	36	NetPerfMeter Protocol Control Channel (NPMP-CONTROL)
6	V5UA	37	NetPerfMeter Protocol Data Channel (NPMP-DATA)
7	H.248	38	Echo (ECHO)
8	BICC/Q.2150.3	39	Discard (DISCARD)
9	TALI	40	Daytime (DAYTIME)
10	DUA	41	Character Generator (CHARGEN)
11	ASAP	42	3GPP RNA
12	ENRP	43	3GPP M2AP
13	H.323	44	3GPP M3AP
14	Q.IPC/Q.2150.3	45	SSH over SCTP
15	SIMCO <draft-kiesel-midcom-simco-sctp-00.txt>	46	Diameter in a SCTP DATA chunk
16	DDP Segment Chunk	47	Diameter in a DTLS/SCTP DATA chunk

Value	SCTP PPID	Value	SCTP PPID
17	DDP Stream Session Control	48	R14P. BER Encoded ASN.1 over SCTP
18	S1 Application Protocol (S1AP)	49	Unassigned
19	RUA	50	WebRTC DCEP
20	HNBAP	51	WebRTC String
21	ForCES-HP	52	WebRTC Binary Partial (deprecated)
22	ForCES-MP	53	WebRTC Binary
23	ForCES-LP	54	WebRTC String Partial (deprecated)
24	SBc-AP	55	3GPP PUA
25	NBAP	56	WebRTC String Empty
26	Unassigned	57	WebRTC Binary Empty
27	X2AP	58-4294967295	Unassigned
28	IRCP - Inter Router Capability Protocol		
29	LCS-AP		
30	MPICH2		

Nokia recommends to limit the SCTP traffic to only those packets with S1 AP PPID. The SCTP filter can be configured to deny all by default and only allow PPID S1 AP (by value 18 or by name *s1-application-protocol*) as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                sctp-filter "SCTP-PPID-Filter" {
                    description "Allow only S1AP PPID"
                    event-log "FW_drops_log"
                    ppid {
                        default-action deny
                        entry 1 {
                            action permit
                            value s1-application-protocol
                        }
                    }
                }
```

This configured SCTP filter is then referenced as an action from within an AQP entry:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
```

```
                    app-qos-policy {
                        entry 100 {
                            admin-state enable
                            description "limit SCTP traffic"
                            match {
                                traffic-direction subscriber-to-network
                                ip-protocol {
                                    eq sctp
                                }
                            }
                            action {
                                sctp-filter "SCTP-PPID-Filter"  # added
                                bandwidth-policer {
                                    single-bucket "sctp_bw_limit"
                                }
                                flow-count-limit-policer {
                                    policer-name "sctp_flow_count"
                                }
                            }
                        }
```

The following command shows the packets allowed or denied by the configured SCTP filter:

```
[/]
A:admin@SeGW-2# show application-assurance group 1:1 sctp-filter "SCTP-PPID-Filter"

===============================================================================
Application Assurance Group 1:1 SCTP Filter "SCTP-PPID-Filter"
===============================================================================
Description           : Allow only S1AP PPID
Maximum PPID          : 4294967295
Minimum PPID          : 0
Default action        : deny
Configured PPIDs      : 1

Packets arrived       : 0
Packets denied
  Malformed packet    : 0
  PPID out of range   : 0
  PPID denied         : 0
Packets permitted     : 0
===============================================================================
```

Note:
The SCTP malformed packet counter shown above increments when an AA SCTP filter encounters an SCTP packet that is malformed, such as:

- IP packet is too small to contain a common SCTP header
- SCTP chunk LEN < 4 bytes: each SCTP chunk header is 4 bytes, so the SCTP chunk cannot be smaller than this
- remaining space in the IP packet is too small to contain a chunk header (for example, your packet has 2 chunks and the 2nd chunk length goes beyond the IP length advertised)
- IP packet is too small to contain the chunk

The SCTP filter statistics cannot be reset while processing without disabling the SCTP filter.

Another way to view the effect of the configured SCTP filter is to check the firewall log, if configured:

```
[/]
```

```
A:admin@SeGW-2# tools dump application-assurance group 1:1 event-log "FW_drops_log" isa 1/2

=====================================================
Application-Assurance event-log "FW_drops_log"
Current Time:          "06/17/2025 13:56:48" (UTC)
  group[:partition]:   1:1
  isa:                 1/2
  admin state:         no shutdown
  buffer-type:         circular
  max-entries:         100000
=====================================================
Event-source             Action        SubType     SubName

  Direction Src-ip  Dst-ip  Ip-protocol Src-port Dst-port Timestamp

Total Records:   0
=====================================================
```

3 Configure AA-ISA to protect SGW (GTP-U traffic)

The steps to configure the AA-ISA in an SeGW to protect against attacks toward the SGW are similar to the steps for SCTP traffic.

- 3.1 Create an AA IP list for SGWs
- 3.2 Allow only GTP-U traffic toward SGWs — no port scanning
- 3.3 DoS protection — limit the number of GTP-U flows from eNBs
- 3.4 DoS protection — limit the GTP-U bandwidth from eNBs
- 3.5 Further GTP filtering and validation

While GTP filtering is very different from SCTP filtering, the configuration to limit the flow counts, bandwidth, and session filter is similar.

3.1 Create an AA IP list for SGWs

In addition to the lists configured in 2.1 Create IP AA lists, the operator can optionally configure a list that contains the SGW IP addresses that are served by the SeGW, in case there is more than one.

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list "SGWs" {
                    description "Serving Gateways IPs"
                    prefix 172.16.111.1/32 {
                    }
                    prefix 172.16.111.2/32 {
                    }
                }
```

3.2 Allow only GTP-U traffic toward SGWs — no port scanning

Similar to 2.2 Allow only SCTP traffic toward MMEs — no port scanning, create an GTP filter to allow only GTP traffic to/from eNBs to SGWs:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter "SeGW_FW" {
                    default-action {
                        action deny
                        event-log "FW_drops_log"
                    }
---snip---
                    entry 2 {
                        description "allow GTP-u to SGWs"
                        match {
                            ip-protocol udp
                            dst-ip {
                                ip-prefix-list "SGWs"
                            }
                            dst-port {
                                eq 2152
                            }
                            src-ip {
                                ip-prefix-list "ALL_eNBs"
                            }
                        }
                        action {
                            permit
                        }
                    }
```

The following session filter needs to be added to the default sub-policy AQP, similar to 2.2 Allow only SCTP traffic toward MMEs — no port scanning:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 500 {
                            admin-state enable
                            description "apply SeGW session filter rules"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                session-filter "SeGW_FW"
                                bandwidth-policer {
                                    single-bucket "limit_eNBs_total_bw"
                                }
                                error-drop {
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_eNBs_total_Flows"
                                }
                                fragment-drop {
                                    drop-scope all
```

```
                                }
                                overload-drop {
                                }
                            }
                        }
```

For AA to recognize GTP traffic and perform sanity packet checking, configure a GTP filter at the group partition level:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                }
```

3.3 DoS protection — limit the number of GTP-U flows from eNBs

AA can be configured to limit the number of GTP flows from an eNB. A GTP-U flow is defined by GTP-U packet destination IP + tunnel ID (TEID).

AA allows the operator to configure two limits: one that applies to the each eNB and one that applies for all GTP-U traffic from all eNBs:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "GTPu-Flow-count-limit" {
                    granularity subscriber
                    limit-gtp-flows true
                    peak-flow-count 800
                }
```

The actual value of the flow count limit is a function of the number of UEs or devices served by an eNB or FAP. In the preceding case, it is assumed that there are 100 devices with a maximum of 8 GTP-U flows per device. For FAP, the number is typically around 32 devices per FAP.

Note: By 3GPP standards, the maximum number of GTP-U tunnels per device is 16.

Assuming that there are 1000 eNBs or FAPs that are served by the SeGW, then to limit the total number of GTP-U flows, the operator can apply the following system policer with granularity system:

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "limit_total_GTPU_Flow_count" {
                    granularity system
                    limit-gtp-flows true
                    peak-flow-count 800000
                }
```

Configure AQPs to execute the policers:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 120 {
                            admin-state enable
                            description "limit GTP-U traffic"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "GTPu-Flow-count-limit"
                                }
                            }
                        }
                        entry 130 {
                            admin-state enable
                            description "limit TOTAL GTPU towards SGWs"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "limit_total_GTPU_Flow_count"
                                }
                            }
                        }
                    }
```

For GTP-U flow count policing, it is important that **aqp-initial-lockup** is enabled:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                aqp-initial-lookup true
```

The preceding configured limits are applied only to upstream traffic, to protect the network. No limit is placed on the downstream traffic toward the eNBs.

Note:
For small cell deployments, the number of GTP-U tunnels per FAP is a function of:

- deployment mode:
  - residential = 32 UEs
  - enterprise = 8 UEs.
- number of guaranteed bit rate (GBR) tunnels (max 8) and non-GBR tunnels (max 8) per UE.

Therefore, the GTP-U tunnel limit per FAP should be set to 32 x 8 = 256 for residential deployments or 8 x 8 = 64 for enterprise deployments.

The operator can view the effect of the configured policers on GTP traffic by running the following **show** command:

```
[/]
A:admin@SeGW-2# show application-assurance group 1:1 gtp

===============================================================================
Application Assurance Group 1:1 GTP
===============================================================================
Admin status      : Up
Event log         : (Not Specified)
Event log action  : deny
Mode              : filtering
GTP-C inspection  : Disabled

-------------------------------------------------------------------------------
GTP Statistics                              sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Incoming packets                                 0                     0
Packets denied
  UDP packet length                              0                     0
  GTP message length                             0                     0
  GTP version                                    0                     0
-------------------------------------------------------------------------------
Packets permitted                                0                     0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
GTP Policing Statistics                     sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                  0                     0
Packets denied
  gtp-traffic flow-count policer                 0                     0
  Other                                          0                     0
-------------------------------------------------------------------------------
Packets permitted                                0                     0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
GTP Filter Statistics                       sub-to-net            net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                  0                     0
Packets denied                                   0                     0
Packets permitted
  gtp-filter                                     0                     0
  no gtp-filter                                  0                     0
-------------------------------------------------------------------------------
Total GTP packets permitted                      0                     0
===============================================================================
```

In the last section shown above, GTP filter statistics are related to GTP filters that are discussed and configured later in 3.5 Further GTP filtering and validation of this chapter.

3.4 DoS protection — limit the GTP-U bandwidth from eNBs

This step is similar to 3.3 DoS protection — limit the number of GTP-U flows from eNBs, but instead of configuring a flow count policer, the operator configures bandwidth policers:

```
configure {
    application-assurance {
```

```
        group 1 {
            policer {
                single-bucket-bandwidth-policer "GTPU_bw_limit" {
                    granularity subscriber
                    mbs 100
                    pir 5000
                }
                single-bucket-bandwidth-policer "limit_total_GTPU_bw" {
                    granularity system
                    mbs 2000
                    pir 2000000
                }
```

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 120 {
                            admin-state enable
                            description "limit GTP-U traffic"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "GTPU_bw_limit"
                                }
                                flow-count-limit-policer {
                                    policer-name "GTPu-Flow-count-limit"
                                }
                            }
                        }
                        entry 130 {
                            admin-state enable
                            description "limit TOTAL GTPU towards SGWs"
                            match {
                                traffic-direction subscriber-to-network
                            }
                            action {
                                bandwidth-policer {
                                    single-bucket "limit_total_GTPU_bw"
                                }
                                flow-count-limit-policer {
                                    policer-name "limit_total_GTPU_Flow_count"
                                }
                            }
                        }
```

The preceding configured limits are applied only to upstream traffic, to protect the network. No limit is placed on downstream traffic toward the eNB.

As a debugging tool, the operator can use the AA **flow-record-search** command to check the status of GTP flows through the ISA:

```
[/]
A:admin@SeGW-2# tools dump application-assurance group 1:1 flow-record-search isa 1/2 flow-
status active protocol "gtp"

=================================================
Application-Assurance flow record search
```

```
Search Start Time:     "06/17/2025 14:06:27" (UTC)
 Search Criteria:
  group[:partition]:   1:1
  isa:                 1/2
  protocol name:       "gtp"
  application name:    none specified
  app-group name:      none specified
  flow-status:         active
  start-flowId:        none specified
  classified:          none specified
  server-ip:           none specified
  server-port:         none specified
  client-ip:           none specified
  bytes-tx:            none specified
  flow-duration:       none specified
  max-count:           none specified
  flow-modified:       none specified
  search-type:         default
=====================================================
FlowId   Init  Src-ip  Dst-ip   Ip-prot   Src-prt Dst-prt Protocol

 Application    Pkts-tx    Bytes-tx    Pkts-disc  Bytes-disc


 Time-ofp(UTC)    Time-olp(UTC)


SEARCH COMPLETED.
Search End Time:     "06/17/2025 14:06:28" (UTC)
Total Records:       0
=====================================================
```

GTP flows that are to be denied by the previous AA configurations should not appear in the search results.

3.5 Further GTP filtering and validation

AA allows the operator to configure a GTP filter to enforce which GTP message types are allowed/denied, as well as the maximum allowed GTP message length:

```
*[ex:/configure application-assurance group 1 partition 1 gtp gtp-filter "test"]
A:admin@SeGW-2# ?

---snip---
 description           - Text description
 gtp-in-gtp            - GTP in GTP packet filtering
 ---snip---
 log                   + Enter the log context
 max-payload-length    - Maximum allowed GTP payload length
 message-type          + Enter the message-type context
---snip---
```

Note:
An AA GTP filter allows the operator to configure a maximum payload size for the GTP traffic. However, in this configuration example, no maximum payload size is configured.

The list of GTP message types are defined by 3GPP standard 3GPP TS 29.281 as per Table 10: GTP messages .

Table 10: GTP messages

Message type value (decimal)	Message	Message type value (decimal)	Message
1	echo-request	55	forward-relocation-complete
2	echo-response	56	relocation-cancel-request
3	version-not-supported	57	relocation-cancel-response
4	node-alive-request	58	forward-sms-context
5	node-alive-response	59	forward-relocation-complete-acknowledge
6	redirection-request	60	forward-sms-context-acknowledge
7	redirection-response	70	ran-information-relay
16	create-pdp-context-request	96	mbms-notification-request
17	create-pdp-context-response	97	mbms-notification-response
18	update-pdp-context-request	98	mbms-notification-reject-request
19	update-pdp-context-response	99	mbms-notification-reject-response
20	delete-pdp-context-request	100	create-mbms-context-request
21	delete-pdp-context-response	101	create-mbms-context-response
22	initiate-pdp-context-activation-request	102	update-mbms-context-request
23	initiate-pdp-context-activation-response	103	update-mbms-context-response
26	error-indication	104	delete-mbms-context-request
27	pdu-notification-request	105	delete-mbms-context-response
28	pdu-notification-response	112	mbms-registration-request
29	pdu-notification-reject-request	113	mbms-registration-response
30	pdu-notification-reject-response	114	mbms-de-registration-request
31	supported-extension-headers-notification	115	mbms-de-registration-response
32	send-routing-information-for-gprs-request	116	mbms-session-start-request

Message type value (decimal)	Message	Message type value (decimal)	Message
33	send-routing-information-for-gprs-response	117	mbms-session-start-response
34	Failure-report-request	118	mbms-session-stop-request
35	failure-report-request	119	mbms-session-stop-response
36	note-ms-gprs-present-request	120	mbms-session-update-request
37	note-ms-gprs-present-response	121	mbms-session-update-response
48	identification-request	128	ms-info-change-notification-request
49	identification-response	129	ms-info-change-notification-response
50	sgsn-context-response	240	data-record-transfer-request
51	sgsn-context-request	241	data-record-transfer-response
52	sgsn-context-acknowledge	254	end-marker
53	forward-relocation-request	255	g-pdu
54	forward-relocation-response		

Of the 67 GTP message types shown in Table 10: GTP messages , only 6 are allowed, by the standards, for GTP-U:

- **echo-request**
- **echo-response**
- **echo-indication**
- **g-pdu**
- **end-marker**
- **supported-extension-headers-notification**

If these message types are permitted by the configured GTP filter, AA performs extensive GTP-U header checking on these six types.

Note:
If no GTP filter is configured, no extensive GTP-U header checks are performed. For example, if the operator wants to allow all GTP-U packets and perform all GTP header sanity checks, then a GTP filter that allows all message types needs to be configured, with the default action of permit and with no values, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
```

```
                    gtp-filter "allow-all" {
                        message-type {
                            default-action permit
                        }
                    }
```

Because AA FW in an SeGW is protecting an S1-U interface running GTP-U, the GTP filter only needs to allow the six GTP messages that are permitted for GTP-U:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                gtp {
                    admin-state enable
                    gtp-filter "filter-gtp-msgs" {
                        description "allow only certain msg types"
                        message-type {
                            default-action deny
                            entry 1 {
                                action permit
                                value echo-request
                            }
                            entry 2 {
                                action permit
                                value echo-response
                            }
                            entry 3 {
                                action permit
                                value error-indication
                            }
                            entry 4 {
                                action permit
                                value supported-extension-headers-notification
                            }
                            entry 5 {
                                action permit
                                value end-marker
                            }
                            entry 6 {
                                action permit
                                value g-pdu
                            }
                        }
                    }
```

This GTP filter is then referenced from within an AQP entry action, as follows, in order for it to take effect:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                policy {
                    app-qos-policy {
                        entry 120 {
                            admin-state enable
                            description "limit GTP-U traffic"
                            match {
                                traffic-direction subscriber-to-network
                                dst-ip {
                                    eq {
                                        ip-prefix-list "SGWs"
```

```
                                    }
                                }
                            }
                            action {
                                gtp-filter "filter-gtp-msgs"
                                bandwidth-policer {
                                    single-bucket "GTPU_bw_limit"
                                }
                                flow-count-limit-policer {
                                    policer-name "GTPu-Flow-count-limit"
                                }
                            }
                        }
```

The operator can view the effect of the configured GTP filter on S1-U traffic using the following show routine:

```
[/]
A:admin@SeGW-2# show application-assurance group 1:1 gtp gtp-filter
                                                        gtp-filter-name "filter-gtp-msgs"

===============================================================================
Application Assurance Group 1:1 GTP Filter "filter-gtp-msgs"
===============================================================================
Description                 : allow only certain msg types
Maximum payload length      : (Not Specified)
Event log                   : (Not Specified)
Event log action            : deny
Default action              : deny
Default GTPv2 action        : permit
Default IMSI-APN action     : permit
GTP in GTP action           : permit
Validate GTP tunnels        : disabled
Validate sequence number    : disabled
Validate source IP address  : disabled
GTP tunnel endpoint limit   : (Not Specified)
Configured messages         : 6
Configured GTPv2 messages   : 0
Configured IMSI-APN filters : 0

Packets arrived             : 0
Packets denied
  Payload length            : 0
  Message type              : 0
  GTPv2 message type        : 0
  IMSI-APN filter           : 0
  Mandatory header          : 0
  Extension header          : 0
  Information element       : 0
  Invalid TEID              : 0
  Invalid sequence number   : 0
  Invalid source IP address : 0
  Missing mandatory IE      : 0
  GTP in GTP                : 0
  No tunnel resource        : 0
  Tunnel endpoint limit     : 0
Packets permitted           : 0
===============================================================================
```

The preceding output is in addition to the information provided by the overall GTP show command:

```
show application-assurance group 1:1 gtp
```

4 Configure AA-ISA to protect NMS (OAM traffic)

The following steps are described in this section:

- 4.1 Create an IP AA prefix list that contains the NMS server IP addresses
- 4.2 Allow eNBs to initiate FTP- and ICMP-only traffic toward NMS, block port scanning

4.1 Create an IP AA prefix list that contains the NMS server IP addresses

The following command configures IP prefix list "NMSs" with IP prefix 172.16.120.0/30:

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                ip-prefix-list "NMSs" {
                    description "Network Management-OAM subnet"
                    prefix 172.16.120.0/30 {
                    }
                }
```

Note:
In the case of small cell deployments, different NMS servers need to be configured

4.2 Allow eNBs to initiate FTP- and ICMP-only traffic toward NMS, block port scanning

The following entries are added to the session filter "SeGW_FW":

```
configure {
    application-assurance {
        group 1 {
            partition 1 {
                session-filter "SeGW_FW" {
                    default-action {
                        action deny
                        event-log "FW_drops_log"
                    }
---snip---
                    entry 3 {
                        description "allow FTP to NMS"
                        match {
                            ip-protocol tcp
                            dst-ip {
                                ip-prefix-list "NMSs"
                            }
                            dst-port {
                                eq 22
                            }
                            src-ip {
                                ip-prefix-list "ALL_eNBs"
                            }
                        }
                        action {
                            permit
```

```
                    }
                }
                entry 4 {
                    description "allow ICMP to NMS"
                    match {
                        ip-protocol icmp
                        dst-ip {
                            ip-prefix-list "NMSs"
                        }
                        src-ip {
                            ip-prefix-list "ALL_eNBs"
                        }
                    }
                    action {
                        permit
                    }
                }
```

The operator can view the effect of the session filter on traffic, in terms of how many times it is applied, using the following show routine:

```
[/]
A:admin@SeGW-2# show application-assurance group 1:1 session-filter

===============================================================================
AA Session Filter Table
===============================================================================
Name                            Default Action    Referenced           Entries
-------------------------------------------------------------------------------
SeGW_FW                         deny              aqp                        4
-------------------------------------------------------------------------------
No. of session filters: 1
===============================================================================
```

```
[/]
A:admin@SeGW-2# show application-assurance group 1:1 session-filter "SeGW_FW"

===============================================================================
AA Session Filter Instance "SeGW_FW"
===============================================================================
Description    : (Not Specified)
Default Action : deny
    Event Log  : FW_drops_log
AQP Entries    :
          500
-------------------------------------------------------------------------------
Filter Match Criteria
-------------------------------------------------------------------------------
Entry          : 1
Description    : allow SCTP to MMEs
IP Protocol    : sctp
Source IP List : ALL_eNBs
Dest IP List   : MMEs
Dest Port      : eq 6005
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 2
Description    : allow GTP-u to SGWs
IP Protocol    : udp
Source IP List : ALL_eNBs
```

