



7450 Ethernet Service Switch 7750 Service Router Virtualized Service Router

Releases up to 25.3.R2

Multiservice ISA and ESA Advanced Configuration Guide for MD CLI

3HE 20799 AAAD TQZZA
Edition: 01
July 2025

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.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

List of tables..... 4

List of figures.....5

Preface..... 6

Application Assurance — Security Gateway Stateful Firewall.....7

Application Assurance — Stateful Firewall.....38

Multi-Chassis IPSec Redundancy.....60

N:M MC-IPsec Redundancy.....98

List of tables

Table 1: SCTP PPIDs.....21

Table 2: GTP messages.....31

List of figures

Figure 1: LTE SeGW firewall deployment.....8

Figure 2: SeGW in small cells architecture..... 8

Figure 3: Configuration topology..... 10

Figure 4: Block unsolicited traffic..... 39

Figure 5: SFW — allow gaming.....40

Figure 6: ALG support example — FTP.....41

Figure 7: Configuration topology..... 43

Figure 8: MC-IPSec architecture.....61

Figure 9: Example topology..... 62

Figure 10: Three-node redundancy domain with a 2 DA + 1 DS model..... 99

Figure 11: SDP full mesh..... 107

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 25.3.R2. 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 — Security Gateway Stateful Firewall

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

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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 1: 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 1: 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 2: SeGW in small cells architecture](#):

Figure 2: 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.

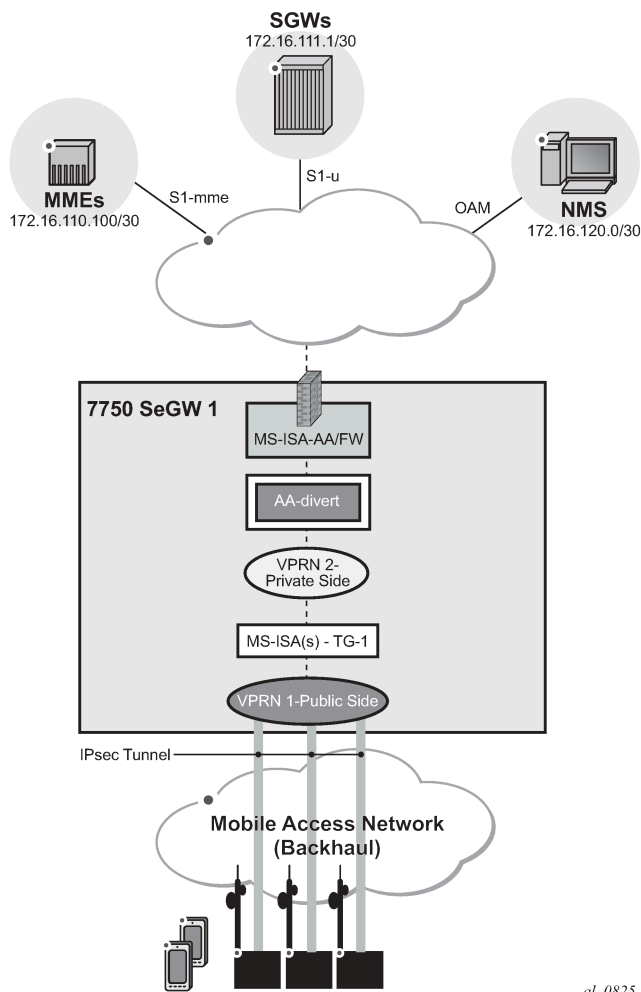
See the configuration described in the [Application Assurance — Stateful Firewall](#) chapter for this context of OAM protection in SeGW.

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 3: 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 3: Configuration topology



al_0825

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](#)
- [2.2 Allow only SCTP traffic toward MMEs — no port scanning](#)
- [2.3 DoS protection — limit the number of SCTP flows from eNBs](#)
- [2.4 DoS protection — limit the SCTP bandwidth from eNB](#)
- [2.5 Allow only specified SCTP PPIDs toward the MMEs](#)

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 1: SCTP PPIDs](#).

Table 1: 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 SIP 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-lockup 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 $32 \times 8 = 256$ for residential deployments or $8 \times 8 = 64$ 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 2: GTP messages](#).

Table 2: 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 2: 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](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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 4: Block unsolicited traffic



To support the example shown in [Figure 4: 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 5: SFW — allow gaming


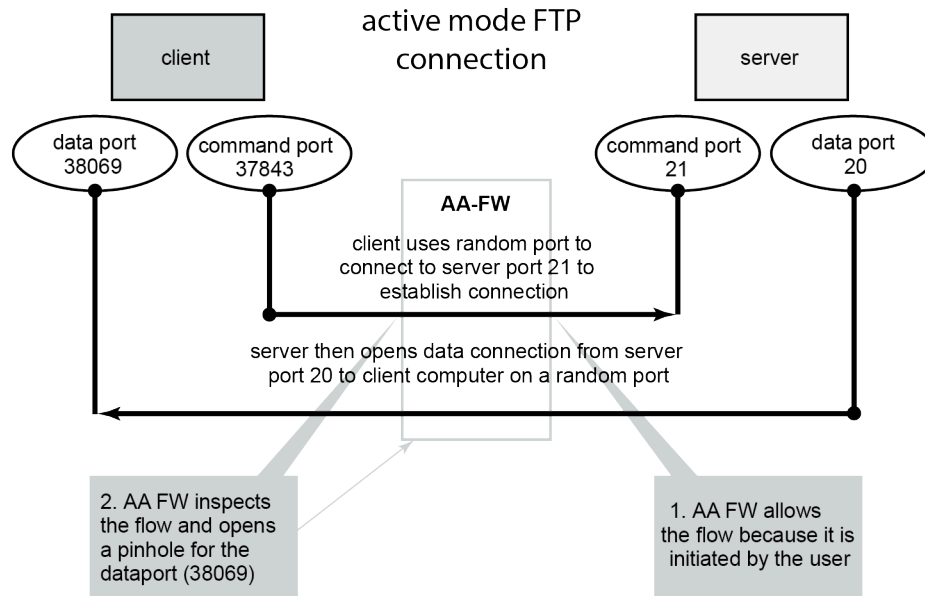


Figure 5: SFW — allow gaming shows a similar concept. It is used to allow UDP traffic for peer to peer applications, such as gaming. When the traffic from one peer is seen by AA-ISA, a pin-hole is opened in the reverse direction to allow for the corresponding UDP traffic from the peer.

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 6: ALG support example — FTP



al_0256

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 6: 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 7: Configuration topology](#).

Figure 7: 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 7: Configuration topology](#).

Subscribers who are not protected (for example subscriber 3 in [Figure 7: 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 and 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 "denyUnsolicitedwMgntCntrl" is configured as follows:

```
configure {
  application-assurance {
    group 1 {
      partition 2 {
        session-filter "denyUnsolicitedwMgntCntrl" {
          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 "denyUnsolicitedwMgntCntrl"
                        }
                    }
                }
            }
        }
    }
}
```

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 0N

Action :
  Session Filter         : denyUnsolicitedWmgntCntrl
=====
```

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
denyUnsolicitedwMgntCntrl          deny             aqp              4
-----
No. of session filters: 2
=====
```

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

=====
AA Session Filter Instance "denyUnsolicitedwMgntCntrl"
=====
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
=====

-----
Packet Validation Statistics
-----
              Admitted    Denied    Admitted    Denied
              Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
              (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

-----
Session Filter Statistics
-----
              Admitted    Denied    Admitted    Denied
              Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
              (Packets)   (Packets)   (Packets)   (Packets)
Session Filter: protectISPlan2
Entry: 10                      0          0          0          0
Default Action                 0          0          0          0

Session Filter: denyUnsolicitedwMgntCntrl
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

-----
Flow Policer Statistics
-----
              Admitted    Denied    Admitted    Denied
              Sub-To-Net  Sub-To-Net  Net-To-Sub  Net-To-Sub
              (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.

Multi-Chassis IPsec Redundancy

This chapter provides information about multi-chassis IPsec redundancy configurations.