```
Dest IP List   : SGWs
Dest Port      : eq 2152
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 3
Description    : allow FTP to NMS
IP Protocol    : tcp
Source IP List : ALL_eNBs
Dest IP List   : NMSs
Dest Port      : eq 22
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 4
Description    : allow ICMP to NMS
IP Protocol    : icmp
Source IP List : ALL_eNBs
Dest IP List   : NMSs
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
No. of entries   : 4
===============================================================================
```

Note:
The preceding configuration is generic and may need to be modified to suit the deployment requirements. For example, in the case of small cell SeGW deployment, traffic on other ports needs to be allowed to/from different NMS type servers, such as allowing TCP port 7003 and port 7013 to Home Device Manager (HDM) servers. This can be accomplished by configuring additional entries in the preceding session filter.

Note:
By allowing port 22 for FTP, the AA FW automatically opens and closes the associated data channel ports. For more information about AA FW capabilities, with regard to OAM FW protection, see Application Assurance — Stateful Firewall .

Conclusion

The SR OS AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW).

AA stateful packet filtering, combined with AA layer 7 classification and control, provides advanced, next-generation firewall functionality, protecting mobile network core infrastructure, such as MMEs, SGWs, and NMSs.

Application Assurance — Stateful Firewall

This chapter describes Application Assurance stateful firewall (FW) configurations for protecting residential and WiFi subscribers.

Topics in this chapter include:

Applicability

Initially, this chapter was written based on classic CLI for SR OS Release 11.0.R1. The TCP validation section was added for SR OS Release 14.0.R4. The MD-CLI in the current edition corresponds to SR OS Release 25.3.R2.

Overview

The AA stateful FW feature extends AA-ISA application level analysis to provide an in-line integrated stateful service that protects subscribers from malicious attacks. AA stateful packet filtering combined with AA layer 7 classification and control, empowers operators with advanced, next generation firewall functionality that is integrated within the Service Router. The AA stateful firewall (FW) and application firewall runs on AA-ISA. Using stateful inspection, the AA firewall not only inspects packets at layers 3-7, but also monitors and keeps track of the connection state. If the operator configures a **deny** action within a session filter, then the matching packets (matching both the AA Application QoS policy (AQP) and associated session filter match conditions) are dropped and no flow session state or context is created.

AA FW can be used in all deployments of AA-ISA; mobile (MG OS) and fixed (SR OS); however, the configuration examples here, while still very applicable (and almost 100% identical in mobile deployments) are focused on AA-ISA deployments in fixed networks.

The AA-ISA FW enabled solution provides:

- stateful (and stateless) packet filtering and inspection with application-level gateway (ALG) support
- denial of service (DoS) attack protection.

The objective of this chapter is to describe the required configuration within AA-ISA (divert to AA-ISA basic knowledge is assumed) to enable AA FW and protect AA subscribers from attacks (unsolicited attacks and DoS attacks), while still allowing pin-holing through the firewall, so that applications such as peer to peer gaming and various ALGs (such as FTP) are not affected.

Stateful filtering

By performing stateful inspection, AA-ISA takes into account which side initiated a session and acts accordingly. Stateful flow processing and inspection uses IP layer 3 and layer 4 header information to build a state of the flow within AA-ISA. Layer 7 inspection is used to provide ALG support. Stateful flow and session processing takes note of the originator of the session and therefore can allow traffic to be initiated from the subscriber, while denying (when configured) traffic originating from the network. Packets received from the network are inspected against the session filter and only those that are part of a subscriber-initiated session are allowed.

Figure 22: Block unsolicited traffic

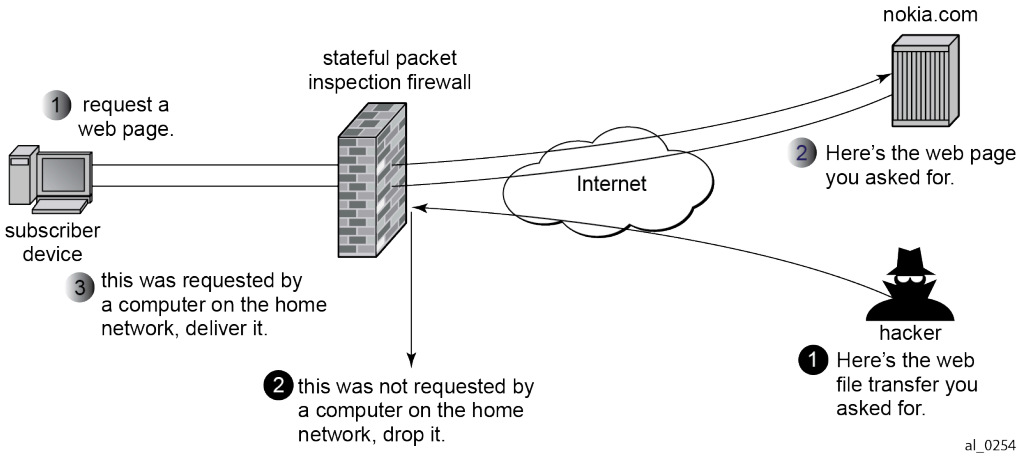

To support the example shown in Figure 22: Block unsolicited traffic, AA is configured with an action to block unsolicited traffic; traffic that is not requested by the subscriber. The direction field in match criteria of AQPs is used to enable this functionality.

Figure 23: SFW — allow gaming

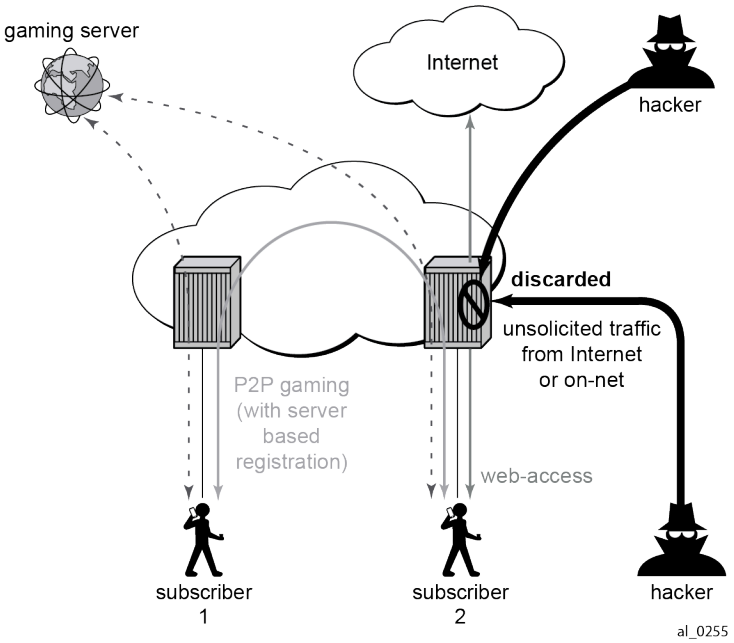

Figure 23: SFW — allow gaming shows a similar concept. It is used to allow UDP traffic for peer to peer applications, such as gaming. When the traffic from one peer is seen by AA-ISA, a pin-hole is opened in the reverse direction to allow for the corresponding UDP traffic from the peer.

Stateless packet filtering on the other hand does not take note of the session initiator. It discards or allows packets independently of any previous packets. In addition to the AA-ISA support for stateless (and stateful) filtering, stateless packet filtering can be performed in the system using line card ACLs (and MGISM PCC rules in mobile gateway deployments).

Application layer gateway filtering

Figure 24: ALG support example — FTP

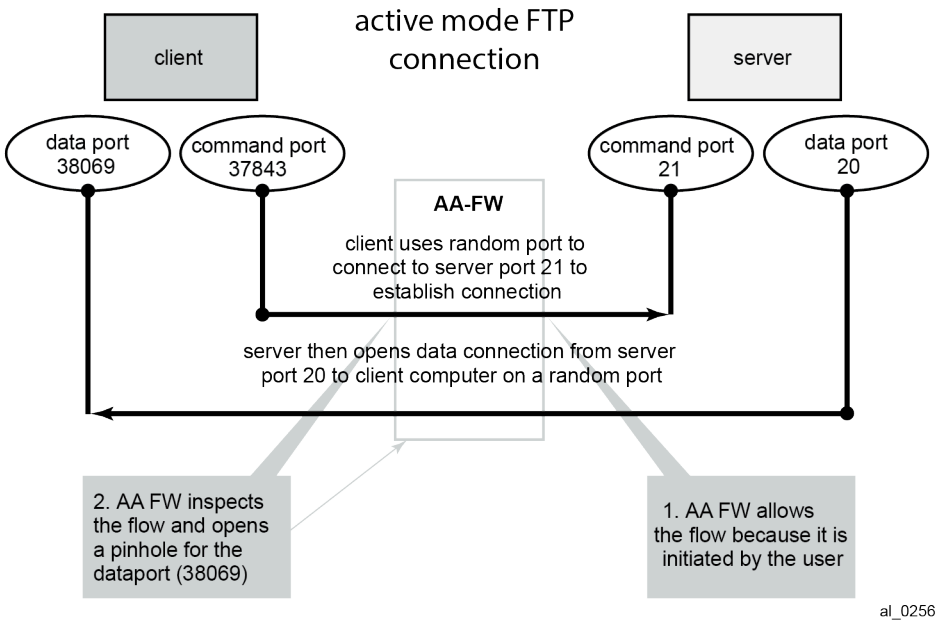

AA FW inspection of packets at Layer 7 offers application layer gateway functionality for a large list of applications (for example, FTP, SIP, RTSP, PPTP, IRC, and so on). These applications make use of control channels or flows that spawn other flows. AA FW inspects the payload of these control flows so it can open a pinhole in advance for unspawned data flows. Figure 24: ALG support example — FTP depicts an example of AA ALG support for FTP traffic.

Denial of Service (DoS) protection

DoS attacks work by consuming network and system resources, making them unavailable for legitimate network applications. Network flooding attacks, malformed packets and port scans are examples of such DoS attacks.

The aim of AA FW DoS protection is to protect subscribers and prevent any abuse of network resources.

Using AA FW stateful session filters, operators can protect their subscribers from any port scan scheme. This can be done by configuring the session filters to disallow any traffic that is initiated from the network.

Furthermore, AA ISA provides configurable flow policers. These policers prevent a wide range of flooding attacks (such as ICMP PING flooding, UDP flooding, SYN flood attack...and so on.). These policers provide protection at multiple levels; per system per application or application group and per subscriber per application or application group.

There are two types of AA ISA flow policers: flow setup rate policers and flow count policers. Flow setup rate policers limit the number of new flows, while flow count policers limit the total number of active flows.

To protect hosts and network resources, AA FW checks different fields in the packet header (checksum, TCP flag, and so on) and if any fails, it declares the packet to be invalid. This complements the SR OS

subscriber management enhanced security features, such as IP (or MAC) anti-spoofing protection (such as protecting against LAND attacks) and network protocol DoS protections. The cut-through-drop AQP action must be configured in order to drop these types of invalid packets.

Virtual FW or zone-based FW

AA FW can provide up to 128 virtual FWs, each with its own FW policies. This is achieved through the use of AA partitions.

In addition, AA subscribers within the same AA partition can have different application profiles with different Application Service Options (ASO) values. This provides a further control mechanism to enable or disable firewall rules.

For example, the operator may want to have some subscribers possess full firewall protection, while other subscribers not subscribed to this service have a partial firewall protection that focuses on protecting network resources, instead of network and subscribers resources.

Configuration

AA-ISA AQP actions provide session filtering functionality. AQPs have partition level scope, which allows different FW policies to be implemented by utilizing AA partitions concepts within the same AA-ISA group. Therefore, multiple virtual AA FW instances can be realized, without the need for multiple physical instances of FWs to implement different FW policies.

The AA FW stateful session filter consists of multiple entries (similar to ACLs) with a match and action per entry. A **deny** action results in packets being discarded without creating a session or flow context. Match conditions include IP protocol types, source and destination IP addresses, and source and destination ports. An overall default action is also configurable in case of no match to any session filter entry.

AQPs with session filter actions need to have — as a matching condition — traffic direction, ASOs, and/or subscriber name. These AQP match rules cannot have any references to applications or application groups.

An AQP action to drop malformed or errored packets is also available.

Statistics are incremented when packets are dropped by a session filter. These are accounted against:

- protocol: denied by default policy
- application: unknown
- application group: unknown.

The configuration topology is shown in Figure 25: Configuration topology.

Figure 25: Configuration topology

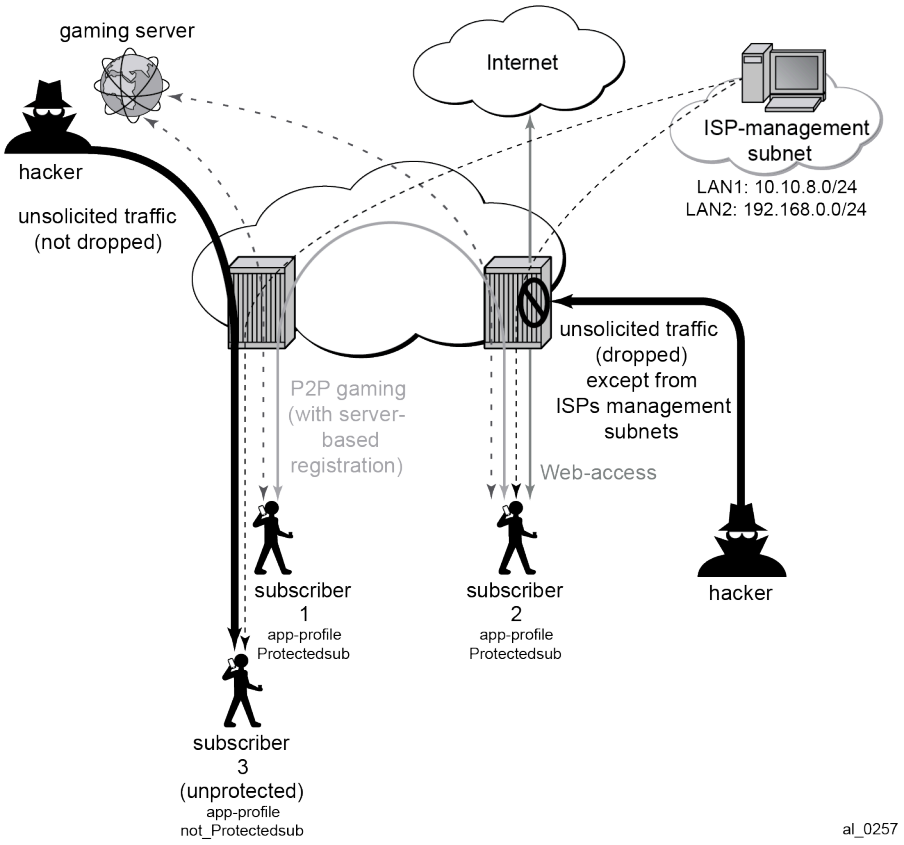

Application profiles

Nothing new is introduced in application profiles to support FW. This section describes how to configure the application profile to allow differentiated FW services for different subscribers. In a nutshell, the AA common building construct or attribute for differentiated policy is ASO.

To configure an ASO for FW protection:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-service-options {
                        characteristic "DOS-Protection" {
                            default-value "None"
                            value "None" { }
                            value "ON" { }
                        }
                        characteristic "FW-Protection" {
                            default-value "None"
                            value "None" { }
                            value "ON" { }
                        }
```

```
                        characteristic "ISP-Protection" {
                            default-value "None"
                            value "None" { }
                            value "ON" { }
                        }
                    }
                }
```

In the preceding example:

- ASO "FW-protection" allows the operator to select if the subscriber is FW protected or not.
- ASO "DOS-protection" refers to if the subscriber is protected from DoS attacks.
- ASO "ISP-protection" is different from the preceding two because it protects the ISP resources by (in the example that follows) not allowing unsolicited traffic. This should be ON for all subscribers (it is then arguable if someone needs it to be defined in the ASO list, instead of merely configuring an AQP to protect ISP resources all the time).

These ASOs are referenced in appProfiles (and later in AQPs) as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-profile "Protected" {
                        divert true
                        characteristic "DOS-Protection" {
                            value "ON"
                        }
                        characteristic "FW-Protection" {
                            value "ON"
                        }
                        characteristic "ISP-Protection" {
                            value "ON"
                        }
                    }
```

The preceding application profile "Protected" is assigned to subscribers who opted/subscribed to the firewall protection service; for example, subscriber 1 and subscriber 2 in the example shown in Figure 25: Configuration topology.

Subscribers who are not protected (for example subscriber 3 in Figure 25: Configuration topology) are assigned a different profile:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-profile "unProtected" {
                        divert true
                        characteristic "DOS-Protection" {
                            value "None"
                        }
                        characteristic "FW-Protection" {
                            value "None"
                        }
                        characteristic "ISP-Protection" {
                            value "ON"
                        }
```

```
                }
            }
```

An alternative method to using application profiles/ASOs to provide differentiated services is to configure multiple partitions with different AQPs/session filters. One partition, for example, is for subscribers who are provided with firewall protection, while another is used for subscribers who are not protected. This configuration is simpler and provides statistics per partition. This example however covers the more complex case using ASOs/appProfiles.

Flow count policers

The following configuration limits the number of flows a subscriber can have at any time to 500. This is done to protect against DoS attacks. The value 500 is arbitrary and requires tuning for each deployment.

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "Dos_police_Flow_count" {
                    granularity subscriber
                    peak-flow-count 500
                }
            }
```

The following configuration limits the total number of flows that matches the configured AQP matching condition. It is used for ICMP applications to prevent mass port scanning.

```
configure {
    application-assurance {
        group 1 {
            policer {
                flow-count-limit-policer "Dos_Police_ICMPFlows" {
                    granularity system
                    peak-flow-count 5000
                }
            }
```

TCP protocol validation

The following configuration allows the operator to call the "TCP_protect" policy from within an AQP action entry.

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                tcp-validate "TCP_protect" {
                }
            }
```

The operator can also configure the policy to be strict, in which case the AA checks for valid sequence and acknowledgements numbers. In this example, the strict option is not used.

Applications

The following application configuration is standard with AppDB. It is shown here for reference.

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    application "ICMP" {
                    }
                    app-filter {
                        entry 1540 {
                            admin-state enable
                            application "ICMP"
                            ip-protocol {
                                eq icmp
                            }
                            protocol {
                                eq "non_tcp_udp"
                            }
                        }
                        entry 35500 {
                            admin-state enable
                            application "ICMP"
                            ip-protocol {
                                eq ipv6-icmp
                            }
                            protocol {
                                eq "non_tcp_udp"
                            }
                        }
                    }
                }
```

Session filters

The following displays session filter configuration commands to be used.

```
configure {
  application-assurance {
    group 1 {
      partition 2 {
        session-filter <name> {
          description  <description>
          default-action {
              action permit|deny    # default=deny
          }
          entry <...> {
              description <entry-description>
              match {
                ip-protocol <ip-protocol-number> | <ip-protocol-name>
                src-ip  <ip4_or_v6-address/mask> | ip-prefix-list <name>
                dst-ip  <ip4_or_v6-address/mask> | dns-ip-cache <name>| ip-prefix-list <name>
                src-port {eq|gt|lt} <port-num> | range <start> <end> | port-list <name>
                dst-port {eq|gt|lt} <port-num> | range <start> <end> | port-list <name>
              }
              action {
                permit | deny | http-redirect | l3-l4-redirect | tcp-optimizer
```

```
            }
          }
    ---snip---
```

Parameters

- **entry**

  A session filter can have multiple match-action rules, each of these match-action rules represents an entry within the session filter. The entries are executed in order. If a match is found, within one entry, the subsequent entries within the session filter are skipped (not evaluated).

- **default-action > action [permit | deny]**

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> default-action]
A:admin@PE-1# action ?

 action <keyword>
 <keyword> - (deny|permit)
 Default   - deny

    Default action for packets not matching filter entries
```

  The default action is performed if no match is found for any of the configured entries within the session filter. Default is deny.

  – A **deny** action drops the packet and does not allow a flow record to be allocated for that flow. A **drop** action within AA AQP drops the packet but it still creates flow record.

  – A **permit** action allows the packet to flow through the system. A flow record is also allocated. The packet may get dropped by other configured AQP actions (because of header check failures).

- **description** *description-string*

  This configures a text string which can be used to describe the use of the session-filter.

- **match**

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10 match]
A:admin@PE-1# ?

 dst-ip                + Enter the dst-ip context
 dst-port              + Enter the dst-port context
 ip-protocol           - IP protocol as a match criterion
 src-ip                + Enter the src-ip context
 src-port              + Enter the src-port context
```

  Keywords to perform the action specified under the **action** keyword only if the conditions in the match section are met.

  – **ip-protocol: <protocol-number> | <protocol-name>**

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10
 match]
A:admin@PE-1# ip-protocol ?

 ip-protocol (<number> | <keyword>)
 <number>  - <0..255>
 <keyword> - (tcp-udp|icmp|igmp|ip|tcp|egp|igp|udp|rdp|ipv6|ipv6-route|ipv6-frag|
             idrp|rsvp|gre|ipv6-icmp|ipv6-no-nxt|ipv6-opts|iso-ip|eigrp|ospf-igp|
             ether-ip|encap|pnni|pim|vrrp|l2tp|stp|ptp|isis|crtp|crudp|sctp)
```

```
    IP protocol as a match criterion
```

– **src-ip>/dst-ip** defines the source an destination IP address within the packet header.

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10
 match]
A:admin@PE-1# src-ip ?

 src-ip

 Choice: match-addr-choice
 ip-prefix           :- Source IP address prefix as match criterion
 ip-prefix-list      :- Source IP address prefix list as match criterion
```

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10
 match]
A:admin@PE-1# dst-ip ?

 dst-ip

 Choice: match-addr-choice
 dns-ip-cache        :- Destination IPs in specified DNS IP Cache
 ip-prefix           :- Destination IP address prefix as match criterion
 ip-prefix-list      :- IP address prefix list as match criterion
```

– **src-port/dst-port**

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10
 match]
A:admin@PE-1# src-port ?

 src-port

 Choice: match-op-choice
 eq                  :- Match criterion used for destination or source port
 gt                  :- Greater than match criterion for the port number
 lt                  :- Less than match criterion for the port
 port-list           :- Destination or source port list
 range               :+ Enable the range context
```

```
[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10
 match]
A:admin@PE-1# dst-port ?

 dst-port

 Choice: match-op-choice
 eq                  :- Match criterion used for destination or source port
 gt                  :- Greater than match criterion for the port number
 lt                  :- Less than match criterion for the port
 port-list           :- Destination or source port list
 range               :+ Enable the range context
```

src-port/dst-port {eq | gt | lt} *number*

- **eq** — equal, exact match
- **gt** — match port numbers that are greater than the one specified.
- **lt** — match port numbers that are lower than the one specified.

- *number* — 0..65535 (port number applicable to TCP, UDP and SCTP protocols only). Default: 0.

- **action**: only executed if a match is found.

```
*[ex:/configure application-assurance group 1 partition 2 session-filter <..> entry 10 action]
A:admin@PE-1# ?

 event-log             - Event log name used to log the action

 Choice: action-choice
 deny                 :- Deny sessions matching the criteria
 http-redirect        :- HTTP redirect for matching sessions
 l3-l4-redirect       :+ Enter the l3-l4-redirect context
 permit               :- Permit sessions that match the criteria
 tcp-optimizer        :- TCP optimizer to handle sessions matching the criteria
```

- **deny** action drops the packet and does not create a flow record.
- **permit** action allows the packet to go through (unless another different action is found that causes it to be dropped).
- **http-redirect** action refers to a HTTP redirect policy.
- **l3-l4-redirect** action redirects sessions matching the criteria to a different destination using a layer 3 and layer 4 redirect.
- **tcp-optimizer** action refers to a TCP optimization policy.

The session filter "denyUnsolictedwMgntCntrl" is configured as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                session-filter "denyUnsolictedwMgntCntrl" {
                    default-action {
                        action deny
                    }
                    entry 10 {
                        description "allow ICMP access from ISP LAN1"
                        match {
                            ip-protocol icmp
                            src-ip {
                                ip-prefix 10.10.8.0/24
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 20 {
                        description "allow ICMP access from ISP LAN2"
                        match {
                            ip-protocol icmp
                            src-ip {
                                ip-prefix 192.168.0.0/24
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 30 {
```

```
                        description "allow all TCP (e.g. FTP/telnet)access from ISP LAN2"
                        match {
                            ip-protocol tcp
                            src-ip {
                                ip-prefix 192.168.0.0/24
                            }
                        }
                        action {
                            permit
                        }
                    }
                    entry 40 {
                        description "allow TCP on port 80 /HTTP access from ISP LAN1"
                        match {
                            ip-protocol tcp
                            dst-port {
                                eq 80
                            }
                            src-ip {
                                ip-prefix 10.10.8.0/24
                            }
                        }
                        action {
                            permit
                        }
                    }
                }
```

The session filter "protectISPLan2" is used to protect systems located in LAN2. It drops all unsolicited traffic except for FTP coming from LAN1.

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                session-filter "protectISPLan2" {
                    description "S-FW to deny all unsolicited requests to LAN2"
                    default-action {
                        action deny
                    }
                    entry 10 {
                        description "allow ftp access from ISP LAN1"
                        match {
                            ip-protocol tcp
                            dst-port {
                                eq 21
                            }
                            src-ip {
                                ip-prefix 10.10.8.0/24
                            }
                        }
                        action {
                            permit
                        }
                    }
```

AQPs

```
configure {
    application-assurance {
```

```
        group 1 {
            partition 2 {
                policy {
                    app-qos-policy {
                        entry 100 {
                            admin-state enable
                            description "protecting ISP1 from DoS attacks from subs"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "ISP-Protection" {
                                    eq "ON"
                                }
                                dst-ip {
                                    eq {
                                        ip-prefix 10.10.8.0/24
                                    }
                                }
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "Dos_police_Flow_count"
                                }
                            }
                        }
                        entry 105 {
                            admin-state enable
                            description "protecting ISP2 from DoS attacks from subs"
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "ISP-Protection" {
                                    eq "ON"
                                }
                                dst-ip {
                                    eq {
                                        ip-prefix 192.168.0.0/24
                                    }
                                }
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "Dos_police_Flow_count"
                                }
                            }
                        }
```

These AQPs protect the ISP network by limiting the number of concurrent flows. Dropping malformed packets is done by entry 130 (see further).

To guard against ICMP flooding attacks, a flow count policer (defined earlier) is used as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-qos-policy {
                        entry 107 {
                            admin-state enable
                            match {
                                traffic-direction subscriber-to-network
                                application {
                                    eq "ICMP"
                                }
```

```
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "Dos_Police_ICMPFlows"
                                }
                            }
                        }
```

To guard against attacks exploiting TCP handshake mechanisms, TCP validate policy is used as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                tcp-validate "TCP_protect" {
                }
                policy {
                    app-qos-policy {
                        entry 108 {
                            admin-state enable
                            match {
                                characteristic "ISP-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                tcp-validate "TCP_protect"
                            }
                        }
                        entry 109 {
                            admin-state enable
                            match {
                                characteristic "FW-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                tcp-validate "TCP_protect"
                            }
                        }
```

TCP validation works on both directions and needs to be called in from within a sub-default AQP entry. Therefore, this AQP action cannot be restricted to one ISP versus another because no destination IP address can be specified. The configuration shown runs TCP validation policy when ISP-Protection or FW-protection ASOs are enabled.

The preceding configuration ensures, for example, that no TCP session starts without the correct handshake message exchanges.

To protect ISP LAN2 from all unsolicited incoming traffic, entry 120 is configured:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-qos-policy {
                        entry 120 {
                            admin-state enable
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "ISP-Protection" {
```

```
                                    eq "ON"
                                }
                            }
                            action {
                                session-filter "protectISPLan2"
                            }
                        }
```

The session filter "ProtectISPLan2" drops all unsolicited traffic to LAN2 (highly secure) except for access to FTP services coming from ISP LAN1.

To enable stateful protection for opted-in subscribers:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-qos-policy {
                        entry 110 {
                            admin-state enable
                            description "FW for managed opted-in subs"
                            match {
                                traffic-direction network-to-subscriber
                                characteristic "FW-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                session-filter "denyUnsolictedwMgntCntrl"
                            }
                        }
```

The preceding AQP protects opt-in subscribers from unsolicited traffic but still allows unsolicited traffic from ISP subnets to manage the subscriber network.

Dropping malformed or illegal packets and protecting against DOS attacks is done via the following entries 130 and 131.

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                policy {
                    app-qos-policy {
                        entry 130 {
                            admin-state enable
                            match {
                                traffic-direction subscriber-to-network
                                characteristic "DOS-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                flow-count-limit-policer {
                                    policer-name "Dos_police_Flow_count"
                                }
                            }
                        }
                        entry 131 {
                            admin-state enable
                            match {
```

```
                                characteristic "DOS-Protection" {
                                    eq "ON"
                                }
                            }
                            action {
                                error-drop {
                                }
                                fragment-drop {
                                    drop-scope all
                                }
                                overload-drop {
                                }
                            }
                        }
```

Threshold crossing alerts

Operators can configure AA to generate TCAs for various firewall related parameters, such as error-drop, session-filter hits, TCP-validate, fragment-drop-all and so on, as well as flow count policers. An example of a TCA used for TCP validation policy is as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                threshold-crossing-alert {
                    tcp-validate "TCP_protect" direction from-sub {
                        high-watermark 50
                        low-watermark 40
                    }
                }
            }
```

Unlike the other TCAs, to configure TCAs for flow count policers, operators need first to configure AA admit-deny to allocate ISA resources, such as:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                statistics {
                    aa-admit-deny {
                        policer-stats-resources true
                    }
```

A TCA can be configured for any flow based policer in the system, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                threshold-crossing-alert {
                    policer "Dos_police_Flow_count" direction from-sub {
                        high-watermark 300
                        low-watermark 199
                    }
```

The system allows the various AA-admit-deny statistics to be exported via XML according to the configured accounting policy on the system. Analytics systems can then use these statistics to generate the right reports and alerts.

As a prerequisite, an accounting policy is configured for aa-admit-deny statistics:

```
configure {
    log {
        accounting-policy 5 {
            admin-state enable
            record aa-admit-deny
        }
```

Then, the operator can configure AA to export the statistics related to various firewall functions configured in the system, as follows:

```
configure {
    application-assurance {
        group 1 {
            partition 2 {
                statistics {
                    aa-admit-deny {
                        accounting-policy 5
                        collect-stats true
                        policer-stats-resources true
                        session-filter-stats true
                        tcp-validate-stats true
                    }
```

GPRS tunneling protocol (GTP) and stream control transmission protocol (SCTP) admit deny stats are related to firewall deployment within a SeGW, which is not covered within the scope of this chapter.

Show commands

Show routines — AQP:

```
[/]
A:admin@PE-1# show application-assurance group 1:2 policy app-qos-policy 110

===============================================================================
Application QOS Policy Entry 110 (Default Subscriber Policy)
===============================================================================
Description : FW for managed opted-in subs
Admin State : in-service
Hits        : 0 flows
Conflicts   : 0 flows

Match :
    Traffic Direction        : network-to-subscriber
    ASO Characteristics      :
        FW-Protection                 : eq ON

Action :
    Session Filter               : denyUnsolictedwMgntCntrl
===============================================================================
```

Show routines — session filter:

```
[/]
A:admin@PE-1# show application-assurance group 1:2 session-filter

===============================================================================
AA Session Filter Table
===============================================================================
Name                            Default Action    Referenced           Entries
-------------------------------------------------------------------------------
protectISPLan2                  deny              aqp                        1
denyUnsolictedwMgntCntrl        deny              aqp                        4
-------------------------------------------------------------------------------
No. of session filters: 2
===============================================================================
```

```
[/]
A:admin@PE-1# show application-assurance group 1:2 session-filter "denyUnsolictedwMgntCntrl"

===============================================================================
AA Session Filter Instance "denyUnsolictedwMgntCntrl"
===============================================================================
Description    : (Not Specified)
Default Action : deny
    Event Log  : (Not Specified)
AQP Entries    :
          110
-------------------------------------------------------------------------------
Filter Match Criteria
-------------------------------------------------------------------------------
Entry          : 10
Description    : allow ICMP access from ISP LAN1
IP Protocol    : icmp
Source IP      : 10.10.8.0/24
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 20
Description    : allow ICMP access from ISP LAN2
IP Protocol    : icmp
Source IP      : 192.168.0.0/24
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 30
Description    : allow all TCP (e.g. FTP/telnet)access from ISP LAN2
IP Protocol    : tcp
Source IP      : 192.168.0.0/24
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 40
Description    : allow TCP on port 80 /HTTP access from ISP LAN1
IP Protocol    : tcp
Source IP      : 10.10.8.0/24
Dest Port      : eq 80
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
```

```
No. of entries   : 4
===============================================================================
```

Show routines — TCP validation:

```
[/]
A:admin@PE-1# show application-assurance group 1:2 tcp-validate "TCP_protect"

===============================================================================
Application Assurance Group 1:2 tcp-validate "TCP_protect"
===============================================================================
Description      : (Not Specified)
Event log        : (Not Specified)
Strict Validation: No
AQP referenced   : Yes

-------------------------------------------------------------------------------
Decision Statistics                        sub-to-net              net-to-sub
-------------------------------------------------------------------------------

Total
-------------------------------------------------------------------------------
Allowed
  Octets                                              0                       0
  Packets                                             0                       0
Dropped
  Octets                                              0                       0
  Packets                                             0                       0

Dropped Reason
-------------------------------------------------------------------------------
Bad Flags
  Octets                                              0                       0
  Packets                                             0                       0
Bad Options
  Octets                                              0                       0
  Packets                                             0                       0
Bad Sequence Number
  Octets                                              0                       0
  Packets                                             0                       0
Bad Acknowledgment Number
  Octets                                              0                       0
  Packets                                             0                       0
No Establishment
  Octets                                              0                       0
  Packets                                             0                       0
SYN After Conn Establishment
  Octets                                              0                       0
  Packets                                             0                       0
Asymmetric Traffic
  Octets                                              0                       0
  Packets                                             0                       0
Traffic After Reset (RST)
  Octets                                              0                       0
  Packets                                             0                       0
Fragmented
  Octets                                              0                       0
  Packets                                             0                       0
===============================================================================
```

```
[/]
A:admin@PE-1# show application-assurance threshold-crossing-alert detail
```

```
===============================================================================
Application Assurance Threshold Crossing Alerts
===============================================================================

-------------------------------------------------------------------------------
policer "Dos_police_Flow_count" from-sub
-------------------------------------------------------------------------------
Group:Part     : 1:2                   Trigger on     : denied-traffic
High watermark : 300                   Low watermark  : 199
Last raised    : N/A                   Last cleared   : N/A
State          : cleared

-------------------------------------------------------------------------------
tcp-validate "TCP_protect" from-sub
-------------------------------------------------------------------------------
Group:Part     : 1:2                   Trigger on     : denied-traffic
High watermark : 50                    Low watermark  : 40
Last raised    : N/A                   Last cleared   : N/A
State          : cleared

No. of TCAs : 2
===============================================================================
```

The following output is slightly modified to make the wide table fit on the page.

```
[/]
A:admin@PE-1# tools dump application-assurance group 1:2 admit-deny-stats

===============================================================================
Application-Assurance Group 1:2 Admit-Deny Statistics
===============================================================================

-------------------------------------------------------------------------------
                               Admitted     Denied     Admitted      Denied
                             Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
Packet Validation Statistics  (Packets)   (Packets)   (Packets)   (Packets)
-------------------------------------------------------------------------------
Error                                0           0           0           0
Fragments: Out-Of-Order              0           0           0           0
Fragments: All                       0           0           0           0
Overload                           N/A           0         N/A           0

-------------------------------------------------------------------------------
                               Admitted     Denied     Admitted      Denied
                             Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
Session Filter Statistics     (Packets)   (Packets)   (Packets)   (Packets)
-------------------------------------------------------------------------------
Session Filter: protectISPLan2
 Entry: 10                           0           0           0           0
 Default Action                      0           0           0           0

Session Filter: denyUnsolictedwMgntCntrl
 Entry: 10                           0           0           0           0
 Entry: 20                           0           0           0           0
 Entry: 30                           0           0           0           0
 Entry: 40                           0           0           0           0
 Default Action                      0           0           0           0

-------------------------------------------------------------------------------
                               Admitted     Denied     Admitted      Denied
                             Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
Flow Policer Statistics       (Packets)   (Packets)   (Packets)   (Packets)
-------------------------------------------------------------------------------
```

```
System Flow Count Policers
   Dos_Police_ICMPFlows                0           0           0           0

Subscriber Flow Count Policers
   Dos_police_Flow_count               0           0           0           0

-------------------------------------------------------------------------------
                               Admitted     Denied     Admitted      Denied
                             Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
TCP Validation Statistics     (Packets)   (Packets)   (Packets)   (Packets)
-------------------------------------------------------------------------------
TCP_protect                          0           0           0           0
-------------------------------------------------------------------------------
```

Conclusion

The AA stateful packet filtering feature combined with AA layer 7 classification and control empowers operators with an advanced, next generation firewall functionality that is integrated within SR OS. This chapter focuses on traditional stateful and stateless session firewall functionality.

Deterministic Large Scale NAT44

This chapter provides information about deterministic large scale NAT44 configurations.

Topics in this chapter include:

Applicability

This chapter was initially based on SR OS Release 11.0.R3, but the MD-CLI in the current edition corresponds to SR OS Release 25.7.R2.

Overview

Deterministic Network Address Translation (NAT) is a mode of operation where mappings between the NAT subscriber and the outside IP address and port range are allocated at the time of configuration. In deterministic NAT for Large Scale NAT IPv4-to-IPv4 (LSN44) subscribers, each LSN44 subscriber is permanently mapped to an outside IP address and a dedicated deterministic port-block based on a specific algorithm. Logging is not needed in this case because the reverse mapping can be obtained using the reverse of the preceding algorithm.

Figure 26: Deterministic NAT mapping shows the mapping between inside IP addresses and outside IP addresses, where the deterministic port blocks use the port range after the static port forwarding range. A deterministic LSN44 subscriber can have only one deterministic port block that can optionally be extended by one or multiple dynamic port blocks in case all ports in deterministic port block are exhausted.

Figure 26: Deterministic NAT mapping

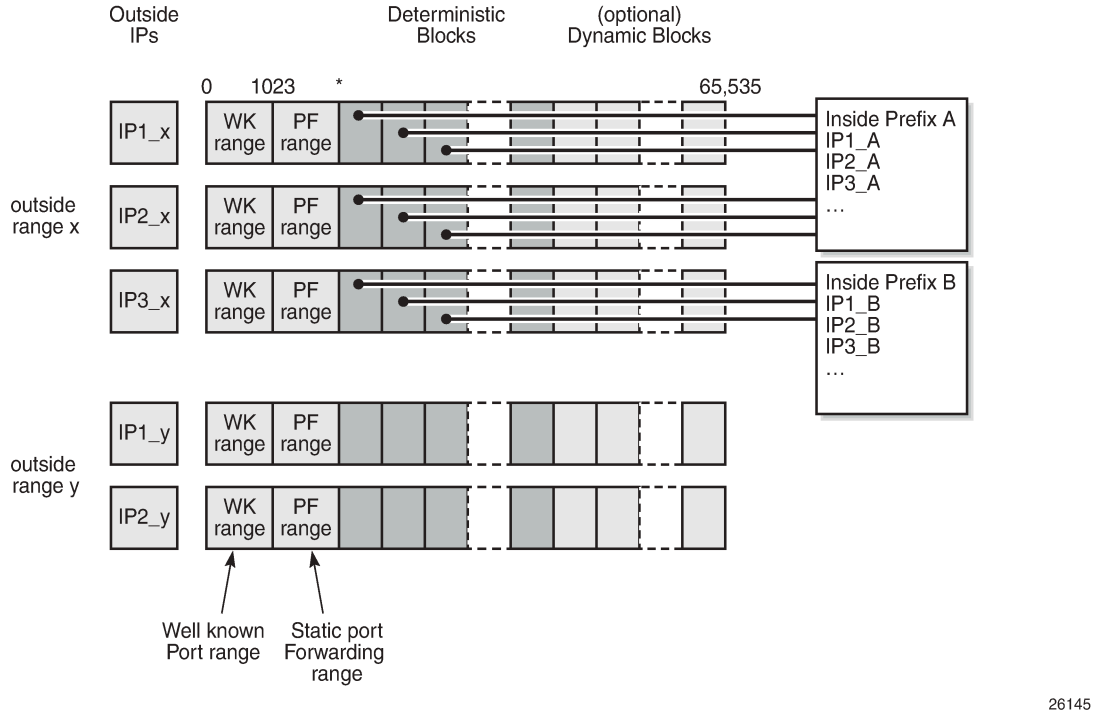

In case an LSN44 subscriber has been assigned both deterministic and dynamic port blocks, logging for the dynamic port block allocation and de-allocation is required. A scalable logging solution for dynamic port blocks is achievable using RADIUS or IPFIX. Logging for dynamic port blocks is out of the scope of this chapter.

Algorithm

The deterministic NAT algorithm makes a predictable mapping between the (inside IP, routing instance) and the (outside IP, routing instance, deterministic port block), as shown in Figure 27: Deterministic NAT algorithm. The algorithm is revertive, meaning that a specific (outside IP, routing instance, deterministic port block) derives a specific (inside IP, routing instance).

The algorithm is loosely based on *draft-donley-behave-deterministic-cgn-00*, which allows for the dynamic expansion of the port blocks when the ports in the original deterministic port block are exhausted.

Figure 27: Deterministic NAT algorithm

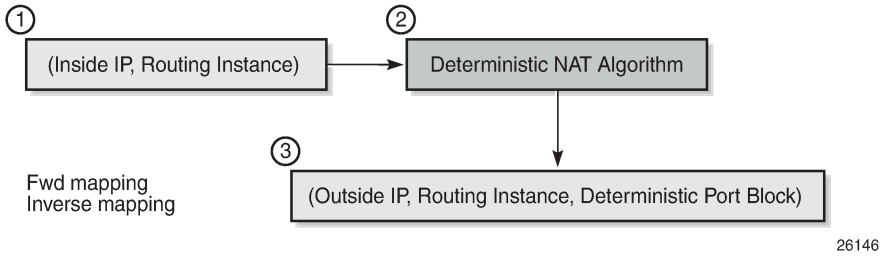

Deterministic mapping

Any inside prefix in any routing instance can be mapped to any pool in any routing instance.

In deterministic NAT, prefixes from multiple routing instances can be mapped to the same outside pool, also prefixes from a single inside routing instance can be selectively mapped to different outside pools, as shown in Figure 28: Deterministic mapping: inside -> outside routing instances .

Figure 28: Deterministic mapping: inside -> outside routing instances

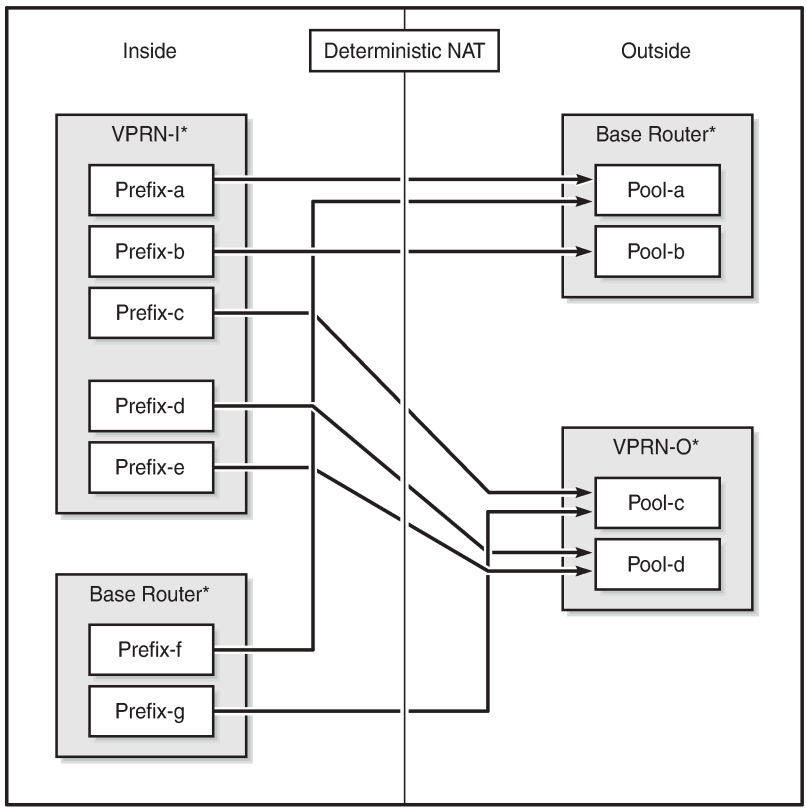

Mapping rules

A deterministic LSN44 subscriber is mapped to only one deterministic block which can further be extended to multiple dynamic blocks if ports within the deterministic block are exhausted.

The subscriber limit is the number of subscribers that can be deterministically mapped to one outside IP address (that is, compression ratio) and must be a power of 2.

The total number of deterministic ports (DetP) per outside IP address is determined by this subscriber limit and the number of deterministic ports per subscriber.

The remaining ports (DynP) beyond the deterministic port range up to 65535 are dedicated for dynamic use when a deterministic block is exhausted.

Every host using an inside prefix is guaranteed one dedicate block in the deterministic port ranges.

If the inside prefix length is m < 32-n, where 2^n=subscriber limit, then the prefix must be broken into pieces so that all hosts (subscriber limit) in each piece maps exactly to one outside IP address.

For example, if there is an inside prefix 192.168.0.0/23, with m=23 and a maximum number of 256 hosts; and the subscriber limit set to 256, then n=8. This results in 23 < 24 (32-8) and so this inside prefix must be broken into 2 pieces, in other words, this inside prefix fits into 2 outside IP addresses, each of 256 port blocks.

In case that the prefix length is m ≥ 32-n, where 2^n=subscriber limit, then all hosts from the configured prefix are mapped to the same outside IP.

For example, if there is an inside prefix 192.168.1.0/25, with m=25 and a maximum number of 128 hosts, and the subscriber limit set to 256, then n=8. This results in 25 > 24 (32-8), so definitely 128 hosts can fit in one outside IP because there are 256 available port blocks, in other words, this inside prefix fits into one outside IP where 128 blocks have been used out of the 256 port-blocks available, and the remaining 128 ports are wasted.

Overbooking of the outside address pool is not supported in deterministic NAT.

Figure 29: Deterministic mapping: outside IP port-blocks/ranges

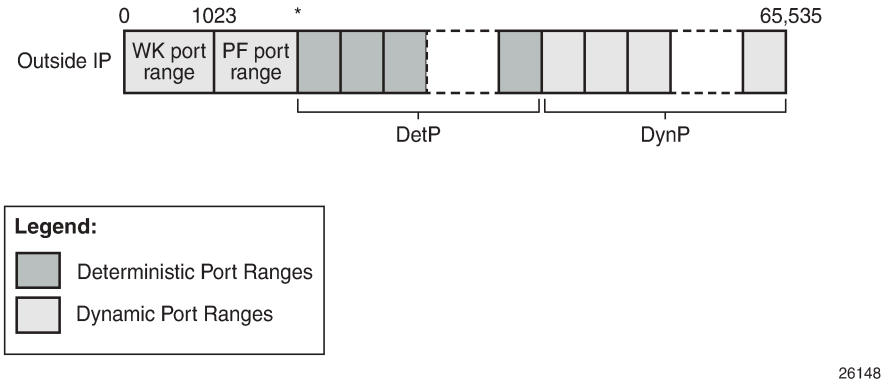

Configuration

Figure 30: Example topology shows a topology with one node connected to two test center ports. Traffic can be sent between tester port 1 and tester port 2.

Figure 30: Example topology

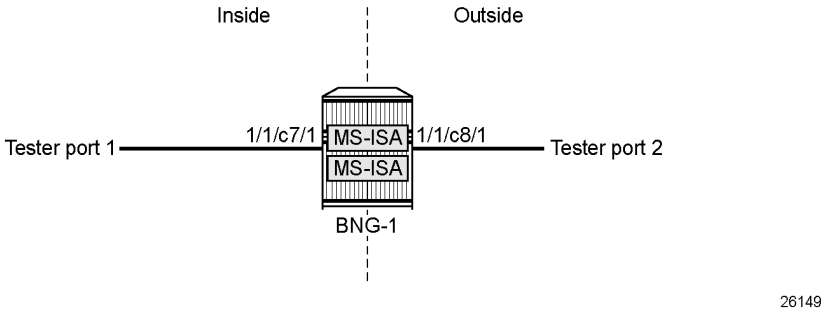

Note:
Private address ranges are used in outside pools in this chapter but normally public address ranges are used.

Configuration prerequisites

The card and MDA configuration is as follows:

```
configure {
    card 1 {
        admin-state enable
        card-type i80-200g-sfpdd+12-400g-qsfpdd-1
        mda 1 {
            admin-state enable
            mda-type m80-200g-sfpdd+12-400g-qsfpdd-1
        }
```

The ESA configuration is as follows:

```
configure {
    esa 1 {
        admin-state enable
        host-port 1/1/c80/1 {
        }
        host-port 1/1/c92/1 {
        }
        vm 1 {
            admin-state enable
            host-port 1/1/c80/1
            vm-type bb
            cores 4
            memory 16
        }
        vm 3 {
            admin-state enable
```

```
            host-port 1/1/c92/1
            vm-type bb
            cores 4
            memory 16
        }
```

NAT group 1 is configured as follows:

```
configure {
    isa {
        nat-group 1 {
            admin-state enable
            description "NAT group 1"
            redundancy {
                active-mda-limit 1
            }
            esa 1 vm 1 { }
            esa 1 vm 3 { }
        }
```

Configuration commands

A NAT outside pool is configured using the following command:

```
configure  {router | service vprn <service-id>}
  nat {
    outside {
      pool <nat-pool-name>
        type large-scale
        nat-group <nat-group-id>
        port-forwarding-range {
           range-end <range-end>
        }
        port-reservation {
           {port-blocks <num-blocks> | ports <num-ports>}
        }
        large-scale {
          subscriber-limit <subscriber-limit>
          deterministic
            port-reservation <det-num-ports>
          }
        }
        address-range <start-ip-address> end <end-ip-address> {
        }
      }
    }
```

where:

- **nat-pool-name** — Specifies the name of the NAT pool up to 32 characters max.
- **nat-group-id** — Specifies the NAT group ID. The values are 1 — 4.
- **type** — Specifies the pool type (**large-scale**).
- **num-blocks** — Specifies the number of dynamic port blocks per outside IP address. The values are 1 — 64512.
- **num-ports** — Specifies the number of ports per dynamic block. The values are 1 — 64512.

- **range-end** — Specifies the upper limit of the port range available for static port forwarding. The values are 1023 — 65535.
- **subscriber-limit** — Specifies the maximum number of subscribers per outside IP address.
  - A power of 2 (2^n) number for deterministic NAT: [1,2,4,8,16,32,64,128,256,512,1024,2048, 4096, 8192,16348, 32768]
  - 1..65535 for non-deterministic NAT with default value: 65535
- **det-num-ports** — Specifies the number of ports in a deterministic port block that is allocated and dedicated to a single subscriber during the configuration phase. The values are 1..65535.
- **start-ip-address** — Specifies the first outside IP address in the a.b.c.d format.
- **end-ip-address** — Specifies the last outside IP address in the a.b.c.d format.

Note:

- When the subscriber limit equals 1, each subscriber is mapped to a single outside IP address, though the NAPT port translation function is still performed.
- 1:1 NAT mode in combination with deterministic NAT is not supported.

A NAT policy is configured using the following command:

```
configure {
  service {
    nat {
      nat-policy <nat-policy-name> {
         block-limit <[1..40]>
         pool {
            router-instance {<router-instance | vprn-service-name>}
            name <nat-pool-name>
         }
```

where:

- **nat-policy-name** — Specifies the NAT policy name up to 32 characters max.
- **block-limit** —The maximum number of deterministic plus dynamic port blocks that can be assigned to a single inside IP address. In other words, the maximum number of dynamic port blocks that can be assigned to an inside IP address when the deterministic port block is exhausted equals (block-limit – 1).
- **nat-pool-name** — Specifies the NAT pool name up to 32 characters max.
- **router-instance** — Specifies the router instance the pool belongs to, either by router name or VPRN service ID. The router instance values are "Base" or a VPRN service name of maximum 64 characters.

A NAT inside prefix is configured using the following command:

```
configure [router| service vprn <service-name>]
    nat {
        inside {
            large-scale {
                nat44 {
                    max-subscriber-limit <max>
                    deterministic {
                        prefix-map <source-prefix> nat-policy <nat-policy-name> {
                            map <lsn-sub-address> to <lsn-sub-address> {
                                first-outside-address <outside-ip-address>
                            }
                        }
                    }
```

```
                }
```

where:

- **max-subscriber-limit <max>** — The power of 2 (2^n) number that must match the largest subscriber limit number in a deterministic pool referenced from this inside routing instance. The range for this command is the same as the subscriber limit command under the pool hierarchy. The values are 1,2,4,8 ... 32768.
  - Should be greater than or equal to the largest subscriber limit of all pools referenced by the NAT policies within the corresponding inside routing instance.
  - Must be configured before any inside prefix configuration.
  - Must be 2^n and affects the ingress hashing of deterministic subscribers and also non-deterministic subscribers in case both are configured under the same inside router instance.
- **source-prefix** — A prefix on the inside encompassing subscribers to be deterministically mapped to an outside IP address and port block in the corresponding pool.
  - **<unicast-ipv4-address-and-prefix>** | <expression>
    - <ipv4-prefix> — a.b.c.d (host bits must be 0)
    - <ipv4-prefix-length> — [0..32]
  - <unicast-ipv4-address-and-prefix> | **<expression>**
    - wildcard (*)
    - regular expression (')
    - range [ ]
- *<nat-policy-name>* — Specifies a NAT policy name up to 32 characters in length.

Three cases are now configured to demonstrate the use of deterministic and dynamic port-block usage:

- Case 1 - mapping multiple prefixes of the same VRF into the same outside pool: Mapping multiple prefixes from the same VRF (VPRN 15001) into the same outside pool, routing instance Base.
- Case 2 - mapping multiple prefixes from the same VRF into different outside NAT pools: Mapping multiple prefixes from the same VRF (VPRN 15001) into different outside pools, routing instance VPRN 15002
- Case 3 - mapping overlapping prefixes from different VRFs into the same outside NAT pool: Mapping overlapping prefixes from different VRFs (VPRN 15001 and VPRN 15002) into the same outside pool, routing instance Base.

Case 1 - mapping multiple prefixes of the same VRF into the same outside pool

Figure 31: Case 1 shows that the inside prefixes 10.0.0.0/24 and 10.10.4.0/22 in VPRN-15001 are mapped to the same outside NAT pool in the Base router. All traffic is NATed.

Figure 31: Case 1

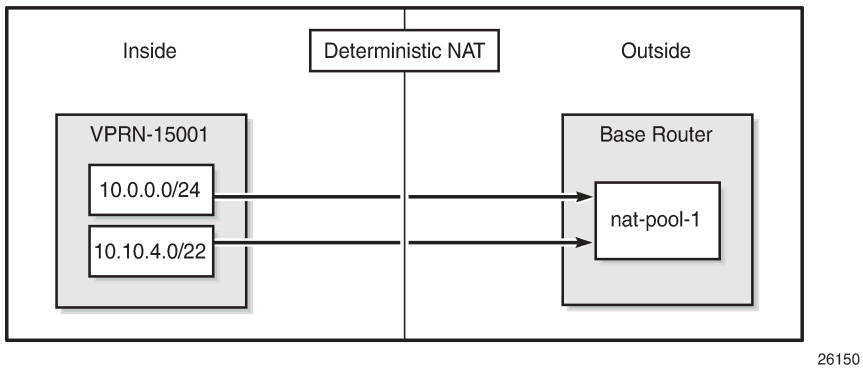

The NAT outside pool is configured as follows:

```
configure {
    router "Base" {
        nat {
            outside {
                pool "nat-pool-1" {
                    admin-state enable
                    type large-scale
                    nat-group 1
                    port-forwarding {
                        range-end 1099
                    }
                    port-reservation {
                        ports 20
                    }
                    large-scale {
                        subscriber-limit 128
                        deterministic {
                            port-reservation 28
                        }
                    }
                    address-range 192.168.0.1 end 192.168.0.100 {
                    }
                }
            }
```

The NAT policy is configured as follows:

```
configure {
    service {
        nat {
            nat-policy "nat-policy-1" {
                block-limit 4
                pool {
                    router-instance "Base"
                    name "nat-pool-1"
                }
            }
        }
```

The NAT inside prefixes are configured as follows:

```
configure {
    service {
        vprn "VPRN-15001" {
            admin-state enable
            service-id 15001
            customer "1"
            nat {
                inside {
                    large-scale {
                        nat44 {
                            max-subscriber-limit 256
                            destination-prefix 0.0.0.0/0 {
                            }
                            deterministic {
                                prefix-map 10.0.0.0/24 nat-policy "nat-policy-1" {
                                    admin-state enable
                                    map 10.0.0.0 to 10.0.0.255 {
                                        first-outside-address 192.168.0.1
                                    }
                                }
                                prefix-map 10.10.4.0/22 nat-policy "nat-policy-1" {
                                    admin-state enable
                                    map 10.10.4.0 to 10.10.7.255 {
                                        first-outside-address 192.168.0.3
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

map statements are automatically created when the prefix is created and it is in **admin-state enable**.

Show commands

The subscriber limit is set to 128 for the 10.0.0.0/24 prefix, so it is broken into two smaller /25 prefixes each. Each of these smaller prefixes are mapped into a specific outside IP address.

The following command shows the first Large Scale NAT (LSN) subscriber of the first /25 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.0.0.0]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824064
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.0.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.1
```

```
-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last subscriber mapping to the same 192.168.0.1 outside IP address has inside address 10.0.0.127.

The following command shows the first LSN subscriber of the second /25 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.128

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.0.0.128]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824192
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.0.0.128/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.2


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last subscriber mapping to the same 192.168.0.2 outside IP address has inside address 10.0.0.255.

The following command shows the base router LSN blocks corresponding to the first inside IP address within the 10.0.0.0/24 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.0

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [1100..1127]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The following command shows the base router LSN blocks corresponding to the last inside IP address within the 10.0.0.0/24 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.255
```

```
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.2 [4656..4683]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The subscriber limit is 128 for the 10.10.4.0/22 prefix, so it is broken into eight /25 prefixes. Each of these smaller prefixes are mapped into a specific outside IP address.

The following command shows the first LSN subscriber of the first /25 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.10.4.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.10.4.0]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824320
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.10.4.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.3


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last subscriber mapping to the same 192.168.0.3 outside IP address has inside address 10.10.4.127.

The following command shows the first LSN subscriber of the last /25 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.10.7.128

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.10.7.128]
NAT policy                  : nat-policy-1
Subscriber ID               : 276825216
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.10.7.128/32
```

```
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.10


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The following command shows the base router LSN blocks corresponding to the first inside IP within 10.10.4.0/24 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.10.4.0

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.3 [1100..1127]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.10.4.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The following command shows the base router LSN blocks corresponding to the last inside IP within 10.10.4.0/24 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.10.7.255

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.10 [4656..4683]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.10.7.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

According to this configuration, each inside IP address has one deterministic block of 28 ports and can have up to three dynamic blocks (block-limit = 4) each of 20 ports, allowing a maximum of 28+3*20 = 88 flows.

Figure 32: Case 1 results

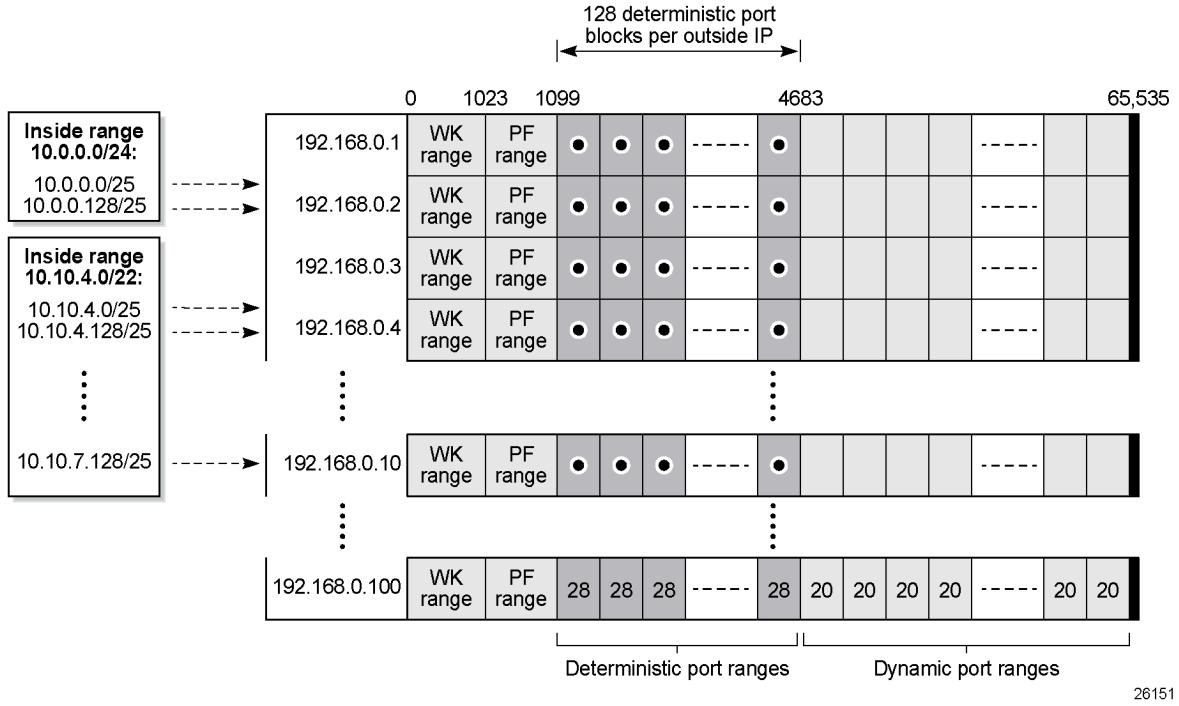

Sending flows

For the inside IP 10.0.0.1, several UDP flows are sent with source IP 10.10.0.1 and different source UDP port numbers. Both the deterministic and dynamic blocks mappings are verified.