Topics in this chapter include:

- [Applicability](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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 8: MC-IPsec architecture](#) shows two redundant IPsec chassis in the middle: a primary chassis and a standby chassis.

Figure 8: MC-IPSec architecture



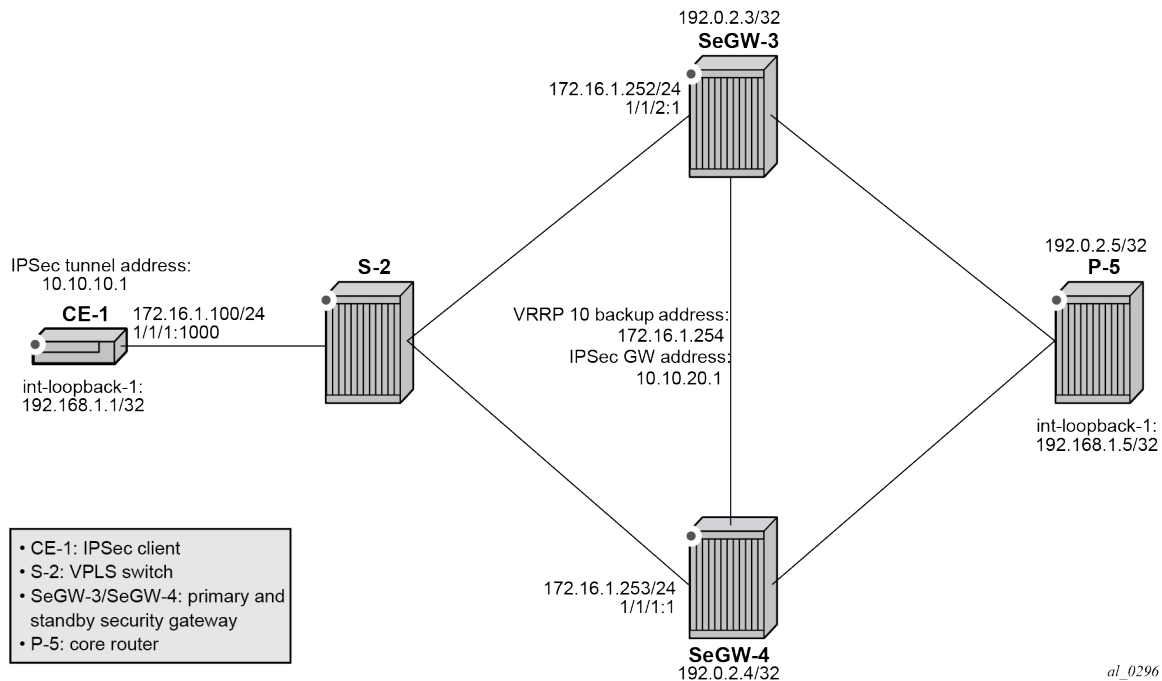
The fundamentals of MC-IPSec are:

- 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 9: Example topology](#).

Figure 9: 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 ["vpn2"]
            }
        }
    }
    default-action {
        action-type accept
        community {
            add ["vpn2"]
        }
    }
}
```

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-nextthop 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](#)
- [Overview](#)
- [Configuration](#)
- [Conclusion](#)

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 10: 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 10: Three-node redundancy domain with a 2 DA + 1 DS model



38339

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 or 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 liveness.
- (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 liveness 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 Router-Id	Oper State	Remote Activity State	Remote Designated 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 Router-Id	Oper State	Remote Activity State	Remote Designated 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/Wv41hy603H76tg=" 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 liveness 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 11: SDP full mesh](#).

Figure 11: SDP full mesh



38340



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 configuration](#) section).

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 Router-Id      Oper State  Remote Activity State  Remote Designated 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 Router-Id	Oper State	Remote Activity State	Remote Designated Role
-----	-----	-----	-----
192.0.2.1 0.0.0.0	Down	unknown	unknown

```

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 Router-Id	Oper State	Remote Activity State	Remote Designated 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 Router-Id	Oper State	Remote Activity State	Remote Designated 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](#)