Figure 33: Case 1 flows

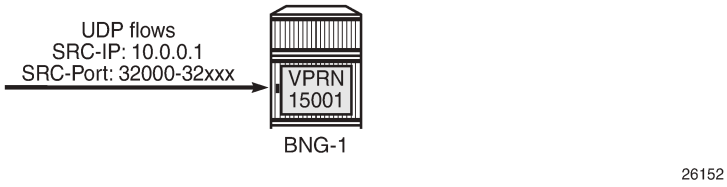

When sending 28 UDP flows or less, all flows are mapped to a single deterministic block because the number of ports in a deterministic block is 28. There is no logging; because no dynamic blocks are used, only the deterministic block is used.

The following command shows the LSN blocks on the outside Base routing instance and the outside ports allocated for the inside IP 10.0.0.1:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
```

```
===============================================================================
192.168.0.1 [1128..1155]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

When increasing the number of flows such that: 29 ≤ number of flows ≤ 48

- In addition to the deterministic block (28 ports), there is an extension by 1 dynamic block of 20 ports (port-reservation=20).
- Logging occurs for the dynamic port block.

With the following command, NAT tmnxNatPlBlockAllocationLsn (2012) events are generated with the default severity:

```
configure {
    log {
        log-events {
            nat event tmnxNatPlBlockAllocationLsn {
                generate true
            }
```

The following command shows the base router LSN blocks and the outside ports allocated to the inside IP address 10.0.0.1: the first block is the deterministic block of 28 ports and the second block is a dynamic block of 20 ports.

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [1128..1155]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4684..4703]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:33:42
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 2
===============================================================================
```

With NAT event 2012 (tmnxNatPlBlockAllocationLsn) generated, logging can be verified using Log 99 which shows the mapping details to the new dynamic block as follows:

```
308 2025/10/13 09:33:42.025 UTC MINOR: NAT #2012 Base NAT
```

```
"{3} Map  192.168.0.1 [4684-4703] ESA-VM 1/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
 at 2025/10/13 09:33:42"
```

When increasing the number of flows such that: 49 ≤ number of flows ≤ 68

- In addition to the deterministic block (28 ports), there is an extension by 2 dynamic blocks of 20 ports each.
- Logging occurs for the dynamic port blocks.

The following command shows LSN blocks on the outside Base routing instance and the outside ports allocated for the inside IP 10.0.0.1:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [1128..1155]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4684..4703]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:33:42
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4704..4723]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:34:24
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 3
===============================================================================
```

Logging is verified using Log 99 (in case event-control NAT events are generated) which shows the mapping details to the new dynamic block as follows:

```
309 2025/10/13 09:34:24.957 UTC MINOR: NAT #2012 Base NAT
"{4} Map  192.168.0.1 [4704-4723] ESA-VM 1/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
 at 2025/10/13 09:34:24"
```

When increasing the number of flows such that: 69 ≤ number of flows ≤ 88

- In addition to the deterministic block (28 ports), there is an extension by 3 dynamic blocks of 20 ports each.
- Logging occurs for the dynamic port blocks.

The following command shows LSN blocks on the outside Base routing instance and the outside ports allocated for the inside IP 10.0.0.1:

```
[/]
```

```
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [1128..1155]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:04:36
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4684..4703]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:33:42
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4704..4723]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:34:24
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [4724..4743]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2025/10/13 09:35:18
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 4
===============================================================================
```

Logging is verified using Log 99 which shows the mapping details to the new dynamic block as follows:

```
310 2025/10/13 09:35:18.679 UTC MINOR: NAT #2012 Base NAT
"{5} Map  192.168.0.1 [4724-4743] ESA-VM 1/1 -- 276824065 classic-lsn-sub %1 vprn15001 10.0.0.1
 at 2025/10/13 09:35:18"
```

When increasing the number of flows such that the number of flows > 88

- No more extension by dynamic blocks (block limit = 4) allowed.
- Any flows more than 88 are dropped and the relevant NAT statistics incremented.

To verify NAT statistics, first check the NAT group, member, and ESA (or MS-ISA) associated with the outside IP 192.168.0.1/32:

```
[/]
A:admin@BNG-1# show router route-table 192.168.0.1/32

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                          Type    Proto     Age        Pref
      Next Hop[Interface Name]                                  Metric
-------------------------------------------------------------------------------
192.168.0.1/32                              Remote  NAT       00h36m01s  0
       NAT outside to esa 1 vm 1                                    0
```

```
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

To check which group and member this ESA belongs to, the following command can be used:

```
[/]
A:admin@BNG-1# show isa nat-group 1 members

===============================================================================
ISA Group 1 members
===============================================================================
Group Member State          MDA/VM   Addresses  Blocks     Se-% Hi Se-Prio
-------------------------------------------------------------------------------
1     1      active         esa-1/1  10         1490       < 1  N  0
-------------------------------------------------------------------------------
No. of members: 1
===============================================================================
```

The following command can be used to verify relevant statistics for this NAT group/member:

```
[/]
A:admin@BNG-1# show isa nat-group 1 statistics esa-vm 1/1 | match flow
no matching flow                                    : 0
max flow exceeded                                   : 0
TCP no flow for RST                                 : 0
TCP no flow for FIN                                 : 0
TCP no flow                                         : 0
flow log failed                                     : 0
new flow                                            : 80
found flow                                          : 122107
flow create logged                                  : 0
flow delete logged                                  : 0
flow log pkt tx                                     : 0
flow create failed, key ambiguous                   : 0
flow create failed, conflicting policies            : 0
too many frag. lists for flow                       : 0
flow log pkt rate limit dropped                     : 0
```

Case 2 - mapping multiple prefixes from the same VRF into different outside NAT pools

Figure 34: Case 2 shows that prefix 10.1.0.0/23 in VPRN-15001 is mapped to outgoing NAT pool "nat-pool-2" in VPRN-15002 while prefix 10.2.0.0/22 in VPRN-15001 is mapped to "nat-pool-3" in VPRN-15002. All traffic is NATed.

Figure 34: Case 2

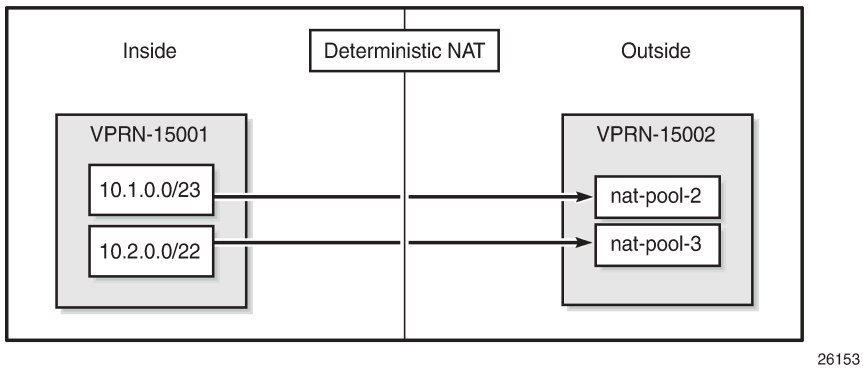

The NAT outside pools are configured in VPRN-15002, as follows:

```
configure {
    service {
        vprn "VPRN-15002" {
            admin-state enable
            service-id 15002
            customer "1"
            nat {
                outside {
                    pool "nat-pool-2" {
                        admin-state enable
                        type large-scale
                        nat-group 1
                        port-reservation {
                            ports 12
                        }
                        large-scale {
                            subscriber-limit 256
                            deterministic {
                                port-reservation 15
                            }
                        }
                        address-range 192.168.2.1 end 192.168.2.200 {
                        }
                    }
                    pool "nat-pool-3" {
                        admin-state enable
                        type large-scale
                        nat-group 1
                        port-forwarding {
                            range-end 1149
                        }
                        port-reservation {
                            ports 50
                        }
                        large-scale {
                            subscriber-limit 64
                            deterministic {
                                port-reservation 60
                            }
                        }
                        address-range 192.168.3.1 end 192.168.3.200 {
                        }
                    }
```

```
                }
            }
```

The NAT policies are configured as follows:

```
configure {
    service {
        nat {
            nat-policy "nat-policy-2" {
                block-limit 4
                pool {
                    router-instance "VPRN-15002"
                    name "nat-pool-2"
                }
            }
            nat-policy "nat-policy-3" {
                block-limit 2
                pool {
                    router-instance "VPRN-15002"
                    name "nat-pool-3"
                }
            }
```

The NAT inside prefixes are configured as follows:

```
configure {
    service {
        vprn "VPRN-15001" {
            admin-state enable
            service-id 15001
            customer "1"
            nat {
                inside {
                    large-scale {
                        nat44 {
                            max-subscriber-limit 256
                            destination-prefix 0.0.0.0/0 {
                            }
                            deterministic {
                                prefix-map 10.1.0.0/23 nat-policy "nat-policy-2" {
                                    admin-state enable
                                    map 10.1.0.0 to 10.1.1.255 {
                                        first-outside-address 192.168.2.1
                                    }
                                }
                                prefix-map 10.2.0.0/22 nat-policy "nat-policy-3" {
                                    admin-state enable
                                    map 10.2.0.0 to 10.2.3.255 {
                                        first-outside-address 192.168.3.1
                                    }
                                }
                            }
                        }
                    }
                }
```

Show commands

The subscriber limit corresponding to the 10.1.0.0/23 prefix is 256, so the 10.1.0.0/23 prefix is broken into two /24 prefixes. Each of these smaller prefixes are mapped into a specific outside IP address.

The following command shows the first LSN subscriber of the first /24 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.1.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.1.0.0]
NAT policy                  : nat-policy-2
Subscriber ID               : 276825344
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.1.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.2.1


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last subscriber mapping to the same 192.168.2.1 outside IP address has inside address 10.1.0.255.

The following command shows the first LSN subscriber of the second /24 prefix for inside routing instance 15001:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.1.1.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.1.1.0]
NAT policy                  : nat-policy-2
Subscriber ID               : 276825600
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.1.1.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.2.2


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last subscriber mapping to the same 192.168.2.2 outside IP address has inside address 10.1.1.255.

The following command shows the VPRN-15002 LSN blocks corresponding to the first inside IP address within the 10.1.0.0/23 prefix:

```
[/]
A:admin@BNG-1# show router 15002 nat lsn-blocks inside-ip 10.1.0.0
```

```
===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.1 [1024..1038]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2025/10/13 09:45:31
Inside router                       : vprn15001
Inside IP address                   : 10.1.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The following command shows the VPRN-15002 LSN blocks corresponding to the last inside IP address within the 10.1.0.0/23 prefix:

```
[/]
A:admin@BNG-1# show router 15002 nat lsn-blocks inside-ip 10.1.1.255

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.2 [4849..4863]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2025/10/13 09:45:31
Inside router                       : vprn15001
Inside IP address                   : 10.1.1.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The subscriber limit corresponding to the 10.2.0.0/22 prefix is 64,so the 10.2.0.0/22 prefix is broken into sixteen /26 prefixes. Each of these /26 prefixes is mapped to a specific outside IP address.

The following command shows the first LSN subscriber for the inside routing instance 15001 for the first /26 prefix:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.2.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.2.0.0]
NAT policy                  : nat-policy-3
Subscriber ID               : 276825856
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.2.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.3.1


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
```

```
===============================================================================
```

The last inside address mapping to the 192.168.3.1 outside address is 10.2.0.63.

The following command shows the first LSN subscriber for the inside routing instance 15001 for the last /26 prefix:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.2.3.192

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.2.3.192]
NAT policy                  : nat-policy-3
Subscriber ID               : 276826816
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.2.3.192/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.3.16


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last inside address mapping to the 192.168.3.16 outside address is 10.2.3.255.

The following command shows the VPRN-15002 LSN blocks corresponding to the first inside IP address within the 10.2.0.0/22 prefix:

```
[/]
A:admin@BNG-1# show router 15002 nat lsn-blocks inside-ip 10.2.0.0

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.3.1 [1150..1209]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2025/10/13 09:45:31
Inside router                       : vprn15001
Inside IP address                   : 10.2.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The following command shows the VPRN-15002 LSN blocks corresponding to the last inside IP within the 10.2.0.0/22 prefix:

```
[/]
A:admin@BNG-1# show router 15002 nat lsn-blocks inside-ip 10.2.3.255

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
```

```
192.168.3.16 [4930..4989]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2025/10/13 09:45:31
Inside router                       : vprn15001
Inside IP address                   : 10.2.3.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

According to this configuration, for the 10.1.0.0/23 prefix, each inside IP address has one deterministic block of 15 ports and can have up to three dynamic blocks (block limit =4) each of 12 ports, allowing for a maximum of 15+3*12 = 51 flows.

Figure 35: Case 2: Prefix 10.1.0.0/23 results

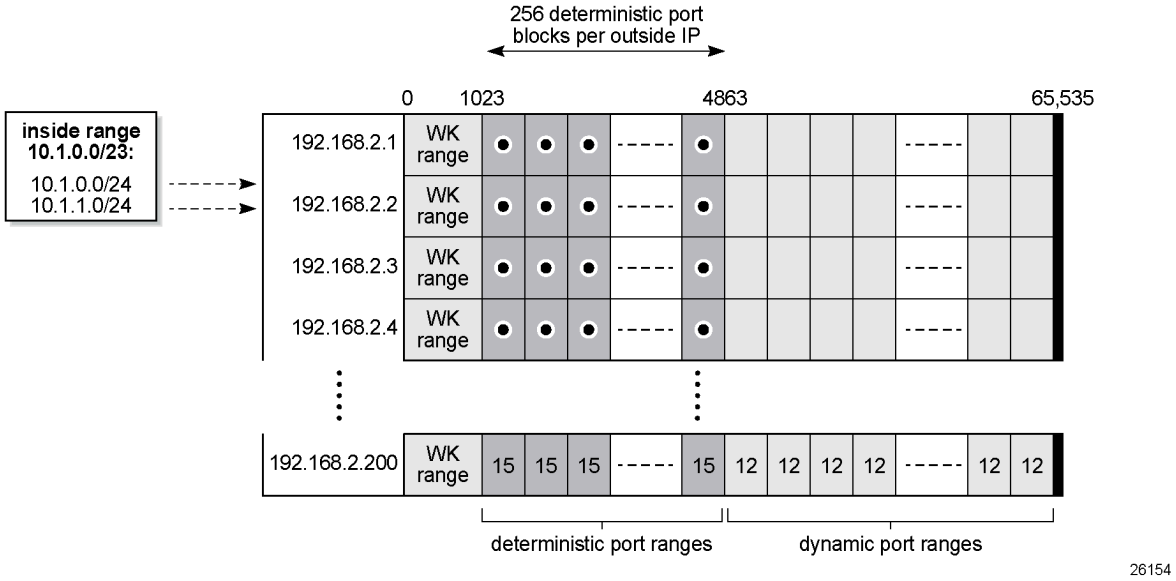

According to this configuration, for the 10.2.0.0/22 prefix, each inside IP address has one deterministic block of 60 ports, and can have up to one dynamic block (block limit =2) of 50 ports, allowing for a maximum of 60+50 = 110 flows.

Figure 36: Case 2: Prefix 10.2.0.0/22 results

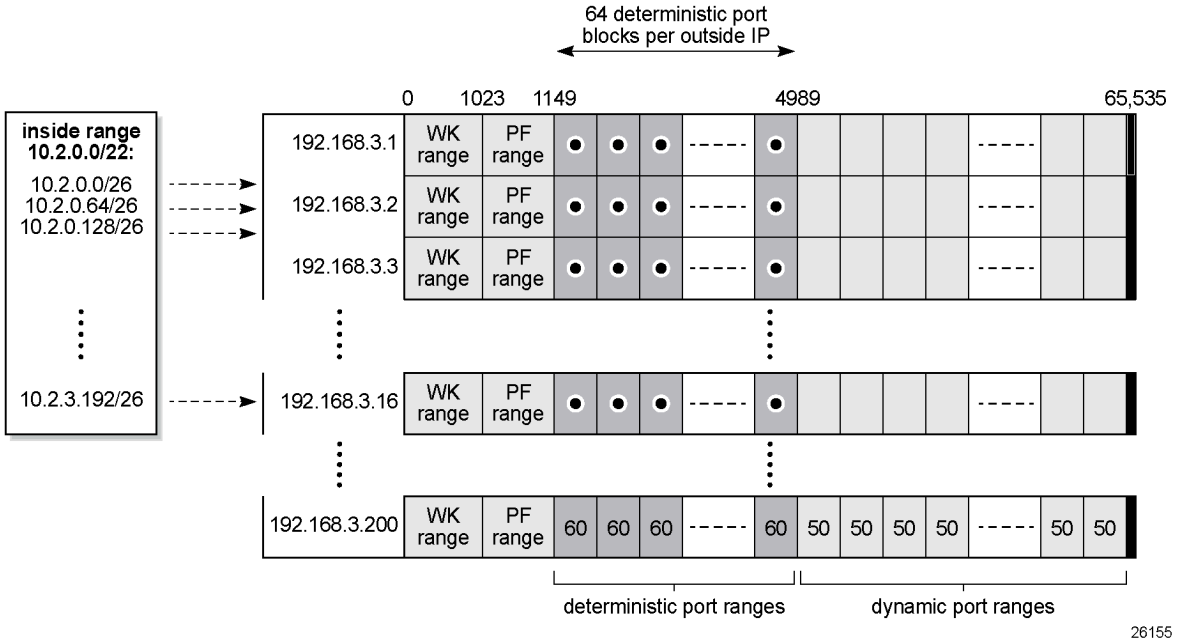

Case 3 - mapping overlapping prefixes from different VRFs into the same outside NAT pool

Figure 37: Case 3 shows that overlapping prefixes from VPRN-15001 and VPRN-15002 are mapped to the same outside NAT pool in the Base router. All traffic is NATed.

Figure 37: Case 3

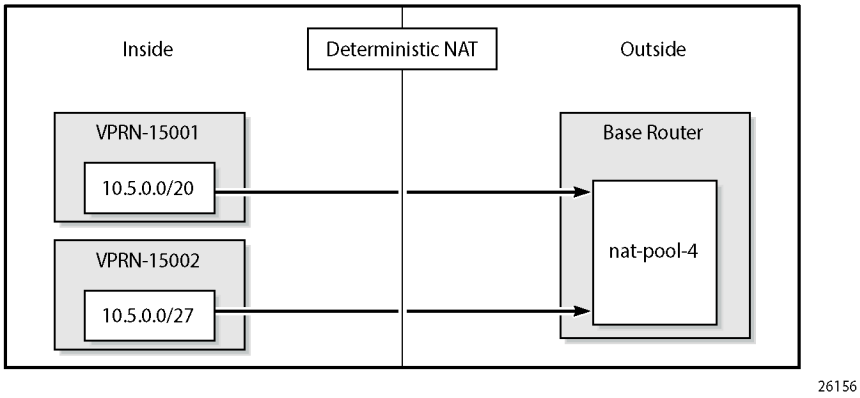

The NAT outside pool is configured as follows:

```
configure {
    router "Base" {
        nat {
```

```
            outside {
                pool "nat-pool-4" {
                    admin-state enable
                    type large-scale
                    nat-group 1
                    port-forwarding {
                        range-end 1049
                    }
                    port-reservation {
                        ports 50
                    }
                    large-scale {
                        subscriber-limit 64
                        deterministic {
                            port-reservation 62
                        }
                    }
                    address-range 192.168.4.1 end 192.168.4.100 {
                    }
                }
```

The NAT policy is configured as follows:

```
configure {
    service {
        nat {
            nat-policy "nat-policy-4" {
                block-limit 4
                pool {
                    router-instance "Base"
                    name "nat-pool-4"
                }
            }
```

The NAT inside prefix is configured as follows:

```
configure {
    service {
        vprn "VPRN-15001" {
            admin-state enable
            service-id 15001
            customer "1"
            nat {
                inside {
                    large-scale {
                        nat44 {
                            max-subscriber-limit 256
                            destination-prefix 0.0.0.0/0 {
                            }
                            deterministic {
                                prefix-map 10.5.0.0/20 nat-policy "nat-policy-4" {
                                    admin-state enable
                                    map 10.5.0.0 to 10.5.15.255 {
                                        first-outside-address 192.168.4.1
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
        vprn "VPRN-15002" {
```

```
            admin-state enable
            service-id 15002
            customer "1"
            nat {
                inside {
                    large-scale {
                        nat44 {
                            max-subscriber-limit 128
                            destination-prefix 0.0.0.0/0 {
                            }
                            deterministic {
                                prefix-map 10.5.0.0/27 nat-policy "nat-policy-4" {
                                    admin-state enable
                                    map 10.5.0.0 to 10.5.0.31 {
                                        first-outside-address 192.168.4.65
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
```

Show commands

For the 10.5.0.0/20 prefix on VPRN 15001, the subscriber-limit is 64. The 10.5.0.0/20 prefix is broken into 64 smaller /26 prefixes, each to be mapped into a specific outside IP address.

The following command shows the first LSN subscriber for the inside routing instance 15001 of the first /26 prefix:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15001

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.0.0]
NAT policy                  : nat-policy-4
Subscriber ID               : 276826880
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.5.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.1


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last inside address mapping to the 192.168.4.1 outside address is 10.5.0.63.

The following command shows the first Large Scale NAT (LSN) subscriber for the inside routing instance 15001 of the last /26 prefix:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.5.15.192 inside-router 15001

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.15.192]
NAT policy                  : nat-policy-4
Subscriber ID               : 276830912
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.5.15.192/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.64


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The last inside address mapping to the 192.168.4.64 outside address is 10.5.15.255.

The following command shows the base router LSN blocks corresponding to the first inside IP address within the 10.5.0.0/20 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15001

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.1 [1050..1111]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2025/10/13 09:52:33
Inside router                       : vprn15001
Inside IP address                   : 10.5.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show the base router LSN blocks corresponding to the last inside IP address within the 10.5.0.0/20 prefix, the following command can be used:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.5.15.255 inside-router 15001

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.64 [4956..5017]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2025/10/13 09:52:33
Inside router                       : vprn15001
```

```
Inside IP address                    : 10.5.15.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

For the 10.5.0.0/27 prefix in VPRN 15002, the subscriber limit is 64. The 10.5.0.0/27 prefix is mapped into one outside IP address.

The following command shows the first LSN subscriber for the inside routing instance 15002 of the 10.5.0.0/27 prefix:

```
[/]
A:admin@BNG-1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15002

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.0.0]
NAT policy                  : nat-policy-4
Subscriber ID               : 276830976
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15002
Inside IP address prefix    : 10.5.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.65


-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
```

The following command shows the LSN blocks corresponding to the first inside IP address within the 10.5.0.0/27 prefix:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15002

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [1050..1111]
Pool                                 : nat-pool-4
Policy                               : nat-policy-4
Started                              : 2025/10/13 09:52:33
Inside router                        : vprn15002
Inside IP address                    : 10.5.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

The following command shows the LSN blocks for the last inside IP address within the 10.5.0.0/27 prefix

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.5.0.31 inside-router 15002

===============================================================================
```

```
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [2972..3033]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2025/10/13 09:52:33
Inside router                       : vprn15002
Inside IP address                   : 10.5.0.31

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

According to this configuration, each inside IP address within VPRN 15001 has one deterministic block of 62 ports and can have up to three dynamic blocks (block-limit =4) of 50 ports each, allowing a maximum of 62+3*50 = 212 flows.

According to this configuration each inside IP address within VPRN 15002 has one deterministic block of 62 ports and can have up to three dynamic blocks (block limit =4) of ports each, allowing a maximum of 62+3*50 = 212 flows.

For VPRN 15002, because the number of LSN subscribers (32) is less than the number of deterministic blocks (64), 32 deterministic blocks are wasted, specifically 32*62 = 1984 ports are wasted which is not good in terms of capacity planning.

Figure 38: Case 3 results

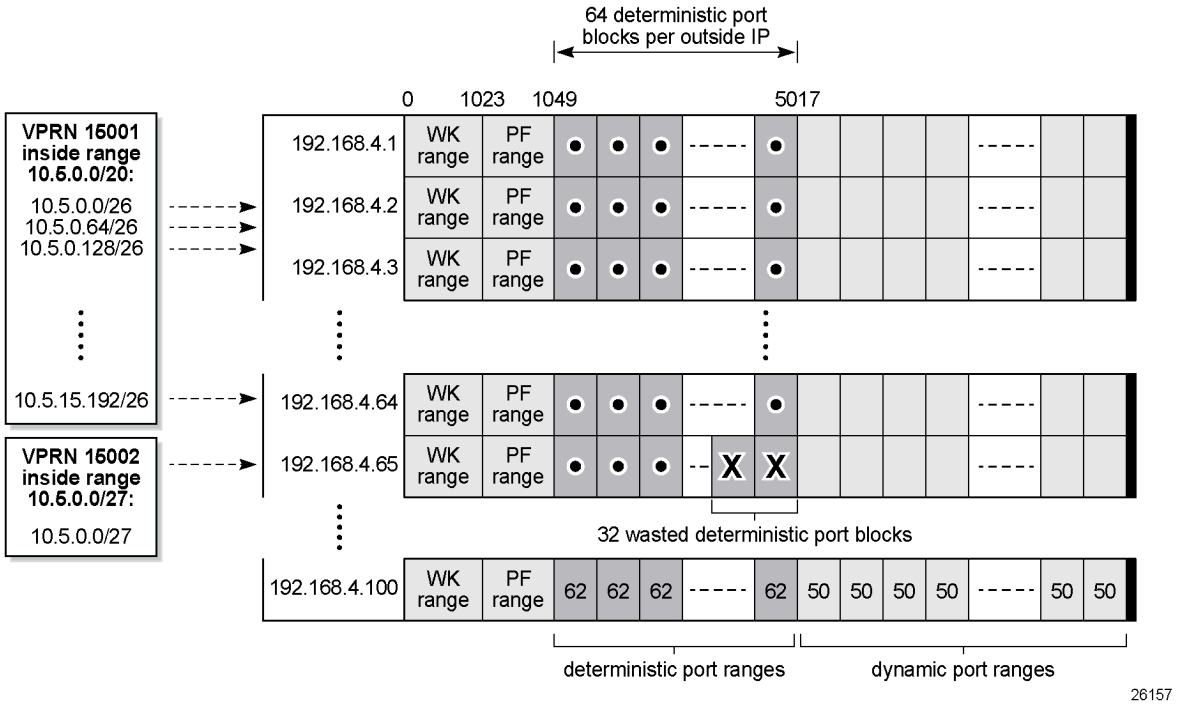

Inverse mapping

In deterministic LSN44, the inside IP addresses are mapped to outside IP addresses and corresponding port blocks based on a deterministic algorithm. The inverse mapping that reveals the subscriber identity behind the NAT is based on the reversal of this algorithm.

Figure 39: Inverse mapping approach shows that inverse mappings can be done either online or offline:

- Online — Locally on the SR OS node, via CLI (MIB)
- Offline — Externally, via a Python script. The purpose of such an offline approach is to provide fast queries without accessing the SR OS node.

Figure 39: Inverse mapping approach

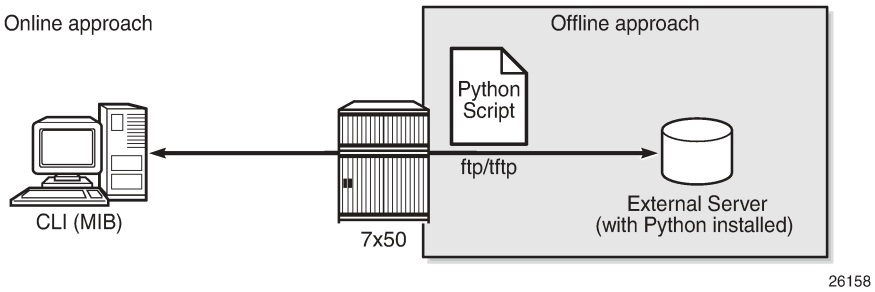

Online approach

A **tools** command is available which shows the reverse mapping (outside to inside) for deterministic NAT instead of using logging.

```
tools dump nat deterministic-mapping outside-ip <ipv4-address | expression> router <router-
instance> outside-port <[1..65535]>
```

Using Case 3 - mapping overlapping prefixes from different VRFs into the same outside NAT pool as an example, to obtain (inside IP, inside routing instance), the inverse mapping for a specific (outside IP, outside routing instance, outside port) is done as follows:

```
[/]
A:admin@BNG-1# tools dump nat deterministic-mapping outside-ip 192.168.4.1 router "Base"
 outside-port 1050
classic-lsn-sub inside router 15001 ip 10.5.0.0 -- outside router Base ip 192.168.4.1 port 1050
 at Mon Oct 13 10:00:37 UTC 2025
```

```
[/]
A:admin@BNG-1# tools dump nat deterministic-mapping outside-ip 192.168.4.65 router "Base"
 outside-port 1050
classic-lsn-sub inside router 15002 ip 10.5.0.0 -- outside router Base ip 192.168.4.65 port
 1050 at Mon Oct 13 10:00:37 UTC 2025
```

Offline approach

The purpose of such an offline approach is to provide fast queries without the need to directly query the SR OS node. This is achieved by generating and exporting a Python script for reverse querying, which is a manual operation that needs to be repeated every time there is configuration change in deterministic NAT.

The script is exported (manually) to the external system.

To configure remotely the location for the Python script, the following command is used:

```
configure {
    service {
        nat {
            deterministic-script {
                location <remote-url>
```

The remote URL is a remote location where the script is stored, such as:

```
[{ftp://|tftp://}<login>:<pswd>@ <remote-location>/][<file-path>]
```

The remote URL has a maximum length of 180 characters.

When the script location is specified, the script can be exported to that location using the following command:

```
admin nat save-deterministic-script
```

Using the following command the status of the script can be checked, and whether it is necessary to re-save (export) the script or not:

```
[/]
A:admin@BNG-1# show service nat deterministic-script

===============================================================================
Deterministic NAT script data
===============================================================================
Location                    : ftp://*:*@192.99.1.249/./images/detnat.py
Save needed                 : no
Last save result            : success
Last save time              : 2025/10/13 10:01:13
===============================================================================
```

The external system must have the Python scripting language installed with the following modules: getopt, math, os, socket, and sys.

The Python script can then be run on the external server; the parameters are as follows:

```
(ANPM2VM1-HOST):~/images > ./detnat.py
Error: need exactly one of --forward or --backward arguments

Usage: detnat.py DIRECTION PARAMETERS
Perform forward or backard NAPT according to the configured deterministic rules.

DIRECTION:
  -f, --forward             Translate from inside to outside address/port
  -b, --backward            Translate from outside to inside address/port

PARAMETERS:
```

```
  -a, --address=IP-ADDRESS  The address to translate. IPv6 addresses can be
                            specified in shorthand or full notation.
  -p, --port=PORT           The outside port in case of backward translation.
  -s, --service=SERVICE-ID  The service where the IP-ADDRESS originates from.
                            This is the inside service in case of forward
                            translation and the outside service in case of
                            backward translation.
                            To specify the base router, this option must be
                            omitted.
  -P, --policy=POLICY-NAME  The nat-policy used by the subscriber.

  -h, --help                Show this help message
```

where detnat.py is the name of the python script previously exported.

As an example of a forward query:

```
(ANPM2VM1-HOST):~/images > ./detnat.py -f -s 15001 -a 10.0.0.1
classic-lsn-sub is using policy nat-policy-1, has public ip address 192.168.0.1 from base
 router and is using ports [1128 - 1155]
```

As an example of a reverse query:

```
(ANPM2VM1-HOST):~/images > ./detnat.py -b -s 0 -a 192.168.4.1 -p 1525
classic-lsn-sub has private ip address 10.5.0.7 from service 15001 and is using policy nat-
policy-4
```

Simultaneous support of deterministic and non-deterministic NAT

Deterministic NAT can be used simultaneously with non-deterministic NAT within the same inside routing instance. However, they cannot share the same pool.

An outside pool can be only deterministic (although expandable by dynamic ports blocks) or non-deterministic at any time. A non-deterministic pool is a pool that contains dynamic port blocks only.

The following shows a configuration using deterministic NAT simultaneously with non-deterministic NAT.

The NAT outside pools are configured as follows:

```
configure {
    router "Base" {
        nat {
            outside {
                pool "nat-pool-1" {
                    admin-state enable
                    type large-scale
                    nat-group 1
                    port-forwarding {
                        range-end 1099
                    }
                    port-reservation {
                        ports 20
                    }
                    large-scale {
                        subscriber-limit 128
                        deterministic {
                            port-reservation 28
                        }
                    }
                    address-range 192.168.0.1 end 192.168.0.100 {
```

```
                    }
                }
                pool "nat-pool-Non-Deterministic" {
                    admin-state enable
                    type large-scale
                    nat-group 1
                    address-range 192.168.7.1 end 192.168.7.100 {
                    }
                }
            }
```

The NAT policies are configured as follows:

```
configure {
    service {
        nat {
            nat-policy "nat-policy-1" {
                block-limit 4
                pool {
                    router-instance "Base"
                    name "nat-pool-1"
                }
            }
            nat-policy "nat-policy-Non-Deterministic" {
                pool {
                    router-instance "Base"
                    name "nat-pool-Non-Deterministic"
                }
            }
```

The NAT inside prefixes are configured as follows:

```
configure {
    service {
        vprn "VPRN-15001" {
            admin-state enable
            service-id 15001
            customer "1"
            nat {
                inside {
                    large-scale {
                        nat-policy "nat-policy-Non-Deterministic"
                        nat44 {
                            max-subscriber-limit 256
                            destination-prefix 0.0.0.0/0 {
                            }
                            deterministic {
                                prefix-map 10.0.0.0/24 nat-policy "nat-policy-1" {
                                    admin-state enable
                                    map 10.0.0.0 to 10.0.0.255 {
                                        first-outside-address 192.168.0.1
                                    }
                                }
                            }
                        }
                    }
                }
            }
```

In this example, the inside IP prefixes that do not match any of the deterministic prefixes are NATed using a non-deterministic pool.

Figure 40: Sending flows: deterministic + non-deterministic NAT

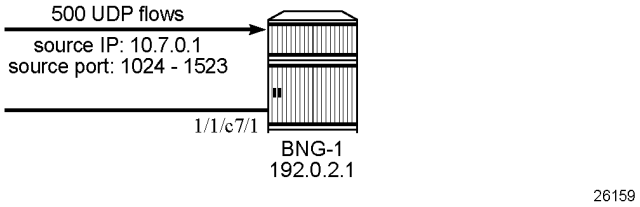

The following command shows which NAT pool and NAT policy are used for NATing the inside IP 10.7.0.1:

```
[/]
A:admin@BNG-1# show router nat lsn-blocks inside-ip 10.7.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.7.100 [1024..1055]
Pool                                : nat-pool-Non-Deterministic
Policy                              : nat-policy-Non-Deterministic
Started                             : 2025/10/13 10:06:06
Inside router                       : vprn15001
Inside IP address                   : 10.7.0.1

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Conclusion

This example provides the commands required for configuring deterministic LSN44 NAT. Both deterministic as well as non-deterministic NAT are supported, with simultaneous operation being possible.

Inverse query can be done online or offline to retrieve the NAT mappings. Logging is not needed as long as there are no dynamic blocks assigned to LSN44 subscribers.

IP/GRE Termination

This chapter provides configuration and troubleshooting commands for IP/GRE termination.

Topics in this chapter include:

Applicability

The chapter was initially written for SR OS Release 9.0.R8. The MD-CLI in the current edition corresponds to SR OS Release 25.3.R2.

The IP GRE tunnel termination configuration described in this chapter requires an MS-ISA. IP GRE tunnels without ISA are beyond the scope of this chapter.

Overview

A common use case for IP/GRE tunneling is remote access to a VPRN over a public IP network because IP/GRE tunneling allows encapsulated packets to follow a path based on the outer IP header which is useful when the inner IP packet cannot or should not be forwarded natively over this path.

GRE allows packets of one protocol, the payload protocol, to be encapsulated by packets of another protocol, called the delivery protocol. Figure 41: GRE packet format shows the GRE packet format with an outer delivery header, GRE header, and payload packet:

Figure 41: GRE packet format

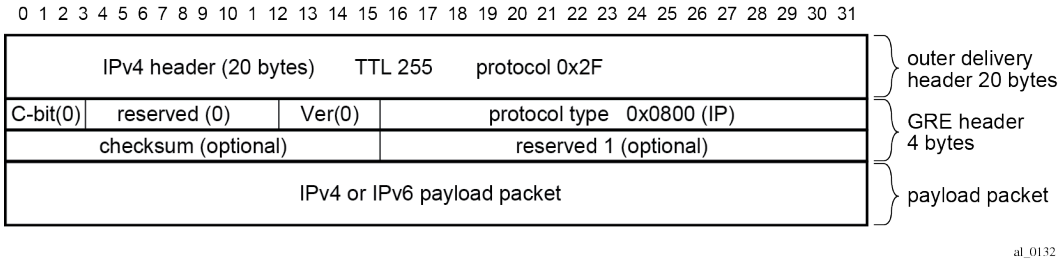

The outer delivery and GRE header for outgoing traffic is as follows.

- Outer delivery header
  - The source address in the IPv4 delivery header is the configured source address.

  - The destination address in the IPv4 delivery header is the configured remote IP (or the backup remote IP) address.
  - The IP protocol value in the IPv4 delivery header is 0x2F or 47 (GRE).
  - The DSCP in the IPv4 outer delivery header is:
    - set to the value configured for the tunnel
    - otherwise, the DSCP value from the payload packet is copied into the outer delivery header.
  - The TTL in the IPv4 outer delivery header is set to 255.
- GRE header
  - The checksum (C) bit in the GRE header is set to 0 (no checksum present).
  - The version in the GRE header is 0.
  - The protocol type is 0x0800 for IPv4.

The outer delivery and GRE header for incoming traffic is as follows:

- Outer delivery header
  - If the packet is a fragment (more fragments=1, non-zero fragment offset), it is dropped.
  - If the checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
  - If the version in the GRE header is not 0, the packet is discarded.
  - If the source and destination address pair in the IPv4 delivery header is any other combination, the packet is dropped.
- GRE header
  - If the checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
  - If the version in the GRE header is not 0, the packet is discarded.

Implementation

Encapsulation, de-encapsulation and other datapath operations related to IP/GRE are handled by the ISA-tunnel MDA.

For GRE tunnels configured as SDPs (which are not covered in this chapter), no ISA-tunnel MDA is required.

Figure 42: Implementation

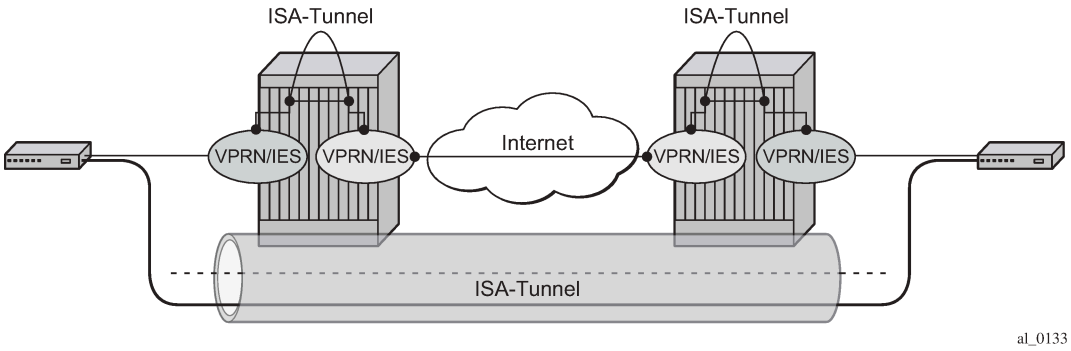

SR OS nodes initially supported the IP/GRE tunnels with static routes and BGP. IP/GRE tunnels have been enhanced by adding support for OSPF and BFD on private tunnel interfaces (used with static routes, OSPF, or BGP) and GRE protection by tunneling into an IPSec tunnel.

Figure 43: IP/GRE over IPSec tunnel

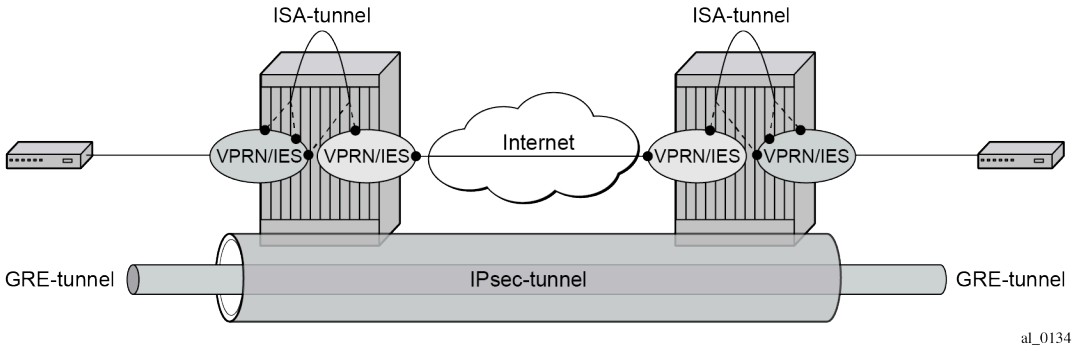

Configuration

ISA-tunnel MDA

The ISA-tunnel MDA supports IP/GRE and IPSec tunnels and is configured as follows:

```
# on PE-1:
configure {
    card 1 {
        mda 2 {
            mda-type isa2-tunnel
        }
    }
    card 2 {
        mda 2 {
            mda-type isa2-tunnel
        }
    }
```

The following command checks the MDA configuration:

```
[/]
A:admin@PE-1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned Type                          Admin     Operational
               Equipped Type (if different)           State     State
-------------------------------------------------------------------------------
1     1     p10-10g-sfp                               up        up
      2     isa2-tunnel                               up        up
               p-isa2-ms
2     1     p10-10g-sfp                               up        up
      2     isa2-tunnel                               up        up
               p-isa2-ms
===============================================================================
```

Tunnel groups and tunnel group restrictions

The first step of the GRE tunnel configuration is to configure a tunnel group.

A tunnel group can have one tunnel ISA designated primary and optionally one tunnel-ISA designated backup. When a GRE tunnel is created, it is assigned to the primary tunnel-ISA in its tunnel group. If the primary tunnel-ISA fails, the backup tunnel-ISA (if not already claimed by another tunnel group) takes over for the failed card.

```
[ex:/configure isa]
A:admin@PE-1# tunnel-group 1 ?

 tunnel-group

 Immutable fields      - isa-scale-mode

 admin-state           - Administrative state of the ISA tunnel group
 apply-groups          - Apply a configuration group at this level
 apply-groups-exclude  - Exclude a configuration group at this level
 description           - Text description
 ipsec-responder-only  - Act as an IKE responder except upon MC-IPsec switchover
 isa-scale-mode        ^ Tunnel limit on each ISA for the tunnel group
 reassembly            + Enter the reassembly context
 spi-range-index       - SPI range index to use for the tunnel group
 stats-collection      + Enter the stats-collection context
 strict-esp-sequence-  - Enable strict ESP sequence number ordering
  number-ordering

 Choice: redundancy
 multi-active         :+ Enable the multi-active context
 backup               :- IPsec module configured in the slot to the IPsec group
 primary               - Primary ISA IPsec module assigned for the tunnel group



# on PE-1:
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
```

```
            primary 1/2
            backup 2/2
        }
```

The failed tunnels are re-established using a cold-standby on the backup tunnel-ISA. Cold-standby means the backup tunnel-ISA has no state or configuration information about the tunnels before the failure.

A tunnel ISA cannot be primary for more than one tunnel group:

```
*[ex:/configure isa tunnel-group 2]
A:admin@PE-1# primary 1/2

*[ex:/configure isa tunnel-group 2]
A:admin@PE-1# commit
MINOR: IPSECGRPMGR #1003: configure isa tunnel-group 2 primary
- The specified MDA is primary in another Tunnel Group - configure isa tunnel-group 1 primary
MINOR: IPSECGRPMGR #1003: configure isa tunnel-group 1 primary
- The specified MDA is primary in another Tunnel Group - configure isa tunnel-group 2 primary
```

A tunnel ISA cannot be primary in one tunnel group and backup in another tunnel group:

```
*[ex:/configure isa tunnel-group 2]
A:admin@PE-1# primary 2/2

*[ex:/configure isa tunnel-group 2]
A:admin@PE-1# backup 1/2

*[ex:/configure isa tunnel-group 2]
A:admin@PE-1# commit
MINOR: IPSECGRPMGR #1003: configure isa tunnel-group 2 backup
- The specified MDA is primary in another Tunnel Group - configure isa tunnel-group 1 primary
MINOR: IPSECGRPMGR #1003: configure isa tunnel-group 1 backup
- The specified MDA is primary in another Tunnel Group - configure isa tunnel-group 2 primary
```

The following commands shows the ISA tunnel group (after tunnel group 2 has been removed):

```
[/]
A:admin@PE-1# show isa tunnel-group

===============================================================================
ISA Tunnel Groups
===============================================================================
Tunnel    PrimaryIsa            BackupIsa    ActiveIsa    Admin     Oper
GroupId                                                   State     State
-------------------------------------------------------------------------------
1         1/2                   2/2          1/2          Up        Up
-------------------------------------------------------------------------------
No. of ISA Tunnel Groups: 1
===============================================================================
```

The following command shows the number of the IP (GRE) tunnels, after configuring IES and VPRN services with tunnel interfaces:

```
[/]
A:admin@PE-1# show ip tunnel count
-------------------------------------------------------------------------------
IP Tunnels: 2
-------------------------------------------------------------------------------
```

The following command shows all IP tunnels:

```
[/]
A:admin@PE-1# show ip tunnel

===============================================================================
IP Tunnels
===============================================================================
TunnelName                     SapId                          SvcId      Admn
 Local Address                                                DlvrySvcId Oper
  OperRemoteAddress
-------------------------------------------------------------------------------
gre-tunnel-1                   tunnel-1.private:1             1          Up
 192.168.1.1                                                  IES 2      Up
  192.168.2.1
protected-gre-tunnel           tunnel-1.private:5             3          Up
 192.168.11.1                                                 VPRN 3     Up
  192.168.22.1
-------------------------------------------------------------------------------
IP Tunnels: 2
===============================================================================
```

The detailed tunnel information is as follows:

```
[/]
A:admin@PE-1# show ip tunnel "gre-tunnel-1"

===============================================================================
IP Tunnel Configuration Detail
===============================================================================
Service Id       : 1                   Sap Id           : tunnel-1.private:1
Tunnel Name      : gre-tunnel-1
Description      : None
GRE Header       : Yes
Delivery Service : IES 2
GRE Keys Set     : False
GRE Send Key     : N/A                 GRE Receive Key  : N/A
Admin State      : Up                  Oper State       : Up
Source Address   : 192.168.1.1
Remote Address   : 192.168.2.1
Backup Address   : (Not Specified)
Oper Remote Addr : 192.168.2.1
DSCP             : None
Reassembly       : inherit
Clear DF Bit     : false               IP MTU           : max
Encap IP MTU     : max
Pkt Too Big      : true
Pkt Too Big Num  : 100                 Pkt Too Big Intvl: 10 secs
Frag Required    : true
Frag Req Count   : 100                 Frag Req Interval: 10 secs
Propagate IPv6 P*: true
Propagate IPv4 P*: true
Oper Flags       : None
Transport Profile: (Not Specified)
Last Oper Changed: 06/25/2025 07:31:30
Host ISA         : 1/2
TCP MSS Adjust
    Public       : Disabled
    Private      : Disabled

-------------------------------------------------------------------------------
Target Address Table
-------------------------------------------------------------------------------
```

```
Destination IP                          IP Resolved Status
-------------------------------------------------------------------------------
10.0.0.2                                Yes
-------------------------------------------------------------------------------

===============================================================================
IP Tunnel Statistics: gre-tunnel-1
===============================================================================
Errors Rx        : 0                    Errors Tx        : 0
Pkts Rx          : 111                  Pkts Tx          : 108
Bytes Rx         : 7963                 Bytes Tx         : 7819
Key Ignored Rx   : 0                    Too Big Tx       : 0
Seq Ignored Rx   : 0
Vers Unsup. Rx   : 0
Invalid Chksum Rx: 0
Key Mismatch Rx  : 0
===============================================================================

===============================================================================
Fragmentation Statistics
===============================================================================
Encapsulation Overhead                 : 24
Temporary Private MTU                  : max
Pre-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
Post-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Interfaces

The interface toward the Internet (or WAN):

- can be a network interface, a VPRN interface, or an IES interface
- provides IP reachability

The tunnel public interface:

- can be a VPRN interface or an IES interface
- represents the public side of the tunnel-ISA

The delivery VPRN or IES service (the service connected to the Internet) must have at least one IP interface associated with a public tunnel SAP to receive and process GRE encapsulated packets.

The public tunnel SAP type has the format **tunnel-***id*.**private|public**:*tag* (where the *id* corresponds to the tunnel group). Figure 45: GRE for remote access to a VPRN service shows the example topology, where CE-2 in customer site A is connected to PE-1.

Figure 44: GRE for remote access to a VPRN service

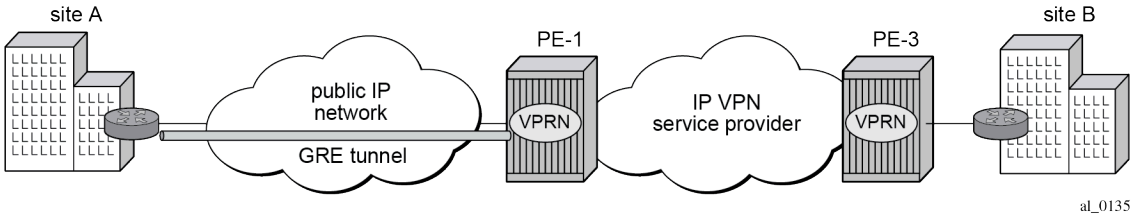

The IES service with public tunnel SAP is configured on PE-1 as follows:

```
[ex:/configure service ies "IES 2" interface "int-tunnel-public"]
A:admin@PE-1# sap ?

 [sap-id] (<port-and-encap> | <expression>)
 <port-and-encap>                              - null                  - (<port-id>|<lag-name>|
<aps-id>)
---snip---
                                                 tunnel-id            - tunnel-<1..64>.
(private|
                                                                         public):<tag>
---snip---


# on PE-1:
configure {
    service {
        ies "IES 2" {
            admin-state enable
            service-id 2
            customer "1"
            interface "int-PE-1-CE-2" {
                sap 1/1/2:2 {
                }
                ipv4 {
                    primary {
                        address 192.168.12.1
                        prefix-length 24
                    }
                }
            }
            interface "int-tunnel-public" {
                # tos-marking-state untrusted         # default
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 192.168.1.2
                        prefix-length 30
                    }
                }
            }
```

PE-1 has address 192.168.1.2/30 assigned to the interface "int-tunnel-public" in IES 2. In a similar way, CE-2 has address 192.168.2.2/30 assigned to the interface "int-tunnel-public" in IES 2.

To reach 192.168.2.0/30 on CE-2, the following static route is configured on PE-1:

```
# on PE-1:
configure {
```

```
    router "Base" {
        static-routes {
            route 192.168.2.0/30 route-type unicast {
                next-hop "192.168.12.2" {
                    admin-state enable
                }
            }
```

In a similar way, a static route is configured on CE-2 to reach 192.168.1.0/30 on PE-1.

Mask /32 is not supported on the public tunnel. When address 192.168.1.2/32 is configured on the interface "int-tunnel-public", the public tunnel cannot be created, as follows:

```
*[ex:/configure service ies "IES 2" interface "int-tunnel-public" ipv4 primary]
A:admin@PE-1# prefix-length 32

*[ex:/configure service ies "IES 2" interface "int-tunnel-public" ipv4 primary]
A:admin@PE-1# commit
MINOR: COMMON #238: configure service ies "IES 2" interface "int-tunnel-public" ipv4
primary prefix-length - Configuration change failed validation
- /32 address can be assigned only to the system or loopback interfaces.
- configure service ies "IES 2" interface "int-tunnel-public" sap tunnel-1.public:1
```

Therefore, the address configured on the interface has mask /30 instead of /32, as shown earlier.

The tunnel private interface:

- can be an IES or VPRN interface
- represents the private side of the tunnel ISA

The private tunnel SAP has the format **tunnel-***id*.**private|public**:*tag* (where the *id* corresponds to the tunnel group) as shown in the following example where an unprotected GRE tunnel is configured in the SAP context.

```
[ex:/configure service vprn "VPRN 1" interface "int-gre-tunnel"]
A:admin@PE-1# sap ?

 [sap-id] (<port-and-encap> | <expression>)
 <port-and-encap>                              - null                  - (<port-id>|<lag-name>|
<aps-id>)
---snip---
                                                 tunnel-id            - tunnel-<1..64>.
(private|
                                                                         public):<tag>
---snip---


# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            ---snip---

            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
```

```
                    }
                }
            }
            sap tunnel-1.private:1 {
                ip-tunnel "gre-tunnel-1" {
                    gre-header {
                        admin-state enable
                    }
                    dest-ip 10.0.0.2 { }
                     ---snip---
```

It is not mandatory to have the same tag (internal dot1q) in private and public GRE tunnels.

```
    sap tunnel-1.private:1 <=> sap tunnel-1.public:2
```

Unprotected GRE tunnel configuration

To associate an unprotected GRE tunnel with a private tunnel SAP, the **ip-tunnel** command is configured in the SAP context.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
        ---snip---
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        ---snip---
                       gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                        ---snip---
```

The **dest-ip** keyword followed by the private IP address of the remote tunnel endpoint is mandatory.

If this remote IP address is not within the subnet of the local private endpoint, then the tunnel does not come up.

The following parameters are configured in the **ip-tunnel** context:

- The local IP address of the GRE tunnel. This is the source IPv4 address of GRE encapsulated packets sent by the delivery service. It must be an address in the subnet of the associated public tunnel SAP interface.
- The remote IP address. If this address is reachable in the delivery service (there is a route), then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.

- The backup remote IP address. If the remote IP address of the tunnel is not reachable, then the backup remote IP address is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The delivery service. This is the identifier or name of the IES or VPRN service where GRE encapsulated packets are injected and terminated. The delivery service can be the same service where the private tunnel SAP interface resides.
- The DSCP marking in the outer IP header of GRE encapsulated packets. If this is not configured, then the default copies the DSCP from the inner IP header to the outer IP header.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        dscp af22
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            ---snip---
```

- A private tunnel SAP can have only one IP/GRE tunnel per SAP.

```
[ex:/configure service vprn "VPRN 1" interface "int-gre-tunnel" sap tunnel-1.private:1]
 A:admin@PE-1# ip-tunnel "gre-tunnel-2" {

*[ex:/configure service vprn "VPRN 1" interface "int-gre-tunnel" sap tunnel-1.private:1
                                                                      ip-tunnel "gre-
tunnel-2"]
 A:admin@PE-1# commit
 MINOR: MGMT_CORE #232: configure service vprn "VPRN 1" interface "int-gre-tunnel"
```

```
sap tunnel-1.private:1 ip-tunnel "gre-tunnel-2"
- Reached maximum number of entries - maximum is 1 but has 2
```

IP/GRE tunneling via static route

A static route can reference the GRE tunnel directly (by next-hop IP address) or the GRE tunnel can be the resolved next-hop for an indirect static route (Figure 45: GRE for remote access to a VPRN service).

Figure 45: GRE for remote access to a VPRN service

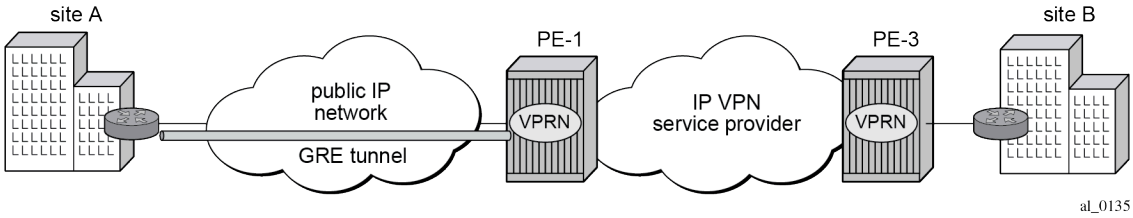

The details of both ends on the GRE tunnel, at site A and PE-1, are shown in Figure 46: IP/GRE tunneling via static route. The node at the lefthand side is CE-2 at site A.

Figure 46: IP/GRE tunneling via static route

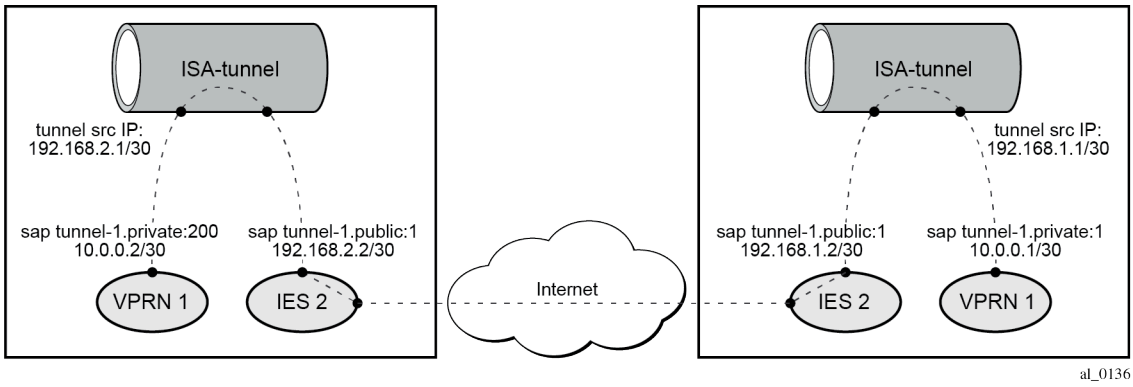

The following shows the configuration of VPRN 1 on PE-1.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            interface "int-gre-tunnel" {
```

```
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
                }
            }
            static-routes {
                route 172.16.2.1/32 route-type unicast {
                    next-hop "10.0.0.2" {
                        admin-state enable
                    }
                }
            }
        }
```

The configuration of the VPRN on CE-2 is similar.

The following command checks the static route status:

```
[/]
A:admin@PE-1# show router 1 static-route

===============================================================================
Static Route Table (Service: 1)  Family: IPv4
===============================================================================
Prefix                                    Nexthop Tag Met    Pref Type Act
   Next Hop                                Interface
-------------------------------------------------------------------------------
172.16.2.1/32                             0           1      5    NH   Y
   10.0.0.2                                int-gre-tunnel
-------------------------------------------------------------------------------
No. of Static Routes: 1
===============================================================================
```

IP/GRE tunneling via BGP peering

In this section, the configuration has BGP running inside the GRE tunnel.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64496
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            bgp {
                group "group-1" {
                    type internal
                    local-address 172.16.1.1
                    export {
                        policy "export-bgp-172.31"
                    }
                }
                neighbor "172.16.2.1" {
                    group "group-1"
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
```

```
            }
        }
        interface "loopback2" {
            loopback true
            ipv4 {
                primary {
                    address 172.31.1.1
                    prefix-length 24
                }
            }
        }
        static-routes {
            route 172.16.2.1/32 route-type unicast {
                next-hop "10.0.0.2" {
                    admin-state enable
                }
            }
        }
    }
```

It is mandatory to configure the autonomous system in the **vprn** context, otherwise the BGP session cannot be established.

The configuration of the VPRN on CE-2 is similar.

The following command on PE-1 shows the summary of the BGP sessions. The BGP session between peers 172.16.1.1 in VPRN 1 on PE-1 and 172.16.2.1 in VPRN 1 on CE-2 is established for address family IPv4.

```
[/]
A:admin@PE-1# show router 1 bgp summary all

===============================================================================
BGP Summary
===============================================================================
Legend : D - Dynamic Neighbor
===============================================================================
Neighbor
Description
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
172.16.2.1
1               64496       7    0 00h01m08s 1/1/1 (IPv4)
                            7    0
-------------------------------------------------------------------------------
```

In this example, PE-1 exports BGP route 172.31.1.0/24 and CE-2 exports BGP route 172.31.2.0/24. The route table for VPRN 1 on PE-1 includes the following BGP route:

```
[/]
A:admin@PE-1# show router 1 route-table protocol bgp

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.31.2.0/24                                 Remote  BGP       00h00m40s  170
       10.0.0.2                                                      1
-------------------------------------------------------------------------------
```

```
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IP/GRE tunneling via OSPFv2 peering

OSPF can be run on IES and VPRN IP interfaces associated with private IP/GRE tunnel SAPs.

All OSPF features are supported, including area 0 and non-area 0 support, virtual links, authentication, BFD, configurable protocol timers.

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
                }
            }
```

```
            ospf 0 {
                admin-state enable
                area 0.0.0.0 {
                    interface "int-gre-tunnel" {
                    }
                    interface "loopback1" {
                    }
                }
            }
```

The configuration on CE-2 is similar.

The following command shows the OSPF neighbors for VPRN 1:

```
[/]
A:admin@PE-1# show router 1 ospf neighbor

===============================================================================
Rtr vprn1 OSPFv2 Instance 0 Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
   Area-Id
-------------------------------------------------------------------------------
int-gre-tunnel                   192.0.2.2       Full       1    0       39
   0.0.0.0
-------------------------------------------------------------------------------
No. of Neighbors: 1
===============================================================================
```

The OSPF routes in the routing table of VPRN 1 are as follows:

```
[/]
A:admin@PE-1# show router service-name "VPRN 1" route-table protocol ospf

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
172.16.2.1/32                                 Remote  OSPF      00h01m02s  10
       10.0.0.2                                                    2
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

IP/GRE tunneling protection using IPSec tunnel mode

To provide protection against potential threats such as spoofing, the GRE packets can be encrypted and authenticated using IPSec.

GRE packets receive IPSec protection by forwarding them, after encapsulation by a tunnel-ISA, into an IPSec tunnel supported by another (or the same) tunnel ISA.

Note that when configuring GRE protection by an IPSec tunnel:

- A GRE tunnel and its protecting IPSec tunnel may belong to the same or different tunnel groups (the same tunnel group is assumed in the following example).
- A GRE tunnel and its protecting IPSec tunnel may be assigned to the same tunnel ISA (if they belong to the same tunnel group) or different tunnel ISAs.
- A single IPSec tunnel can protect one or more GRE tunnels in addition to other IP traffic that meets the IPSec security policy.
- The private IPSec tunnel SAP interface and public GRE tunnel SAP interface are always part of the same VPRN. The private GRE tunnel SAP interface can be part of this same VPRN or a different VPRN.

In the following example, the GRE tunnel and its protecting IPSec tunnel belong to the same tunnel group.

Figure 47: Example GRE over IPSec tunnel

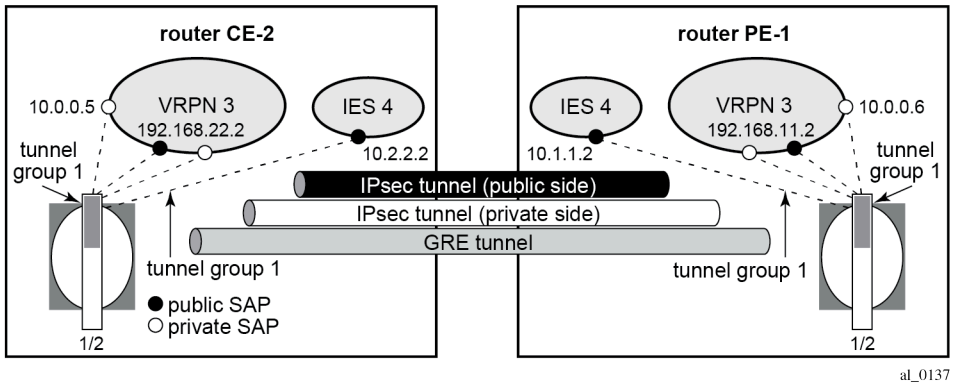

IPSec configuration

An **ike-policy** and **ipsec-transform** must be configured on PE-1 and CE-2, as follows:

```
# on PE-1, CE-2:
configure {
    ipsec {
        ike-policy 1 {
            ike-transform [1]
        }
        ike-transform 1 {
            dh-group group-5
        }
        ipsec-transform 1 {
            esp-encryption-algorithm aes-256
        }
```

The public and private side of the GRE tunnel and the private side of the IPSec tunnel are in the same VPRN, as shown in the following configuration example:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 3" {
            admin-state enable
            service-id 3
```

```
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 1 {
                        local-ip {
                            address 192.168.11.0/24
                        }
                        remote-ip {
                            address 192.168.22.0/24
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:3"
                    vrf-target {
                        community "target:64496:3"
                    }
                }
            }
            interface "int-private-gre-1" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.6 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:5 {
                    ip-tunnel "protected-gre-tunnel" {
                        admin-state enable
                        delivery-service "VPRN 3"
                        remote-ip-address 192.168.22.1
                        local-ip-address 192.168.11.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.5 { }
                    }
                }
            }
            interface "int-private-ipsec-1" {
                tunnel true
                sap tunnel-1.private:3 {
                    ipsec-tunnel "ipsec-tunnel-for-gre-tunnel" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.1.1.1
                            remote-ip-address 10.2.2.1
                            delivery-service "IES 4"
                        }
                        security-policy {
                            id 1
                        }
```

```
                    }
                }
            }
            interface "int-public-gre-1" {
                ipv4 {
                    primary {
                        address 192.168.11.2
                        prefix-length 24
                    }
                }
                sap tunnel-1.public:4 {
                }
            }
            static-routes {
                route 192.168.22.0/24 route-type unicast {
                    ipsec-tunnel "ipsec-tunnel-for-gre-tunnel" {
                        admin-state enable
                    }
                }
            }
        }
```

The following displays a configuration example of the public side of the IPSec tunnel:

```
# on PE-1:
configure {
    service {
        ies "IES 4" {
            admin-state enable
            service-id 4
            customer "1"
            interface "int2-PE-1-CE-2" {
                sap 1/1/2:4 {
                }
                ipv4 {
                    primary {
                        address 192.168.112.1
                        prefix-length 30
                    }
                }
            }
            interface "public-ipsec-1" {
                sap tunnel-1.public:3 {
                }
                ipv4 {
                    primary {
                        address 10.1.1.2
                        prefix-length 24
                    }
                }
            }
        }
```

The following static route is configured in the base router on PE-1:

```
# on PE-1:
configure {
    router "Base" {
        static-routes {
            route 10.2.2.0/24 route-type unicast {
                next-hop "192.168.112.2" {
                    admin-state enable
                }
```

```
            }
```

The configuration is similar on CE-2.

The following command shows that the tunnel "protected-gre-tunnel" with SAP tunnel-1.private:5 is up:

```
[/]
A:admin@PE-1# show ip tunnel

===============================================================================
IP Tunnels
===============================================================================
TunnelName                       SapId                          SvcId      Admn
 Local Address                                                  DlvrySvcId Oper
  OperRemoteAddress
-------------------------------------------------------------------------------
gre-tunnel-1                     tunnel-1.private:1             1          Up
 192.168.1.1                                                    IES 2      Up
  192.168.2.1
protected-gre-tunnel             tunnel-1.private:5             3          Up
 192.168.11.1                                                   VPRN 3     Up
  192.168.22.1
-------------------------------------------------------------------------------
IP Tunnels: 2
===============================================================================
```

The following command shows the IP/GRE tunnel information for this IPSec tunnel:

```
[/]
A:admin@PE-1# show ip tunnel "protected-gre-tunnel"

===============================================================================
IP Tunnel Configuration Detail
===============================================================================
Service Id       : 3                    Sap Id           : tunnel-1.private:5
Tunnel Name      : protected-gre-tunnel
Description      : None
GRE Header       : Yes
Delivery Service : VPRN 3
GRE Keys Set     : False
GRE Send Key     : N/A                  GRE Receive Key  : N/A
Admin State      : Up                   Oper State       : Up
Source Address   : 192.168.11.1
Remote Address   : 192.168.22.1
Backup Address   : (Not Specified)
Oper Remote Addr : 192.168.22.1
DSCP             : None
Reassembly       : inherit
Clear DF Bit     : false                IP MTU           : max
Encap IP MTU     : max
Pkt Too Big      : true
Pkt Too Big Num  : 100                  Pkt Too Big Intvl: 10 secs
Frag Required    : true
Frag Req Count   : 100                  Frag Req Interval: 10 secs
Propagate IPv6 P*: true
Propagate IPv4 P*: true
Oper Flags       : None
Transport Profile: (Not Specified)
Last Oper Changed: 06/25/2025 07:44:20
Host ISA         : 1/2
TCP MSS Adjust
    Public       : Disabled
    Private      : Disabled
```

```
-------------------------------------------------------------------------------
Target Address Table
-------------------------------------------------------------------------------
Destination IP                          IP Resolved Status
-------------------------------------------------------------------------------
10.0.0.5                                Yes
-------------------------------------------------------------------------------

===============================================================================
IP Tunnel Statistics: protected-gre-tunnel
===============================================================================
Errors Rx        : 0                    Errors Tx        : 0
Pkts Rx          : 0                    Pkts Tx          : 0
Bytes Rx         : 0                    Bytes Tx         : 0
Key Ignored Rx   : 0                    Too Big Tx       : 0
Seq Ignored Rx   : 0
Vers Unsup. Rx   : 0
Invalid Chksum Rx: 0
Key Mismatch Rx  : 0
===============================================================================

===============================================================================
Fragmentation Statistics
===============================================================================
Encapsulation Overhead                 : 24
Temporary Private MTU                  : max
Pre-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
Post-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

By default, the IPSec tunnel is down if it is not used by any traffic, as follows:

```
[/]
A:admin@PE-1# show ipsec tunnel

===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                      LocalAddress      SvcId        Admn   Keying
  SapId                           RemoteAddress     DlvrySvcId   Oper   Sec
                                                                        Plcy
-------------------------------------------------------------------------------
ipsec-tunnel-for-gre-tunnel     10.1.1.1          3            Up     Dynamic
  tunnel-1.private:3              10.2.2.1          IES 4        Down   1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

After it is used by traffic, the status is changed to be up.

```
[/]
A:admin@PE-1# ping 10.0.0.5 router-instance "VPRN 3" interval 0.1 output-format summary
PING 10.0.0.5 56 data bytes
!!!!!
---- 10.0.0.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
```

```
round-trip min = 4.58ms, avg = 4.94ms, max = 5.32ms, stddev = 0.245ms
```

The IPSec tunnel is now operationally up, as follows:

```
[/]
A:admin@PE-1# show ipsec tunnel

===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                      LocalAddress      SvcId        Admn   Keying
  SapId                           RemoteAddress     DlvrySvcId   Oper   Sec
                                                                        Plcy
-------------------------------------------------------------------------------
ipsec-tunnel-for-gre-tunnel     10.1.1.1          3            Up     Dynamic
  tunnel-1.private:3              10.2.2.1          IES 4        Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

BFD support on private tunnel interfaces

BFD is supported on IP interfaces associated with private IP/GRE tunnel SAPs. The BFD state of the interface can be used by static routes, OSPFv2, or BGP configured on the interface. It is not used to trigger a switchover to the backup remote IP address of the GRE tunnel.

The following displays a static route example:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    bfd {
                        admin-state enable      # Configure BFD on GRE tunnel interface
                    }
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
```

```
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
                }
            }
            static-routes {
                route 172.16.2.1/32 route-type unicast {
                    next-hop "10.0.0.2" {
                        admin-state enable
                        bfd-liveness        # Enable BFD for static route
                    }
                }
            }
        }
```

The following command shows that the BFD session on interface "int-gre-tunnel" is up for protocol static:

```
[/]
A:admin@PE-1# show router 1 bfd session

===============================================================================
Legend:
  Session Id = Interface Name | LSP Name | Prefix | RSVP Sess Name | Service Id
  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Session Id                                      State      Tx Pkts    Rx Pkts
  Rem Addr/Info/SdpId:VcId                    Multipl     Tx Intvl   Rx Intvl
  Protocols                                      Type     LAG Port     LAG ID
  Loc Addr                                                          LAG name
-------------------------------------------------------------------------------
int-gre-tunnel                                    Up          N/A        N/A
  10.0.0.2                                         3         1000       1000
  static                                       cpm-np          N/A        N/A
  10.0.0.1
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

When no static routes are configured and OSPF is configured instead, the configuration of VPRN 1 on PE-1 is as follows:

```
# on PE-1:
configure {
    service  {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
```

```
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    bfd {
                        admin-state enable      # Configure BFD on GRE tunnel interface
                    }
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
                }
            }
            ospf 0 {
                admin-state enable
                area 0.0.0.0 {
                    interface "int-gre-tunnel" {
                        bfd-liveness {          # Enable BFD on OSPF interface "int-gre-tunnel"
                        }
                    }
                    interface "loopback1" {
                    }
                }
            }
```

The following shows that the BFD session is up for protocol OSPF on interface "int-gre-tunnel":

```
[/]
A:admin@PE-1# show router 1 bfd session

===============================================================================
Legend:
```

```
  Session Id = Interface Name | LSP Name | Prefix | RSVP Sess Name | Service Id
  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Session Id                                        State      Tx Pkts    Rx Pkts
  Rem Addr/Info/SdpId:VcId                      Multipl     Tx Intvl   Rx Intvl
  Protocols                                        Type     LAG Port     LAG ID
  Loc Addr                                                             LAG name
-------------------------------------------------------------------------------
int-gre-tunnel                                      Up          N/A        N/A
  10.0.0.2                                           3         1000       1000
  ospf2                                          cpm-np          N/A        N/A
  10.0.0.1
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

When BGP is configured instead of OSPF, the configuration of VPRN 1 on PE-1 is as follows:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            admin-state enable
            service-id 1
            customer "1"
            autonomous-system 64496
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:1"
                    vrf-target {
                        community "target:64496:1"
                    }
                }
            }
            bgp {
                group "group-1" {
                    type internal
                    local-address 172.16.1.1
                }
                neighbor "172.16.2.1" {
                    bfd-liveness true               # Enable BFD on BGP
                    group "group-1"
                }
            }
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
```

```
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
            interface "loopback1" {
                loopback true
                ipv4 {
                    bfd {
                        admin-state enable          # Configure BFD on loopback interface for
 BGP
                    }
                    primary {
                        address 172.16.1.1
                        prefix-length 32
                    }
                }
            }
            static-routes {
                route 172.16.2.1/32 route-type unicast {
                    next-hop "10.0.0.2" {
                        admin-state enable
                    }
                }
            }
```

The following command shows that the BFD session is up for protocol BGP on interface "loopback1":

```
*A:PE-1# show router 1 bfd session

===============================================================================
Legend:
  Session Id = Interface Name | LSP Name | Prefix | RSVP Sess Name | Service Id
  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
Session Id                                        State      Tx Pkts    Rx Pkts
  Rem Addr/Info/SdpId:VcId                      Multipl     Tx Intvl   Rx Intvl
  Protocols                                        Type     LAG Port     LAG ID
  Loc Addr                                                             LAG name
-------------------------------------------------------------------------------
loopback1                                           Up          N/A        N/A
  172.16.2.1                                         3         1000       1000
  bgp                                            cpm-np          N/A        N/A
  172.16.1.1
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
```

IP/GRE termination – Advanced topics

DSCP value of outer delivery header

- Default behavior — The DSCP value from the payload header is copied into the outer GRE header. This is a one to one copy and no QoS classifications are required. It is performed when no DSCP value is configured under the private GRE tunnel.

- Non default behavior — DSCP is configured under the private SAP (following example using DSCP af41).

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            interface "int-gre-tunnel" {
                tunnel true
                ipv4 {
                    addresses {
                        address 10.0.0.1 {
                            prefix-length 30
                        }
                    }
                }
                sap tunnel-1.private:1 {
                    ip-tunnel "gre-tunnel-1" {
                        admin-state enable
                        delivery-service "IES 2"
                        dscp af41                           # Configure DSCP value
                        remote-ip-address 192.168.2.1
                        local-ip-address 192.168.1.1
                        gre-header {
                            admin-state enable
                        }
                        dest-ip 10.0.0.2 { }
                    }
                }
            }
---snip---
```

The log filter output shows TOS=88 (DSCP=af41) in the public network.

```
[/]
A:admin@PE-1# show filter log 102

===============================================================================
Filter Log
===============================================================================
Admin state : Enabled
Description : (Not Specified)
Destination : Memory
Wrap        : Enabled
-------------------------------------------------------------------------------
Maximum entries configured : 1000
Number of entries logged   : 5
-------------------------------------------------------------------------------
2025/06/25 07:57:46  Ip Filter: 2:10  Desc:
SAP: tunnel-1.private:1  Direction: Egress
Src MAC: 00-03-fe-00-02-c9  Dst MAC: 00-00-00-07-a0-bd  EtherType: 0800
Src IP: 10.0.0.1  Dst IP: 10.0.0.2  Flags: 0  TOS: 88  TTL: 64 Len: 84
Protocol: ICMP  Type: Echo Request  Code: 0
---snip---
```

IP MTU

It is possible to configure the IP MTU of a private tunnel SAP interface. This sets the maximum IP packet size payload (including IP header) that can be sent into the tunnel (it applies to the packet size before the tunnel encapsulation is added).

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            interface "int-gre-tunnel"
                ip-mtu 1476
            ---snip---
```

When an IPv4 packet needs to be forwarded to the tunnel and is larger than IP MTU bytes:

- If the DF bit is clear, the payload packet is IP fragmented to the MTU size before the tunnel encapsulation.
- If the DF bit is set, the payload packet is discarded.

The IP MTU range supported is from 512 to 9000 bytes.

The following command shows the configured IP MTU and the operational IP MTU for the GRE tunnel:

```
[/]
A:admin@PE-1# show router 1 interface "int-gre-tunnel" detail | match MTU
IP MTU           : 1476
IP Oper MTU      : 1476
```

Statistics and accounting

Collect-stats can be configured under public and private SAPs.

For public SAPs:

```
# on PE-1:
configure {
    service {
        ies "IES 2" {
            interface "int-tunnel-public" {
                sap tunnel-1.public:1 {
                    collect-stats true
```

For private SAPs:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            interface "int-gre-tunnel" {
                sap tunnel-1.private:1 {
                    collect-stats true
```

Filtering, policing, and QoS

An IP filter and QoS policy can be applied to the ingress and egress traffic of the private and public SAPs.

Public SAPs:

```
# on PE-1:
configure {
    service {
        ies "IES 2" {
            interface "int-tunnel-public"
                sap tunnel-1.public:1
                    ingress {
                        qos {
                            sap-ingress {
                                policy-name "sap-ingr-10"
                            }
                        }
                        filter {
                            ip "ip-1"
                        }
                    }
                    egress
                        qos {
                            sap-egress {
                                policy-name "sap-egr-20"
                            }
                        }
                        filter {
                            ip "ip-2"
                        }
                    }
            ---snip---
```

Private SAPs:

```
# on PE-1:
configure {
    service {
        vprn "VPRN 1" {
            interface "int-gre-tunnel" {
                sap tunnel-1.private:1 {
                    ingress {
                        qos {
                            sap-ingress {
                                policy-name "sap-ingr-10"
                            }
                        }
                        filter {
                            ip "ip-1"
                        }
                    }
                    egress {
                        qos {
                            sap-egress {
                                policy-name "sap-egr-20"
                            }
                        }
                        filter {
                            ip "ip-2"
                        }
                    }
```

```
                }
            ---snip---
```

Mirroring

The public and private SAPs can be mirrored. The following is in classic CLI:

```
# on PE-1:
debug
    mirror-source 99
        sap tunnel-1.private:3 egress ingress
        sap tunnel-1.public:1 egress ingress
        no shutdown
    exit
exit
```

Conclusion

This chapter provides configuration and show commands for IP/GRE termination.

Multi-Chassis IPSec Redundancy

This chapter provides information about multi-chassis IPSec redundancy configurations.

Topics in this chapter include:

Applicability

This initial version of this chapter was based on SR OS Release 10.0.R8, but the MD-CLI in the current edition corresponds to SR OS Release 22.10.R2.

Overview

Multi-Chassis IPSec redundancy (MC-IPSec) is a stateful inter-chassis IPSec failover mechanism. IPSec tunnel states are synchronized between the primary and standby chassis. A tunnel group failure on the primary chassis or a primary chassis failure could trigger MC-IPSec failover to the standby chassis.

The following are some highlights of this feature:

- Internet Key Exchange version 2 (IKEv2) only
- Multi-active tunnel group only
- The granularity of failover is tunnel group, which means a specific tunnel group could failover to the standby chassis independent of other tunnel groups on the primary chassis
- Both static and dynamic LAN-to-LAN tunnels are supported

This feature has the following building blocks:

- Primary chassis election: MC-IPSec mastership protocol (MIMP) runs between the chassis to elect a primary chassis with independent MIMP runs for each tunnel group
- Synchronization: multi-chassis synchronization (MCS) synchronizes the IPSec states between chassis
- Routing:
  - MC-IPSec-aware routing attracts traffic to the primary chassis
  - Shunting support
  - MC-IPSec-aware virtual router redundancy protocol (VRRP)

The figure Figure 48: MC-IPSec architecture shows two redundant IPSec chassis in the middle: a primary chassis and a standby chassis.

Figure 48: MC-IPSec architecture

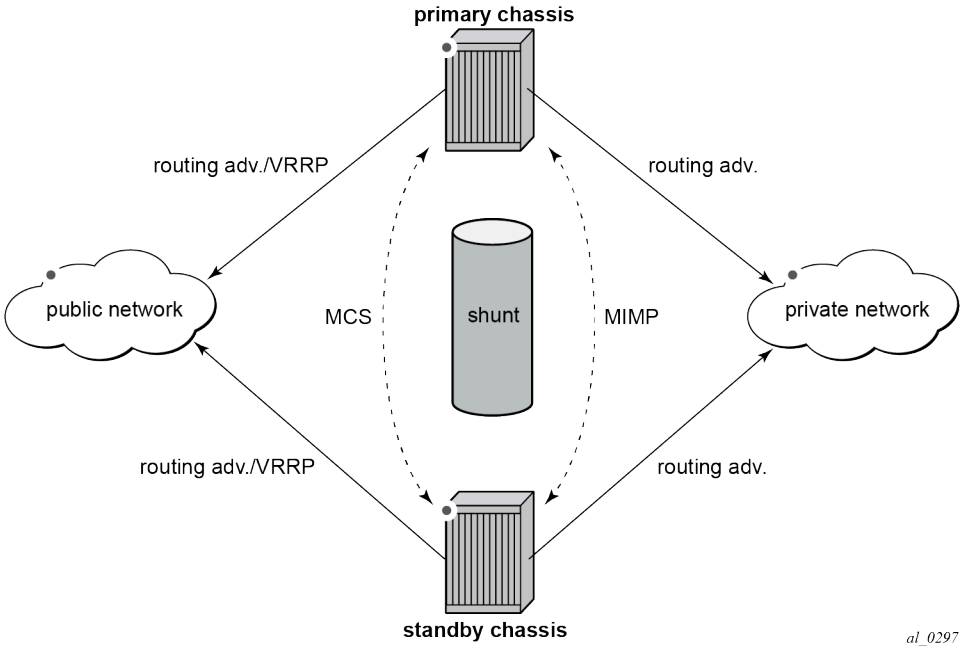

The fundamentals of MC-IPSec are:

- Only the primary chassis processes encapsulating security payload (ESP) and IKE traffic. If the standby chassis receives traffic, it shunts it to the primary chassis, if possible. The traffic is discarded if the standby chassis fails to shunt the traffic.
- The same local gateway address must be provisioned on both chassis.
- MC-IPSec does not synchronize configurations.
- MC-IPSec-aware routing attracts traffic to the primary chassis for both public and private services, which is achieved by exporting the corresponding IPSec routes to the routing protocol using a route policy and setting a different routing metric according to the MC-IPSec state.
- In case of a Layer 2 public network, MC-IPSec-aware VRRP can be used to trigger VRRP switchover upon MC-IPSec switchover.
- MCS synchronizes IPSec states between chassis so that existing IPSec tunnels do not need to be re-established upon switchover.
- MIMP elects mastership between two chassis, and it can also detect chassis failure and tunnel group failure; a central BFD session can be associated with MIMP to achieve fast chassis failure detection.

Configuration

The example topology is shown in the figure Figure 49: Example topology.

Figure 49: Example topology

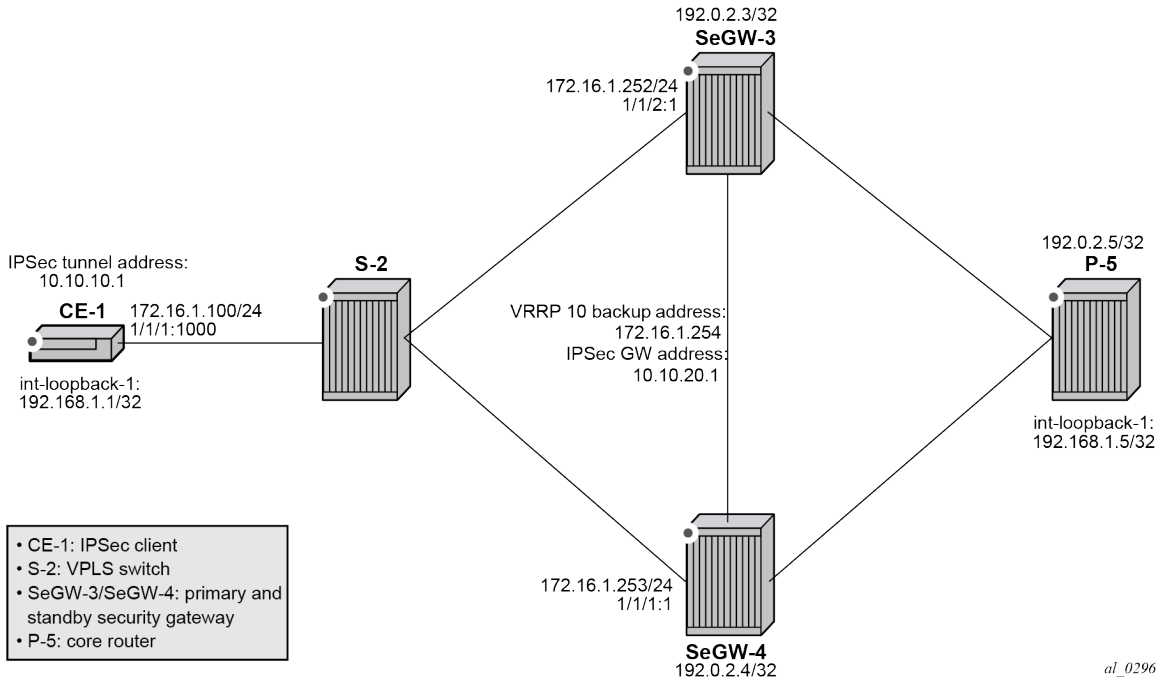

The example setup includes:

- an IPSec tunnel initiated by CE-1 and terminated on the primary chassis of the two SeGWs.
- a public IES service "IES-1" and a private VPRN service "VPRN-2" configured on CE-1, SeGW-3, and SeGW-4.
- VPRN "VPRN-2" (also) configured on P-5.
- a static LAN-to-LAN tunnel with pre-shared key.
- a local VPLS service "VPLS-3" on S-2 to simulate a Layer 2 switch.
- VRRP 10 between SeGW-3 and SeGW-4 to provide a backup address 192.168.1.254, which is the default next hop for CE-1.
- VRRP policy 1 bound to VRRP 10 on the primary chassis SeGW-3 to change the in-use priority upon MC-IPSec switchover.
- OSPF as IGP running in the base routing instance between SeGW-3, SeGW-4, and P-5.
- MP-BGP running between SeGW-3, SeGW-4, and P-5 for the VPN-IPv4 address family.

A ping in VPRN "VPRN-2" between loopback interface address 192.168.1.1 on CE-1 and 192.168.1.5 on P-5 is used to verify the connectivity over the IPSec tunnel.

The MC-IPSec configuration commands are shown below.

```
configure
    redundancy
        multi-chassis
            peer <ip-address>
                sync
                    ipsec
                    tunnel-group <1..64>
```

```
                        sync-tag <string>

                mc-ipsec
                    bfd-liveness <boolean>
                    discovery-interval
                        interval-secs <1..1800>
                        boot <1..1800>
                    hold-on-neighbor-failure <2..25>
                    keep-alive-interval <5..500>     # deciseconds
                    tunnel-group <1..64>
                        admin-state <boolean>
                        peer-group <1..64>
                        priority <0..255>
```

```
configure
    policy-options
        policy-statement <string>
            entry <1..4294967295>
                from
                    state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer
                    protocol
                        name ipsec
```

```
configure
    service
        ies <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>


configure
    service
        vprn <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>
```

```
configure
    isa
        tunnel-group <1..64>
            ipsec-responder-only <boolean>
```

```
configure
    vrrp
        policy <1..9999>
            priority-event
                mc-ipsec-non-forwarding <tunnel-grp-id>
                    hold-clear <1..86400 seconds>
                    hold-set <1..86400 seconds>
                    priority
                        priority-level <1..254>
                        event-type (delta|explicit)
```

The parameters are the following:

- in the **configure redundancy multi-chassis** context:

- **peer** <*ip-address*> — This command creates or enters a multi-chassis peer. The peer address is by default the system address. This can be changed on the peer using the **configure redundancy multi-chassis peer source-address** command.
  - **sync** — This command enters the sync configuration context.
    - **ipsec** <*boolean*> — This command enables MCS to synchronize IPSec states.
    - **tunnel-group** <*tunnel-group-id*> **sync-tag** <*tag-name*> — This command enables MCS to synchronize the IPSec states of the specified tunnel group. The **sync-tag** parameter is used to match the tunnel group of the peer. The tunnel group states with the same **sync-tag** on both chassis will be synchronized.
  - **mc-ipsec** — This command enters the multi-chassis IPSec configuration context.
    - **bfd-liveness** <*boolean*> — The command **bfd-liveness** *true* enables tracking a central BFD session; if the BFD session goes down, then the system considers the peer as down and changes the MC-IPSec status of the configured tunnel group accordingly.

      The BFD session uses the source address of MCS as its source address and the MCS peer address as the destination address. Other BFD parameters are configured in the **bfd** context on the interface that the MCS source address resides on.

      The configuration of BFD is optional for MC-IPSec.
    - **discovery-interval interval-secs** <*interval-1*> **[boot** <*interval-2*>**]** — This command specifies the time interval that the tunnel group stays in **discovery** state. Interval 1 is used as discovery interval when a new tunnel group is added to multi-chassis redundancy (**mp-ipsec**); interval 2 is used as discovery interval after system boot-up. Interval 2 is optional, and when it is not specified, the value for interval 1 is used. Both intervals have a default value of 300 seconds.
    - **hold-on-neighbor-failure** <*2..25*> — This command specifies the number of keep-alive failures before considering the peer to be down. The default value is 3.
    - **keep-alive-interval** <*5..500*> — This command specifies the time interval of the mastership election protocol keep-alive packets in deciseconds. The default value is 10 deciseconds (1 s).
    - **tunnel-group** <*tunnel-group-id*> — This command enables multi-chassis redundancy for the specified tunnel group, or enters an already configured tunnel group context. The configured tunnel groups can failover independently.

      — **peer-group** <*tunnel-group-id*> — This command specifies the corresponding tunnel group ID on the peer node. The peer tunnel group ID is not necessarily equal to local tunnel group ID.

      — **priority** <*priority*> — This command specifies the local priority of the tunnel group, this is used to elect a primary chassis, where the higher number prevails. If the priorities are the same, then the peer which has more active ISAs wins; if the priority and the number of active ISAs are same, then the peer with higher IP address wins. The range is from 0 to 255 and the default value is 100.

- in a **from** statement of a route policy entry:
  - **state ipsec-master-with-peer | ipsec-non-master | ipsec-master-without-peer** — These commands specify the MC-IPSec state in a **from** statement of a route policy entry:
    - **ipsec-master-with-peer**: The tunnel group is the primary chassis with a peer reachable.
    - **ipsec-master-without-peer**: The tunnel group is the primary chassis with peer unreachable.

- **ipsec-non-master**: The tunnel group is not the primary chassis.

– **protocol name ipsec** — This command specifies IPSec as protocol in a **from** statement of a route policy entry. **protocol name ipsec** refers to the /32 local gateway routes (of both static and dynamic tunnels) and reverse route of dynamic tunnel.

- on a public or private IPSec interface in an IES or VPRN service:

– **static-tunnel-redundant-nexthop** <*ip-address*> and **dynamic-tunnel-redundant-nexthop** <*ip-address*> — These commands specify the redundant next hop address on a public or private IPSec interface (with public or private tunnel SAP) for a static and dynamic IPSec tunnel respectively. The specified next hop address is used by the standby chassis to shunt traffic to the primary chassis in case it receives any traffic. The next hop address is resolved in the routing table of the corresponding service.

Note:

- Shunting is supported over:
  - directly connected SAPs
  - spoke SDP terminated IP interfaces
- Shunting over auto-bind tunnel is not supported.
- Shunting does not work if the tunnel group is down.

- in the **isa tunnel-group <id>** context:

– **ipsec-responder-only** <*boolean*> — With the command **ipsec-responder-only** *true* , the system only acts as IKE responder except for the automatic CHILD_SA rekey upon MC-IPSec switchover.

This command is required for MC-IPSec support of static LAN-to-LAN tunnels.

- in the **vrrp policy <id> priority-event** context:

– **mc-ipsec-non-forwarding** <*tunnel-grp-id*> — This command creates a VRRP policy priority event: *mc-ipsec-non-forwarding*, which is triggered whenever the specified tunnel group enters the non-forwarding state.

- **hold-clear** <*seconds*> — This command configures the hold time before clearing the event. The range is from 0 to 86400 seconds and the default value is 0 s.
- **hold-set** <*seconds*> — This command configures the hold time before setting the event. The range is from 0 to 86400 seconds and the default value is 0 s.
- **priority** <*priority-level*> **explicit** — This command sets the VRRP in-use priority to the configured value upon the event. The range is from 0 to 254 and the default value is 0.

The initial configuration must include the following:

- The system time of SeGW-3 and SeGW-4 must be the same for the feature to work. Nokia recommends to use a time synchronization protocol such as NTP or SNTP.
- SeGW-3 and SeGW-4 must be IP reachable in the base routing instance because both MCS and MIMP run in the base routing instance.

Configuration of MC-IPSec

In this section, the following steps are described:

- configure CE-1
- configure S-2
- configure P-5
- configure IPSec tunnel on SeGW-3
- enable MC-IPSec for tunnel group on SeGW-3
- configure MC-IPSec-aware routing on SeGW-3
- configure MC-IPSec-aware VRRP on SeGW-3
- configure SeGW-4

Configure CE-1

On CE-1, the following is configured:

- a public IES service "IES-1" and a private VPRN service "VPRN-2".
- a static default route pointing to the VRRP backup address 172.16.1.254.
- a static IPSec tunnel "tunnel-1" with local address 10.10.10.1 and remote address 10.10.20.1.
- a loopback interface in VPRN "VPRN-2" with address 192.168.1.1/32 to be used as source address for the ping command to verify the connectivity between CE-1 and P-5 over the IPSec tunnel.

The following base router configuration on CE-1 includes a static route with next hop 172.16.1.254, which is the VRRP backup address.

```
# on CE-1:
configure {
    router "Base" {
        interface "int-CE-1-S-2" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 172.16.1.100
                    prefix-length 24
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 172.31.2.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 0.0.0.0/0 route-type unicast {
                next-hop "172.16.1.254" {       # VRRP backup address
                    admin-state enable
                }
            }
        }
```

IPSec is configured as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ike-transform [1]
            ike-version-2 {
            }
            dpd {    # dead peer detection (on peer side; not on MC-IPSec chassis)
            }
        }
        ike-transform 1 {
        }
        ipsec-transform 1 {
        }
```

Tunnel group 1 is configured as follows:

```
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            primary 1/2
        }
```

The public IES service is configured as follows:

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.10.254
                        prefix-length 24
                    }
                }
            }
        }
```

The private VPRN service on CE-1 is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.1/32
                        }
                        remote-ip {
```

```
                        address 192.168.1.5/32
                    }
                }
            }
        }
        interface "int-IPsec-private-1" {
            tunnel true
            sap tunnel-1.private:1 {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                    key-exchange {
                        dynamic {
                            ike-policy 1
                            ipsec-transform [1]
                            pre-shared-key "pass"
                        }
                    }
                    tunnel-endpoint {
                        local-gateway-address 10.10.10.1
                        remote-ip-address 10.10.20.1
                        delivery-service "IES-1"
                    }
                    security-policy {
                        id 1
                    }
                }
            }
        }
        interface "int-loopback-1" {
            loopback true
            ipv4 {
                primary {
                    address 192.168.1.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.1.5/32 route-type unicast {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                }
            }
        }
```

Configure S-2

On S-2, a local VPLS service 3 simulates a Layer 2 switch between CE-1, SeGW-3, and SeGW-4:

```
# on S-2:
configure {
    service {
        vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            sap 1/1/c1/1:1 {
                description "to SAP in IES 1 on SeGW-3"
            }
            sap 1/1/c1/2:1000 {
```

```
                description "to router interface in CE-1"
            }
            sap 1/1/c1/3:1 {
                description "to SAP in IES 1 on SeGW-4"
            }
        }
```

Configure P-5

P-5 simulates the core network router, connecting to SeGW-3 and SeGW-4. The configuration on P-5 includes the following:

- a loopback interface with address 192.168.1.5/32 in VPRN "VPRN-2", which is the destination address of the ping traffic from CE-1.
- an MP-BGP session for the VPN-IPv4 address family between P-5, SeGW-3, and SeGW-4.
- GRE spoke SDPs to connect to SeGW-3 and SeGW-4.

On P-5, the following router interfaces are configured in the base router. OSPF is used as IGP.

```
# on P-5:
configure {
    router "Base" {
        interface "int-P-5-SeGW-3" {
            port 1/1/c1/2:1000
            ipv4 {
                primary {
                    address 192.168.35.2
                    prefix-length 30
                }
            }
        }
        interface "int-P-5-SeGW-4" {
            port 1/1/c1/1:1000
            ipv4 {
                primary {
                    address 192.168.45.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-P-5-SeGW-3" {
                }
                interface "int-P-5-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```

On P-5, the following GRE SDPs are configured toward SeGW-3 and SeGW-4:

```
configure {
    service {
        sdp 53 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 54 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
```

VPRN "VPRN-2" is configured on P-5, as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-loopback-1" {
                loopback true
                ipv4 {
                    primary {
                        address 192.168.1.5
                        prefix-length 32
                    }
                }
            }
            spoke-sdp 53:2 {
            }
            spoke-sdp 54:2 {
            }
        }
```

The BGP configuration on P-5 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
```

```
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
        }
```

Configure IPSec tunnel on SeGW-3

The configuration on SeGW-3 is described in four consecutive sections. In this first section, the following is configured:

- the tunnel group, which must be in multi-active mode before MC-IPSec can be enabled.
- an interface "int-Redundant-1", which is a spoke-SDP terminated interface used for shunting.
- GRE SDP 34 toward SeGW-4 and GRE SDP 35 toward P-5.
- IPSec tunnel "tunnel-1" is the tunnel to CE-1; both SeGW-3 and SeGW-4 use the same local gateway address: 10.10.20.1.

The following configures tunnel group 1 on SeGW-3:

```
# on SeGW-3
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

On SeGW-3, the following router interfaces are configured in the base router. A static route is configured toward CE-1. OSPF is the IGP used between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-3-P-5" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 192.168.35.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-3-SeGW-4" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.1
                    prefix-length 30
                }
            }
```

```
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-3-P-5" {
                }
                interface "int-SeGW-3-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```

The IPSec settings are as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The GRE SDPs are configured as follows:

```
configure {
    service {
        sdp 34 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
        sdp 35 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
```

```
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows. In a later step, a VRRP policy will be configured and applied.

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.2
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                    }
                }
            }
```

The private VPRN service "VPRN-2" is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
```

```
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.2
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.1
                        prefix-length 30
                    }
                }
                spoke-sdp 34:20 {
                    ingress {
                        vc-label 2049
                    }
                    egress {
                        vc-label 2048
                    }
                }
            }
            spoke-sdp 34:2 {
                description "SDP to SeGW-4"
            }
            spoke-sdp 35:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
```

Enable MC-IPSec for tunnel group 1 on SeGW-3

In this section, the following steps are described:

- Create a multi-chassis peer using the system address of SeGW-4.
- Enable MCS for IPSec and tunnel group 1.
- Enable MC-IPSec for the tunnel group with a configured priority 200.
- Bind a central BFD session to MC-IPSec from the system interface.

Multi-chassis peer 192.0.2.4 is configured and MCS and MC-IPSec are enabled for tunnel group 1:

```
# on SeGW-3:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.4 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 200
                    }
                }
            }
```

BFD is enabled for MC-IPSec in the preceding configuration. BFD is configured on the system interface 192.0.2.3:

```
configure {
    router "Base" {
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
```

Configure MC-IPSec-aware routing on SeGW-3

In this step, a route policy is defined and applied to VPRN "VPRN-2".

Route policy "IPsec-to-MPBGP" exports static route 192.168.1.1/32 in VPRN "VPRN-2" to P-5. This policy sets the local preference of the prefix 192.168.1.1/32 according to the MC-IPSec state:

- for the **ipsec-master-with-peer** state: local preference 200
- for the **ipsec-non-master** state: local preference 100
- for the **ipsec-master-without-peer** state: local preference 200

The state **ipsec-master-without-peer** can be used to attract traffic to the designated primary chassis in case of "dual master" (meaning two chassis lose the MIMP connection in the base routing instance). In this example, SeGW-3 has local preference 200 and SeGW-4 has local preference 100 for **ipsec-master-without-peer**.

The route policy is configured as follows:

```
# on SeGW-3:
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
```

```
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-3 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
```

The route policy is applied as **vrf-export** in VPRN "VPRN-2":

```
configure {
    service {
        vprn "VPRN-2" {
            bgp-ipvpn {
                mpls {
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
```

Configure MC-IPSec-aware VRRP on SeGW-3

In this section, a VRRP policy is defined that uses the **mc-ipsec-non-forwarding** priority event to lower the in-use VRRP priority upon MC-IPSec switchover, which ensures VRRP and MC-IPSec have the same primary chassis. The VRRP instance needs to be in preempt mode.

This VRRP policy is only configured on the designated VRRP primary chassis SeGW-3, not on the standby chassis. The VRRP policy is applied to the interface "int-SeGW3-S-2" of IES "IES-1".

VRRP policy 1 is configured as follows:

```
# on SeGW-3:
configure {
    vrrp {
        policy 1 {
            priority-event {
                mc-ipsec-non-forwarding 1 {
                    priority {
                        priority-level 50
```

```
                        event-type explicit
                    }
                }
            }
        }
```

VRRP policy 1 is applied in VRRP instance 10 in the IES service:

```
configure {
    service {
        ies "IES-1" {
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                        policy 1
                    }
                }
            }
---snip---
```

Configure SeGW-4

The configuration on the standby chassis SeGW-4 is similar, but with different priorities and without the VRRP policy.

The tunnel group is configured in multi-active mode:

```
# on SeGW-4:
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

The MCS and MC-IPSec configuration is as follows:

```
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.3 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
```

```
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 150
                    }
                }
            }
        }
```

The base router configuration on SeGW-4 includes the following router interfaces and a static route to CE-1. OSPF is used as IGP between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-4-P-5" {
            port 1/1/2:1000
            ipv4 {
                primary {
                    address 192.168.45.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-4-SeGW-3" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.4
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-4-P-5" {
                }
                interface "int-SeGW-4-SeGW-3" {
                }
                interface "system" {
```

```
                }
            }
        }
```

The IPSec configuration is as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The following route policy is configured on SeGW-4, The local preference is lower for the **ipsec-master-without-peer** state.

```
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
```

```
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-4 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
        }
```

The following GRE SDPs are configured:

```
configure {
    service {
        sdp 43 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 45 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows:

```
configure {
```

```
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.1
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-4-S-2" {
                sap 1/1/1:1 {
                }
                ipv4 {
                    primary {
                        address 172.16.1.253
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        ping-reply true
                    }
                }
            }
        }
```

The private VPRN service is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
                }
```

```
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.1
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.2
                        prefix-length 30
                    }
                }
                spoke-sdp 43:20 {
                    ingress {
                        vc-label 2048
                    }
                    egress {
                        vc-label 2049
                    }
                }
            }
            spoke-sdp 43:2 {
                description "SDP to SeGW-3"
            }
            spoke-sdp 45:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
        }
```

Verification

The following will be verified in this section:

- the MC-IPSec status and VRRP status on SeGW-3 and SeGW-4

- the status of the IPSec tunnel on CE-1
- the status of the IPSec tunnel on the SeGWs

Verify the MC-IPSec status on SeGW-3 and SeGW-4

The following is verified:

- SeGW-3 is the primary chassis (**master**) and SeGW-4 is the standby for tunnel group 1 because SeGW-3 has the higher priority 200.
- SeGW-3 is the primary node for VRRP instance 10 and SeGW-4 is the backup.

SeGW-3 is the primary chassis in tunnel group 1 with priority 200:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is the standby chassis in tunnel group 1 with priority 150:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
```

```
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is backup for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Verify the IPSec tunnel on CE-1

The following is verified in this section:

- the connectivity between CE-1 and P-5
- the IPSec tunnel information

A ping command is launched from the loopback interface in VPRN "VPRN-2" on CE-1 to the loopback interface in VPRN "VPRN-2" on P-5:

```
[/]
A:admin@CE-1# ping 192.168.1.5 router-instance "VPRN-2"
PING 192.168.1.5 56 data bytes
64 bytes from 192.168.1.5: icmp_seq=1 ttl=63 time=2.44ms.
64 bytes from 192.168.1.5: icmp_seq=2 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=3 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=4 ttl=63 time=2.51ms.
64 bytes from 192.168.1.5: icmp_seq=5 ttl=63 time=2.50ms.
```

```
---- 192.168.1.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.38ms, avg = 2.44ms, max = 2.51ms, stddev = 0.053ms
```

The following command shows the IPSec tunnel information.

```
[/]
A:admin@CE-1# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.10.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.20.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

Verify the IPSec tunnel on the SeGWs

In this section, the following is verified:

- the MCS database is in-sync, so the tunnel status is up on both chassis.
- P-5 receives two VPN-IPv4 routes for prefix 192.168.1.1/32: the route from SeGW-3 has local preference 200; the route from SeGW-4 has local preference 100.

On both SeGWs, the IPSec tunnel with local address 10.10.20.1 and remote address 10.10.10.1 is up:

```
[/]
A:admin@SeGW-3# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.10.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

```
[/]
A:admin@SeGW-4# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
```

```
tunnel-1                          10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1                10.10.10.1        IES-1        Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

MCS is in sync on both SeGWs:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.4
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.3
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.3
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.4
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
```

```
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

The following command shows that P-5 received two VPN-IPv4 routes for prefix 192.168.1.1/32: one from SeGW-3 with local preference 200 and one from SeGW-4 with local preference 100:

```
[/]
A:admin@P-5# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.5        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  64496:2:192.168.1.1/32                            200         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*i    64496:2:192.168.1.1/32                            100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.0/30                           100         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*>i   64496:2:192.168.20.0/30                           100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.1/32                           100         0
      192.0.2.3                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.2/32                           100         0
      192.0.2.4                                         None        10
      No As-Path                                                    524286
-------------------------------------------------------------------------------
```

```
Routes : 6
===============================================================================
```

MC-IPSec failover scenarios

Two MC-IPSec failover scenarios are described in this section:

- MC-IPSec failover when MS-ISA is disabled
- MC-IPSec failover when the primary chassis SeGW-3 reboots

Failover when MS-ISA is disabled

Initially, MS-ISA is enabled, so SeGW-3 is the primary chassis and SeGW-4 is the standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Master       200       1
                                 IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3
```

```
===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Backup       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The following command disables the MS-ISA on the primary chassis SeGW-3, which will trigger an MC-IPSec failover.

```
configure {
    card 1 {
        mda 2 {
            admin-state disable
```

With MS-ISA disabled, the MC-IPSec state of tunnel group 1 on SeGW-3 becomes **notEligible**, which means that the tunnel group is down, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for details description of MIMP states.:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10
```

```
===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             notEligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is backup for VRRP instance 10 with in-use priority 50, as per the VRRP policy 1:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is now the primary chassis in tunnel group 1. This is triggered by MC-IPSec failover, as per the **mc-ipsec-non-forwarding** event in VRRP policy 1.

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is primary for VRRP instance 10;

```
[/]
A:admin@SeGW-4# show router vrrp instance
```

```
===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The situation is restored by enabling MS-ISA on SeGW-3:

```
configure {
    card 1 {
        mda 2 {
            admin-state enable
```

MC-IPSec failover when primary chassis reboots

The following **tools** command on SeGW-3 triggers an MC-IPSec switchover:

```
tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1

[/]
A:admin@SeGW-3# tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1
WARNING! Forcing a mastership switchover may significantly impact traffic. Are you sure (y/n)?
 y
```

Before the failure condition takes place, SeGW-3 is the primary chassis for tunnel group 1:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```

SeGW-3 is primary for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is the standby chassis for tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===========================================================================
ID             Peer Group     Priority  Admin State    Mastership
---------------------------------------------------------------------------
1              1              150       Up             standby
---------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===========================================================================
===============================================================================
```

The VRRP state on SeGW-4 is backup:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The following command reboots the primary chassis SeGW-3:

```
[/]
A:admin@SeGW-3# admin reboot card active now
```

While SeGW-3 reboots, the IPSec state of SeGW-4 becomes **eligible**:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-4 is primary (**master**):

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Master       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

When SeGW-3 comes up, the IPSec state of tunnel group 1 is **discovery**, which means that the system has not established the MIMP session with its peer yet.

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
```

```
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             discovery
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

After a while, the preceding **show** command is repeated and the IPSec state for tunnel 1 on SeGW-3 is standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-3 is backup:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is the primary chassis in MC-IPSec tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```

SeGW-4 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Configuration guidelines

The following is a list of guidelines for configuring MC-IPSec:

- To avoid high CPU load and issues in some complex cases, the following are suggestions for configuring the IKEv2 lifetime:
  - Both IKE_SA and CHILD_SA lifetime on MC-IPSec chassis (SeGW-3 and SeGW-4) should be around three times larger than on the IPSec peer CE-1.
  - With the first rule, the lifetime of the side with smaller lifetime (IPSec peer CE-1) should not be too small (these being the default values):
    - IKE_SA: >= 86400 seconds
    - CHILD_SA: >= 3600 seconds

  - With the first rule, on the side with smaller lifetime (IPSec peer CE-1), the IKE_SA lifetime must be at least 3 times larger than CHILD_SA lifetime.
- The IKE protocol is the control plane of IPSec, so IKE packets must be treated as high QoS priority in the end-to-end path of the public service. On the public interface, a SAP ingress QoS policy must be configured to ensure that IKE packets get high QoS priority.
- Configure **ipsec-responder-only true** under **tunnel-group** for static LAN-to-LAN tunnels.
- Enable dead peer detection (DPD) on the IPSec peer side (CE-1); disable DPD (default) on the MC-IPSec chassis side.
- The direct and redundant physical link between MC-IPSec chassis must be configured with sufficient bandwidth for MCS and shunting traffic, and proper QoS configuration to make sure the MIMP and MCS packets are treated as high priority traffic.
- The system time must be same on both MC-IPSec chassis.
- Make sure the protection status is **nominal** on both chassis before provoking a controlled switchover. The protection status can be displayed with the **show redundancy multi-chassis mc-ipsec peer ip-address <addr>** command.
- Wait at least five minutes between two consecutive switchovers if possible, to prevent a second switchover happening before the standby is ready to become the primary chassis.

Conclusion

MC-IPSec provides a stateful multi-chassis IPSec redundancy solution. This is very important in a carrier grade network, especially in applications such as mobile backhaul where high value mobile services run over IPSec tunnels.

N:M MC-IPsec Redundancy

This chapter describes N:M MC-IPsec redundancy.

Topics in this chapter include:

Applicability

The information and MD-CLI configuration in this chapter are based on SR OS Release 22.10.R1.

The IPsec tunnel termination configuration described in this chapter requires an MS-ISA2 or an ESA server configured with a virtual machine. Configuration and setup for ISA2 or ESA are beyond the scope of this chapter; see the Multi-Chassis IPSec Redundancy chapter.

Overview

The N:M MC-IPsec redundancy model is a feature of the multi-chassis (MC) capabilities of SR OS when the router is deployed as Security Gateway (SeGW). N:M aims at enhancing the existing 1:1 redundancy model for IPsec tunnels. For the definition of N:M terminology and a description of its benefits, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide*.

The figure Figure 50: Three-node redundancy domain with a 2 DA + 1 DS model shows a three-node redundancy domain (RD) with the SeGWs SEG1, SEG2, and SEG3. SEG 1 and SEG 2 are designated active (DA) SeGWs and SEG 3 is designated standby (DS) SeGW.

Figure 50: Three-node redundancy domain with a 2 DA + 1 DS model

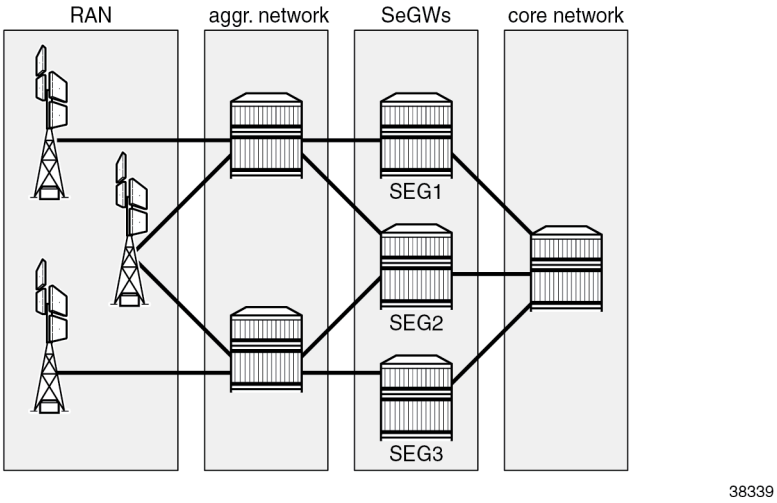

Radio access network (RAN) elements are opening IPsec tunnels toward SeGW cluster tunnel endpoint IP addresses. The RAN, aggregation network, and core network are emulated with standard routing nodes. For this deployment, assume that connectivity between elements is established using routing protocols and, as for a classic SeGW router, the public side where traffic is encrypted is built on top of a public-side VPRN, while private side (clear-text traffic) is associated with another VPRN. ISA2 or ESA resources manage encryption and decryption operations across the VPRN boundary.

This chapter describes configuration of SeGW elements, as well as MD-CLI commands for tracking the functionality of N:M nodes in the same redundancy domain (RD).

Configuration

Assume that IP connectivity is established across the IP network elements in the architecture. It is beyond the scope of this chapter to describe how traffic is carried from the RAN to the SeGW or from SeGW to the mobile packet core. Among the protocols and techniques that are required to speed up convergence of routing, the bidirectional forwarding detection (BFD) protocol is especially useful to keep network convergence time in a range compatible with mobile traffic use case.

ISA2 or ESA setup for N:M

The nodes participating in the IPsec domain have a standard setup for ISA2 or ESA resources.

SEG1 and SEG 2 can each be configured like a classic SeGW, as follows:

```
[gl:/configure isa]
A:admin@SEG1# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            isa 1/2 { }
```

```
            active-isa-number 1
        }
        reassembly {
            max-wait-time 1200
        }
        stats-collection {
            isa-dp-cpu-usage true
        }
    }
```

The **active-isa-number** command specifies the number of active encryption and decryption elements. Nokia recommends implementing the same number of ISA2 and ESA resources among the nodes participating in the RD, which allows for the DS node to activate the same number of ISA2 or ESA resources when failover occurs. However, a failover can occur even if the DS node has a lower number of ISA2 or ESA resources available in its local pool. This allows operators to save costs, but if the ISA2 ore ESA resources on the initial DA nodes were fully loaded, the DS node cannot host all tunnels and the protection is only partial.

N:M redundancy allows DS nodes to cover multiple TGs, and therefore, multiple RDs. DS nodes may have more ISA2 or ESA resources than the DA nodes, because the DS nodes should be able to cover one or more DA node failures, with a maximum of 16.

The output from SEG2 is the same as for SEG1.

SEG3 is configured as the DS node of the domain, where the configuration contains the **tunnel-member-pool** command:

```
[gl:/configure isa]
A:admin@SEG3# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            member-pool "MP1"
        }
        reassembly {
            max-wait-time 1200
        }
    }
    tunnel-member-pool "MP1" {
        isa 1/2 { }
    }
```

The **tunnel-member-pool** option defines the set of ISA2 or ESA resources used by the DS node during failures on active nodes. It is referenced in the tunnel group (TG) configuration, because multiple TGs can use the same tunnel member pool using the same set of ISA2 or ESA resources.

The output of the **show isa tunnel-member-pool** command lists ISA (ISA2 or ESA) members and their states. Under normal conditions, the ISA2 or ESA resource is not active on SEG3.

```
[gl:/configure isa]
A:admin@SEG3# /show isa tunnel-member-pool "MP1" detail

===============================================================================
ISA Tunnel Member Pool : MP1
Description            : (Not Specified)
Associated Tunnel Grps : 1
===============================================================================
Isa Members                   Active In Group     Last Configuration Change
-------------------------------------------------------------------------------
```

```
1/2                                                    11/25/2022 12:10:14
-------------------------------------------------------------------------------
Number of Configured Entries: 1
Number of Active Entries: 0
===============================================================================
```

Redundancy domain configuration

The configuration of MC-IPsec as N:M starts by defining node roles and behavior. The configuration on SEG1 (with system IP address 192.0.2.1) is as follows:

```
[gl:/configure redundancy]
A:admin@SEG1# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
            priority 250
            tunnel-group 1
        }
        peer 192.0.2.2  {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The preceding configuration example shows a multi-chassis IPsec domain, where the following domain characteristics have been specified:

- domain number – must be shared across all the nodes joining the redundancy domain (RD)
- designated role – DA or DS
- priority – required by the multi-chassis IPsec mastership protocol (MIMPv2) when an operationally active (OA) node must be elected. Setting a higher priority for an SeGW increases the likelihood of it being elected as the OA. In this case, SEG1 has the highest priority and DA role, so it is elected OA for RD 1.

- tunnel group – must be defined as per the ISA2 or ESA setup. The TG is always mapped to the RD in a 1:1 relationship
- peers – up to three peers can be added. While full-mesh peering between them is required, Nokia also recommends deploying highly redundant network paths between these peers.

  Each peer has its own CLI tree where the following characteristics must be defined:

  – the domain or domains the peer belongs to

  – the synchronization state for IPsec

  – whether BFD is applied to check peer liveliness.
- (optional) other parameters for keepalives, hold-time, and discovery-interval are configured with default values. Do not change these values unless a different setup is required under specific network conditions.

The configuration for the redundancy domain on SEG2 is the same as on SEG1, but with different IP addresses for peers and different priority:

```
A:admin@SEG2# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
            priority 240
            tunnel-group 1
        }
        peer 192.0.2.1 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The designated role of SEG2 is **active**, which means SEG2 behaves similarly to the 1:1 model where tunnel states are synchronized with SEG1 and immediately pushed to ISA2 or ESA resources. This behavior allows for a very quick failover when SEG1 experiences a failure.

The priority is 240, which is lower than for SEG1. As a result, SEG1 receives node role DA and is operationally active (OA) while SEG2 receives node role DA and is operationally standby (OS).

The RD configuration for DS SEG3 is as follows:

```
[gl:/configure redundancy multi-chassis]
A:admin@SEG3# info
    ipsec-domain 1 {
        admin-state enable
        designated-role standby
        priority 230
        tunnel-group 1
    }
    peer 192.0.2.1 {
        admin-state enable
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
    peer 192.0.2.2 {
        admin-state enable
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
```

The peer configuration is similar to those of other nodes where BFD liveliness is enabled.

The designated role is standby (DS). The default value in the configuration is not shown from the **info** command.

The priority is 230 but the node role is DS. The DS node will not become OA because the DA role of SEG1 and SEG2 always prevails when electing the OA, regardless of priority value. Therefore, a DS node can become OA only if there are no DA nodes available in the domain.

After the setup of MC IPsec RD is completed across all the nodes, **show** commands can be used to track RD behavior and state:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 250
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status  : nominal
Router Id          : 192.0.2.1         Current Active     : 192.0.2.1
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
```

```
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                                 Up      standby      active
  192.0.2.2
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

The output shows important information about the redundancy domain:

- the designated role of the node – active or standby
- the activity state based on fault conditions – active or standby
- the protection status – "nominal" means that the nodes are synchronized.
- the domain adjacencies – list of peers and their activity state and designated role
- the tunnel statistics – in this case, seven dynamic tunnels are established

The same **show** command executed on SEG2 provides similar output, with differences for the priority and the designated role. The seven tunnels are shown in the "Installed" state because SEG2 is a DA node.

The same **show** command on DS SEG3 shows the following:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby           Priority          : 230
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : nominal
Router Id          : 192.0.2.3         Current Active    : 192.0.2.1
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                                 Up      active       active
  192.0.2.1
192.0.2.2                                 Up      standby      active
  192.0.2.2
```

```
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                              Static              Dynamic
-------------------------------------------------------------------------------
Installed                     0                   0
Installing                    0                   0
Standby Dormant               0                   7
Awaiting Config               0                   0
Failed                        0                   0
===============================================================================
```

Relevant information from the SEG3 CLI output, apart from the activity state, the designated role, and the peer's state, is the tunnel state, which is now marked as "Standby Dormant".

Tunnels on SEG3 are not installed on the ISA2 or ESA; rather, they are stored in the router CPM and are kept ready to be offloaded on the ISA2 or ESA resources connected to the router. These tunnels are offloaded as soon as SEG3 becomes OA, following a node reboot, failure, or manual switchover.

Services configuration

The tunnels opened by RAN elements are terminated in a public-side VPRN IP address called TEIP (the public side can also be made on a IES service). Assume that the RAN nodes are using a single tunnel setup with a single IKE_SA, whereas the Child_SA's number is specific to the deployment. The configuration of this public side VPRN is the same for all three nodes and follows the standard SeGW setup:

```
[gl:/configure service vprn "100"]
A:admin@SEG1# info
    vprn "100" {
        admin-state enable
        description "public side"
        customer "1"
        ipsec {
            multi-chassis-shunt-interface "to_SEG2_Shunt" {
                next-hop {
                    address 10.1.12.2
                }
            }
            multi-chassis-shunt-interface "to_SEG3_Shunt" {
                next-hop {
                    address 10.1.13.2
                }
            }
            multi-chassis-shunting-profile "MCSPROF1" {
                peer 192.0.2.2 {
                    multi-chassis-shunt-interface "to_SEG2_Shunt"
                }
                peer 192.0.2.3 {
                    multi-chassis-shunt-interface "to_SEG3_Shunt"
                }
            }
        }
        interface "PUBLIC1" {
            multi-chassis-shunting-profile "MCSPROF1"
            sap tunnel-1.public:100 {
```

```
                ipsec-gateway "IPSECGW1" {
                    admin-state enable
                    default-tunnel-template 1
                    ike-policy 1
                    pre-shared-key "uCLxzS3PxoW0foPjmAKJ/Wv41hy6O3H76tg=" hash2
                    default-secure-service {
                        service-name "200"
                        interface "PRIVATE1"
                    }
                    local {
                        gateway-address 10.51.100.1
                    }
                }
            }
            ipv4 {
                primary {
                    address 198.51.100.2
                    prefix-length 24
                }
            }
        }
        interface "to_SEG2_Shunt" {
            spoke-sdp 2000:1 {
            }
            ipv4 {
                primary {
                    address 10.1.12.1
                    prefix-length 30
                }
            }
        }
        interface "to_SEG3_Shunt" {
            spoke-sdp 3000:1 {
            }
            ipv4 {
                primary {
                    address 10.1.13.1
                    prefix-length 30
                }
            }
        }
        ospf 0 {
            export-policy ["EXPORT_OSPF"]
        }
```

The parts of the configuration that are exclusive of N:M are those related to shunt-link setup.

The **multi-chassis-shunting-profile** command can be found under the **ipsec** configuration for the IES or VPRN service, where the multi-chassis shunting (MCS) profile is required to map each peer to a dedicated shunt interface. The MCS profile is referenced under the interface where the IPsec gateway is configured. In this scenario, peer 192.0.2.2 is reached through the to_SEG2_Shunt interface, which is defined under the same VPRN as an interface built on top of sdp:2000:1.

A full mesh of shunt interfaces is made across the RD, for both public and private side services.

```
A:admin@SEG1# show ipsec multi-chassis-shunt-interface service "100"

===============================================================================
IPsec Multi-Chassis Shunt Interfaces
===============================================================================
Service Id  MC Shunt Interface Name          Next Hop                Resolved
-------------------------------------------------------------------------------
100         to_SEG2_Shunt                    10.1.12.2               Yes
```

```
100         to_SEG3_Shunt                   10.1.13.2               Yes
-------------------------------------------------------------------------------
No. of IPsec MC Shunt Interfaces: 2
===============================================================================
```

The **show ipsec multi-chassis-shunt-interface service** command shows the liveliness of shunt interfaces and information on the next-hop resolution, whereas the **show ipsec multi-chassis-shunting-profile service** command provides a summary of the MCS profile and associated peers:

```
A:admin@SEG1# show ipsec multi-chassis-shunting-profile service "100"

===============================================================================
Multi-Chassis Shunting Profile Params Entries
===============================================================================
Service Id  MC Shunting Profile Name                MC Shunt Interface Name
            Peer
-------------------------------------------------------------------------------
100         MCSPROF1                                to_SEG2_Shunt
            192.0.2.2
100         MCSPROF1                                to_SEG3_Shunt
            192.0.2.3
-------------------------------------------------------------------------------
No. of IPsec MC Shunting Profile Params Entries: 2
===============================================================================
```

The SDP full mesh must be configured on both sides, as shown in the figure Figure 51: SDP full mesh.

Figure 51: SDP full mesh

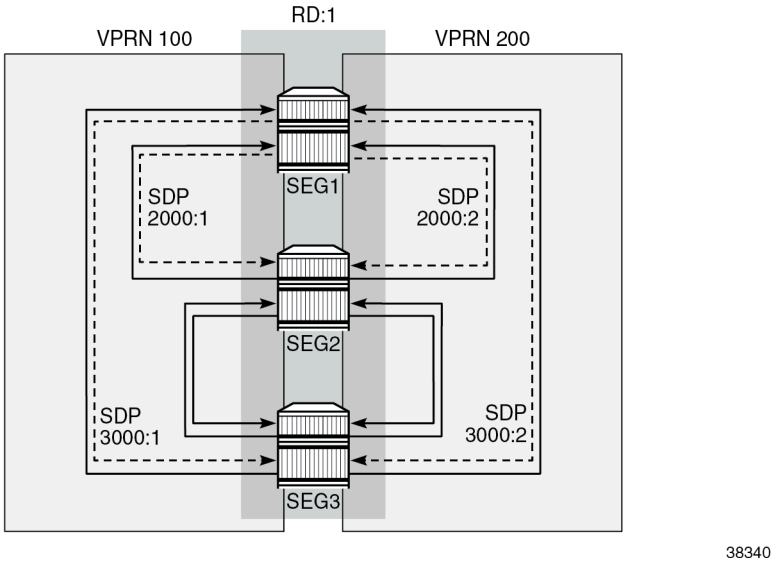

Note: Only the SDPs from SEG1 are shown with IDs.

The shunt link can be built from a standard spoke SDP or from a port-based interface. In this example, the following spoke SDPs are used in the public-side VPRN 100:

```
A:admin@SEG1# show service id "100" sdp
```

```
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:1           Spok     192.0.2.2       Up      Up        524285    524285
3000:1           Spok     192.0.2.3       Up      Up        524283    524285
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The **show** output for the private-side VPRN 200 looks similar to that for the public-side VPRN, except for the SDP IDs and label values:

```
A:admin@SEG1# show service id "200" sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:2           Spok     192.0.2.2       Up      Up        524284    524284
3000:2           Spok     192.0.2.3       Up      Up        524282    524284
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

There are no routing policy changes from the 1:1 MC-IPsec cluster, although this example could have a more complex routing setup, considering that the number of routers in a domain is higher than in the 1:1 model. The following configuration shows the SEG1-2-3 export policy used on the public side where the OSPF protocol is used under VPRN 100:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG1# info
    description "EXPORT TEIP OSPF - PUBLIC SIDE"
    entry 10 {
        from {
            state ipsec-master-with-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 30
            }
        }
    }
    entry 20 {
        from {
            state ipsec-non-master
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
```

```
            tag 100
            metric {
                set 190
            }
        }
    }
    entry 30 {
        from {
            state ipsec-master-without-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 40
            }
        }
    }
    default-action {
        action-type reject
    }
```

On SEG2, only the metrics are different and are aligned with DA priorities:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG2# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 60
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 190
                }
            }
        }
        entry 30 {
            from {
```

```
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 50
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

On SEG3, the export policy is as follows:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG3# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 90
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 195
                }
            }
        }
        entry 30 {
            from {
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
```

```
                tag 300
                metric {
                    set 60
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

The export policy on the private-side VPRN is made with the same concept as the public side, but is not shown here.

Note: Parts of the configuration where the parameters remain the same as those in classic SeGW deployments (either stand-alone or 1:1) have not been added to this chapter. This information is described in the Multi-Chassis IPSec Redundancy chapter.

On the private side of SeGWs, a different VPRN is required, as per standard IPsec configuration. The private-side VPRN configuration on SEG1 is as follows:

```
[gl:/configure service vprn "200"]
A:admin@SEG1# info
    admin-state enable
    description "private segw testing"
    customer "1"
    ipsec {
        multi-chassis-shunt-interface "to_SEG2_Shunt" {
            next-hop {
                address 10.2.12.2
            }
        }
        multi-chassis-shunt-interface "to_SEG3_Shunt" {
            next-hop {
                address 10.2.13.2
            }
        }
        multi-chassis-shunting-profile "MCSPROF1" {
            peer 192.0.2.2 {
                multi-chassis-shunt-interface "to_SEG2_Shunt"
            }
            peer 192.0.2.3 {
                multi-chassis-shunt-interface "to_SEG3_Shunt"
            }
        }
    }
    bgp-ipvpn {
        mpls {
            admin-state enable
            route-distinguisher "300:4"
        }
    }
    interface "PRIVATE1" {
        tunnel true
        multi-chassis-shunting-profile "MCSPROF1"
        sap tunnel-1.private:100 {
        }
    }
    interface "to_SEG2_Shunt" {
        ipv4 {
            primary {
                address 10.2.12.1
```

```
                prefix-length 30
            }
        }
        spoke-sdp 2000:2 {
        }
    }
    interface "to_SEG3_Shunt" {
        ipv4 {
            primary {
                address 10.2.13.1
                prefix-length 30
            }
        }
        spoke-sdp 3000:2 {
        }
    }
```

As the configuration shows, the same setup of shunt links is required on the private side to allow path resiliency in case of faults for the traffic going downstream from core toward the RAN.

Failure scenario – active node experiences a power failure

N:M can be triggered by different fault conditions, such as a complete node failure, an ISA2 or ESA failure, or a manual switchover executed with the **tools** command. In this scenario, complete node failures are simulated. When SEG1 experiences a node failure, SEG2 takes over. When SEG2 fails too, SEG3 takes over and remains the only node with active tunnels.

The initial scenario has SEG1 and SEG2 configured as DA nodes, while SEG3 is the DS node for the domain configured as **ipsec-domain 1**. The state can be verified with the **show redundancy multi-chassis ipsec-domain 1** command (as shown above in the Redundancy domain configurationsection).

As soon as SEG1 experiences a node failure, SEG2 takes over:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 240
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.2         Current Active    : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote      Remote
  Router-Id                             State   Activity    Designated
                                                State       Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown     unknown
  0.0.0.0
192.0.2.3                               Up      standby     standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================
```

```
===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

Although the protection status, as seen from SEG2 and SEG3, is initially "notReady", it changes to "nominal" after few minutes. From the SEG2 and SEG3 point of view, SEG1 is unreachable, and its activity state remains unknown. Log 99 also records the failure event:

```
A:admin@SEG2# show log log-id 99

===============================================================================
Event Log 99 log-name 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=187  (not wrapped)]

186 2022/12/13 14:05:32.534 UTC WARNING: MC_REDUNDANCY #2047 Base MC-IPSEC-DOMAIN 1
"Protection status for the multi-chassis ipsec domain 1 changed to nominal"

185 2022/12/13 14:02:19.611 UTC MINOR: VRTR #2061 Base 192.0.2.1
"BFD: Local Discriminator 1 BFD session on node 192.0.2.1 is down due to noHeartBeat "

---snip---

179 2022/12/13 14:02:19.124 UTC WARNING: MC_REDUNDANCY #2004 Base
"The Sync status of peer 192.0.2.1 changed to outOfSync"

178 2022/12/13 14:02:18.746 UTC WARNING: MC_REDUNDANCY #2046 Base MC-IPSEC-DOMAIN 1
"Multi-chassis ipsec domain 1 local activity state changed from standby to active because an
 inter-chassis link went down. The active router in the domain is 192.0.2.2."
```

Next, SEG2 also experiences a full node failure, and SEG3 takes over:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby           Priority          : 230
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.3         Current Active    : 192.0.2.3
Activity State     : eligible
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown      unknown
  0.0.0.0
```

```
192.0.2.2                                   Down    unknown      unknown
  0.0.0.0
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                             Static              Dynamic
-------------------------------------------------------------------------------
Installed                    0                   7
Installing                   0                   0
Standby Dormant              0                   0
Awaiting Config              0                   0
Failed                       0                   0
===============================================================================
```

Both SEG1 and SEG2 are seen as operationally down with an unknown activity state. On SEG3, the tunnel states have been copied from the CPM to the ISA2 or ESA entities and are now shown as "Installed", rather than "Standby Dormant". As soon as SEG1 or SEG2 are back up, the **revertive** flag configured under the **ipsec-domain** command determines if the tunnels are kept on the current active DS node or if they are moved back to SEG1 ownership.

Failure scenario – using the tools command line

A planned failure condition is commonly seen when executing software upgrades or hardware maintenance on SeGW nodes, which leverages the **tools** command line utility to move tunnels toward other peering nodes.

The initial state is the same as for the previous example where SEG1 is initially the operationally active DA.

The following tools command triggers a switchover and therefore causes all the tunnels installed on the operationally active DA node to move on another node in the domain, selected by the **auto** flag in this case.

```
A:admin@SEG1# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 auto
 now
```

To specify a peer IP address among those available in the domain, the **to <peer_ip>** option could be used instead of **auto**.

The following output shows the domain state as seen from SEG1 after the execution of the tools command:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 250
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.1         Current Active    : 192.0.2.2
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
```

```
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                               Up      active       active
  192.0.2.2
192.0.2.3                               Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

As shown in the output, the current active node is SEG2 (192.0.2.2). The **auto** flag forced all the traffic to move across the second preferred active node in the domain, which is SEG2.

The protection status, as seen from SEG2, changes to "nominal" after a few minutes:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active             Priority           : 240
Tunnel Group       : 1                  Revertive          : false
Admin State        : Up                 Protection Status  : nominal
Router Id          : 192.0.2.2          Current Active     : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                                 Up      standby      active
  192.0.2.1
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
```

```
Installing                    0                   0
Standby Dormant               0                   0
Awaiting Config               0                   0
Failed                        0                   0
===============================================================================
```

After maintenance operations on SEG1 have been completed and the node is operational (which can be verified using the **show** commands described in this chapter), the operator reverts services and traffic back to SEG1. For this purpose and in this specific example, the same **tools** command can be used. The **auto** flag selects SEG1, according to its highest priority in the domain. If more predictability is required in the selection choice, the **to <peer_ip>** flag can be used, as in this example:

```
A:admin@SEG2# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 to
 192.0.2.1 now
```

Conclusion

N:M adds a level of redundancy to an already efficient redundancy model; it ensures that RAN elements stay connected to the core network under a wide range of failure conditions. SR OS uses a full set of commands to implement this feature, available for both classic and MD-CLI. N:M also gives network engineers and architects the capability to deploy SeGW services with greater flexibility; for example, to deploy super-resilient SeGW clusters to serve high-density RAN areas, or to introduce cost-optimized solutions with an acceptable level of automated fault recovery.

Customer document and product support

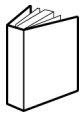

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback