7450 Ethernet Service Switch
7750 Service Router
Virtualized Service Router

Releases up to 24.7.R2

Multiservice ISA and ESA Advanced Configuration Guide for MD CLI

3HE 20799 AAAB TQZZA
Edition: 01
October 2024

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

Table of contents

List of figures

Preface

About This Guide

Each Advanced Configuration Guide is organized alphabetically and provides feature and configuration explanations, CLI descriptions, and overall solutions. The Advanced Configuration Guide chapters are written for and based on several Releases, up to 24.7.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.

Multi-Chassis IPSec Redundancy

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

Topics in this chapter include:

Applicability

This initial version of this chapter was based on SR OS Release 10.0.R8, but the MD-CLI in the current edition corresponds to SR OS Release 22.10.R2.

Overview

Multi-Chassis IPSec redundancy (MC-IPSec) is a stateful inter-chassis IPSec failover mechanism. IPSec tunnel states are synchronized between the primary and standby chassis. A tunnel group failure on the primary chassis or a primary chassis failure could trigger MC-IPSec failover to the standby chassis.

The following are some highlights of this feature:

- Internet Key Exchange version 2 (IKEv2) only
- Multi-active tunnel group only
- The granularity of failover is tunnel group, which means a specific tunnel group could failover to the standby chassis independent of other tunnel groups on the primary chassis
- Both static and dynamic LAN-to-LAN tunnels are supported

This feature has the following building blocks:

- Primary chassis election: MC-IPSec mastership protocol (MIMP) runs between the chassis to elect a primary chassis with independent MIMP runs for each tunnel group
- Synchronization: multi-chassis synchronization (MCS) synchronizes the IPSec states between chassis
- Routing:
  - MC-IPSec-aware routing attracts traffic to the primary chassis
  - Shunting support
  - MC-IPSec-aware virtual router redundancy protocol (VRRP)

The figure Figure 1: MC-IPSec architecture shows two redundant IPSec chassis in the middle: a primary chassis and a standby chassis.

Figure 1: MC-IPSec architecture

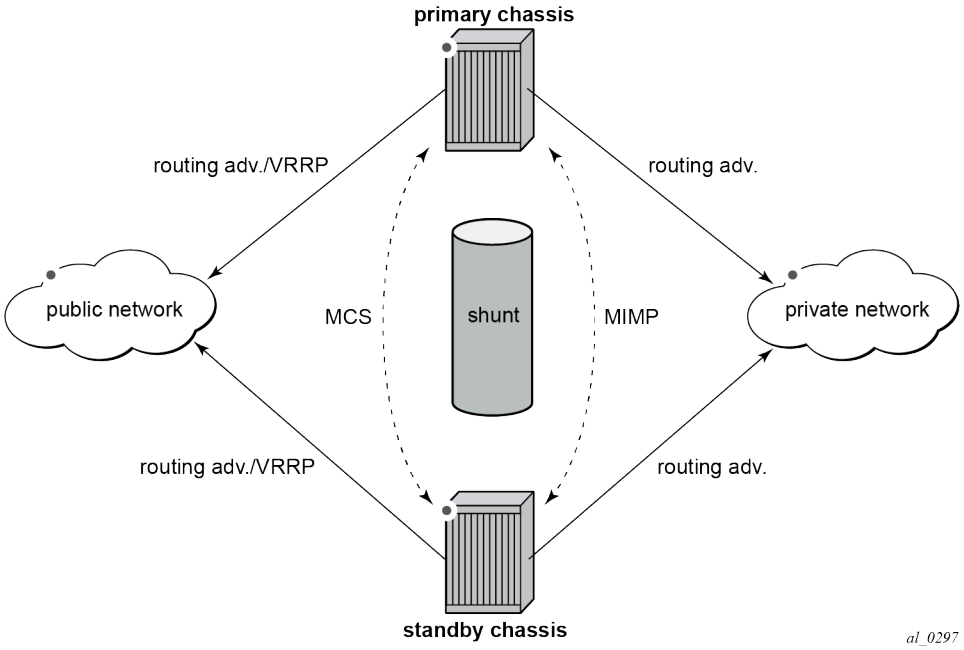

The fundamentals of MC-IPSec are:

- 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 2: Example topology.

Figure 2: Example topology

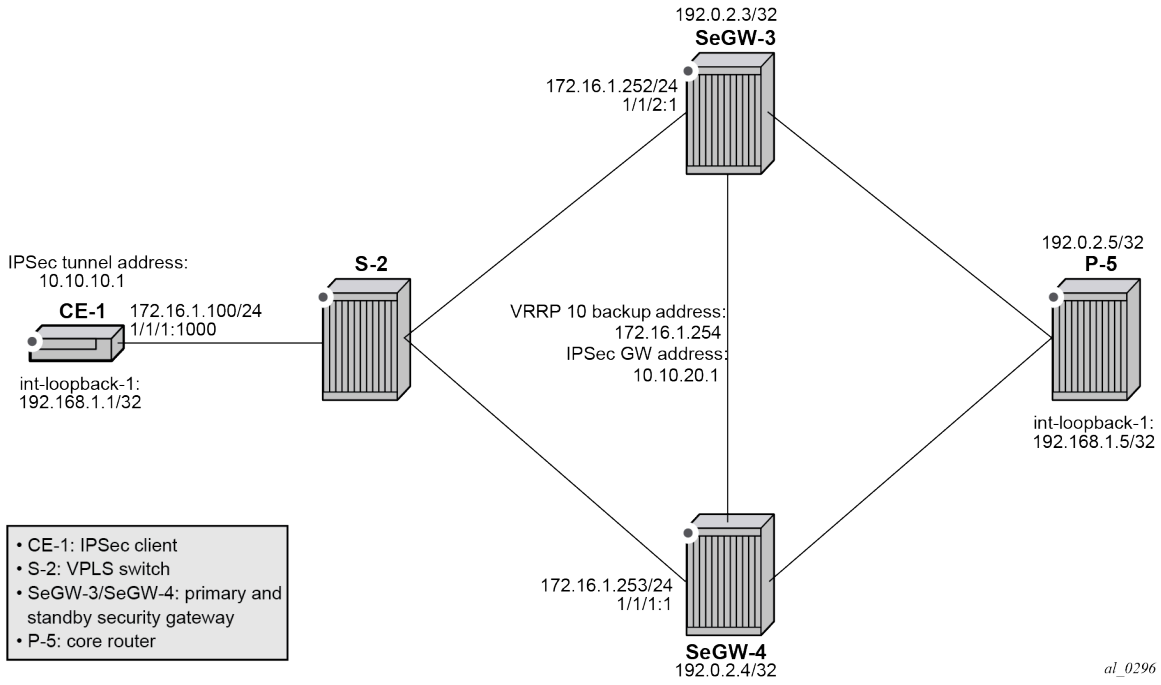

The example setup includes:

- an IPSec tunnel initiated by CE-1 and terminated on the primary chassis of the two SeGWs.
- a public IES service "IES-1" and a private VPRN service "VPRN-2" configured on CE-1, SeGW-3, and SeGW-4.
- VPRN "VPRN-2" (also) configured on P-5.
- a static LAN-to-LAN tunnel with pre-shared key.
- a local VPLS service "VPLS-3" on S-2 to simulate a Layer 2 switch.
- VRRP 10 between SeGW-3 and SeGW-4 to provide a backup address 192.168.1.254, which is the default next hop for CE-1.
- VRRP policy 1 bound to VRRP 10 on the primary chassis SeGW-3 to change the in-use priority upon MC-IPSec switchover.
- OSPF as IGP running in the base routing instance between SeGW-3, SeGW-4, and P-5.
- MP-BGP running between SeGW-3, SeGW-4, and P-5 for the VPN-IPv4 address family.

A ping in VPRN "VPRN-2" between loopback interface address 192.168.1.1 on CE-1 and 192.168.1.5 on P-5 is used to verify the connectivity over the IPSec tunnel.

The MC-IPSec configuration commands are shown below.

```
configure
    redundancy
        multi-chassis
            peer <ip-address>
                sync
                    ipsec
                    tunnel-group <1..64>
```

```
                        sync-tag <string>

                mc-ipsec
                    bfd-liveness <boolean>
                    discovery-interval
                        interval-secs <1..1800>
                        boot <1..1800>
                    hold-on-neighbor-failure <2..25>
                    keep-alive-interval <5..500>     # deciseconds
                    tunnel-group <1..64>
                        admin-state <boolean>
                        peer-group <1..64>
                        priority <0..255>
```

```
configure
    policy-options
        policy-statement <string>
            entry <1..4294967295>
                from
                    state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer
                    protocol
                        name ipsec
```

```
configure
    service
        ies <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>


configure
    service
        vprn <string>
            interface <string>
                dynamic-tunnel-redundant-nexthop <unicast-ipv4-address>
                static-tunnel-redundant-nexthop <unicast-ipv4-address>
```

```
configure
    isa
        tunnel-group <1..64>
            ipsec-responder-only <boolean>
```

```
configure
    vrrp
        policy <1..9999>
            priority-event
                mc-ipsec-non-forwarding <tunnel-grp-id>
                    hold-clear <1..86400 seconds>
                    hold-set <1..86400 seconds>
                    priority
                        priority-level <1..254>
                        event-type (delta|explicit)
```

The parameters are the following:

- in the **configure redundancy multi-chassis** context:

- **peer** <*ip-address*> — This command creates or enters a multi-chassis peer. The peer address is by default the system address. This can be changed on the peer using the **configure redundancy multi-chassis peer source-address** command.
  - **sync** — This command enters the sync configuration context.
    - **ipsec** <*boolean*> — This command enables MCS to synchronize IPSec states.
    - **tunnel-group** <*tunnel-group-id*> **sync-tag** <*tag-name*> — This command enables MCS to synchronize the IPSec states of the specified tunnel group. The **sync-tag** parameter is used to match the tunnel group of the peer. The tunnel group states with the same **sync-tag** on both chassis will be synchronized.
  - **mc-ipsec** — This command enters the multi-chassis IPSec configuration context.
    - **bfd-liveness** <*boolean*> — The command **bfd-liveness** *true* enables tracking a central BFD session; if the BFD session goes down, then the system considers the peer as down and changes the MC-IPSec status of the configured tunnel group accordingly.

      The BFD session uses the source address of MCS as its source address and the MCS peer address as the destination address. Other BFD parameters are configured in the **bfd** context on the interface that the MCS source address resides on.

      The configuration of BFD is optional for MC-IPSec.
    - **discovery-interval interval-secs** <*interval-1*> **[boot** <*interval-2*>**]** — This command specifies the time interval that the tunnel group stays in **discovery** state. Interval 1 is used as discovery interval when a new tunnel group is added to multi-chassis redundancy (**mp-ipsec**); interval 2 is used as discovery interval after system boot-up. Interval 2 is optional, and when it is not specified, the value for interval 1 is used. Both intervals have a default value of 300 seconds.
    - **hold-on-neighbor-failure** <*2..25*> — This command specifies the number of keep-alive failures before considering the peer to be down. The default value is 3.
    - **keep-alive-interval** <*5..500*> — This command specifies the time interval of the mastership election protocol keep-alive packets in deciseconds. The default value is 10 deciseconds (1 s).
    - **tunnel-group** <*tunnel-group-id*> — This command enables multi-chassis redundancy for the specified tunnel group, or enters an already configured tunnel group context. The configured tunnel groups can failover independently.

      — **peer-group** <*tunnel-group-id*> — This command specifies the corresponding tunnel group ID on the peer node. The peer tunnel group ID is not necessarily equal to local tunnel group ID.

      — **priority** <*priority*> — This command specifies the local priority of the tunnel group, this is used to elect a primary chassis, where the higher number prevails. If the priorities are the same, then the peer which has more active ISAs wins; if the priority and the number of active ISAs are same, then the peer with higher IP address wins. The range is from 0 to 255 and the default value is 100.

- in a **from** statement of a route policy entry:
  - **state ipsec-master-with-peer | ipsec-non-master | ipsec-master-without-peer** — These commands specify the MC-IPSec state in a **from** statement of a route policy entry:
    - **ipsec-master-with-peer**: The tunnel group is the primary chassis with a peer reachable.
    - **ipsec-master-without-peer**: The tunnel group is the primary chassis with peer unreachable.

- **ipsec-non-master**: The tunnel group is not the primary chassis.

– **protocol name ipsec** — This command specifies IPSec as protocol in a **from** statement of a route policy entry. **protocol name ipsec** refers to the /32 local gateway routes (of both static and dynamic tunnels) and reverse route of dynamic tunnel.

- on a public or private IPSec interface in an IES or VPRN service:

– **static-tunnel-redundant-nexthop** <*ip-address*> and **dynamic-tunnel-redundant-nexthop** <*ip-address*> — These commands specify the redundant next hop address on a public or private IPSec interface (with public or private tunnel SAP) for a static and dynamic IPSec tunnel respectively. The specified next hop address is used by the standby chassis to shunt traffic to the primary chassis in case it receives any traffic. The next hop address is resolved in the routing table of the corresponding service.

Note:

- Shunting is supported over:
  - directly connected SAPs
  - spoke SDP terminated IP interfaces
- Shunting over auto-bind tunnel is not supported.
- Shunting does not work if the tunnel group is down.

- in the **isa tunnel-group <id>** context:

– **ipsec-responder-only** <*boolean*> — With the command **ipsec-responder-only** *true* , the system only acts as IKE responder except for the automatic CHILD_SA rekey upon MC-IPSec switchover.

This command is required for MC-IPSec support of static LAN-to-LAN tunnels.

- in the **vrrp policy <id> priority-event** context:

– **mc-ipsec-non-forwarding** <*tunnel-grp-id*> — This command creates a VRRP policy priority event: *mc-ipsec-non-forwarding*, which is triggered whenever the specified tunnel group enters the non-forwarding state.

- **hold-clear** <*seconds*> — This command configures the hold time before clearing the event. The range is from 0 to 86400 seconds and the default value is 0 s.
- **hold-set** <*seconds*> — This command configures the hold time before setting the event. The range is from 0 to 86400 seconds and the default value is 0 s.
- **priority** <*priority-level*> **explicit** — This command sets the VRRP in-use priority to the configured value upon the event. The range is from 0 to 254 and the default value is 0.

The initial configuration must include the following:

- The system time of SeGW-3 and SeGW-4 must be the same for the feature to work. Nokia recommends to use a time synchronization protocol such as NTP or SNTP.
- SeGW-3 and SeGW-4 must be IP reachable in the base routing instance because both MCS and MIMP run in the base routing instance.

Configuration of MC-IPSec

In this section, the following steps are described:

- configure CE-1
- configure S-2
- configure P-5
- configure IPSec tunnel on SeGW-3
- enable MC-IPSec for tunnel group on SeGW-3
- configure MC-IPSec-aware routing on SeGW-3
- configure MC-IPSec-aware VRRP on SeGW-3
- configure SeGW-4

Configure CE-1

On CE-1, the following is configured:

- a public IES service "IES-1" and a private VPRN service "VPRN-2".
- a static default route pointing to the VRRP backup address 172.16.1.254.
- a static IPSec tunnel "tunnel-1" with local address 10.10.10.1 and remote address 10.10.20.1.
- a loopback interface in VPRN "VPRN-2" with address 192.168.1.1/32 to be used as source address for the ping command to verify the connectivity between CE-1 and P-5 over the IPSec tunnel.

The following base router configuration on CE-1 includes a static route with next hop 172.16.1.254, which is the VRRP backup address.

```
# on CE-1:
configure {
    router "Base" {
        interface "int-CE-1-S-2" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 172.16.1.100
                    prefix-length 24
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 172.31.2.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 0.0.0.0/0 route-type unicast {
                next-hop "172.16.1.254" {       # VRRP backup address
                    admin-state enable
                }
            }
        }
```

IPSec is configured as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ike-transform [1]
            ike-version-2 {
            }
            dpd {    # dead peer detection (on peer side; not on MC-IPSec chassis)
            }
        }
        ike-transform 1 {
        }
        ipsec-transform 1 {
        }
```

Tunnel group 1 is configured as follows:

```
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            primary 1/2
        }
```

The public IES service is configured as follows:

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.10.254
                        prefix-length 24
                    }
                }
            }
        }
```

The private VPRN service on CE-1 is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.1/32
                        }
                        remote-ip {
```

```
                        address 192.168.1.5/32
                    }
                }
            }
        }
        interface "int-IPsec-private-1" {
            tunnel true
            sap tunnel-1.private:1 {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                    key-exchange {
                        dynamic {
                            ike-policy 1
                            ipsec-transform [1]
                            pre-shared-key "pass"
                        }
                    }
                    tunnel-endpoint {
                        local-gateway-address 10.10.10.1
                        remote-ip-address 10.10.20.1
                        delivery-service "IES-1"
                    }
                    security-policy {
                        id 1
                    }
                }
            }
        }
        interface "int-loopback-1" {
            loopback true
            ipv4 {
                primary {
                    address 192.168.1.1
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 192.168.1.5/32 route-type unicast {
                ipsec-tunnel "tunnel-1" {
                    admin-state enable
                }
            }
        }
```

Configure S-2

On S-2, a local VPLS service 3 simulates a Layer 2 switch between CE-1, SeGW-3, and SeGW-4:

```
# on S-2:
configure {
    service {
        vpls "VPLS-3" {
            admin-state enable
            service-id 3
            customer "1"
            sap 1/1/c1/1:1 {
                description "to SAP in IES 1 on SeGW-3"
            }
            sap 1/1/c1/2:1000 {
```

```
                description "to router interface in CE-1"
            }
            sap 1/1/c1/3:1 {
                description "to SAP in IES 1 on SeGW-4"
            }
        }
```

Configure P-5

P-5 simulates the core network router, connecting to SeGW-3 and SeGW-4. The configuration on P-5 includes the following:

- a loopback interface with address 192.168.1.5/32 in VPRN "VPRN-2", which is the destination address of the ping traffic from CE-1.
- an MP-BGP session for the VPN-IPv4 address family between P-5, SeGW-3, and SeGW-4.
- GRE spoke SDPs to connect to SeGW-3 and SeGW-4.

On P-5, the following router interfaces are configured in the base router. OSPF is used as IGP.

```
# on P-5:
configure {
    router "Base" {
        interface "int-P-5-SeGW-3" {
            port 1/1/c1/2:1000
            ipv4 {
                primary {
                    address 192.168.35.2
                    prefix-length 30
                }
            }
        }
        interface "int-P-5-SeGW-4" {
            port 1/1/c1/1:1000
            ipv4 {
                primary {
                    address 192.168.45.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                primary {
                    address 192.0.2.5
                    prefix-length 32
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-P-5-SeGW-3" {
                }
                interface "int-P-5-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```

On P-5, the following GRE SDPs are configured toward SeGW-3 and SeGW-4:

```
configure {
    service {
        sdp 53 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 54 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
```

VPRN "VPRN-2" is configured on P-5, as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-loopback-1" {
                loopback true
                ipv4 {
                    primary {
                        address 192.168.1.5
                        prefix-length 32
                    }
                }
            }
            spoke-sdp 53:2 {
            }
            spoke-sdp 54:2 {
            }
        }
```

The BGP configuration on P-5 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
```

```
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
        }
```

Configure IPSec tunnel on SeGW-3

The configuration on SeGW-3 is described in four consecutive sections. In this first section, the following is configured:

- the tunnel group, which must be in multi-active mode before MC-IPSec can be enabled.
- an interface "int-Redundant-1", which is a spoke-SDP terminated interface used for shunting.
- GRE SDP 34 toward SeGW-4 and GRE SDP 35 toward P-5.
- IPSec tunnel "tunnel-1" is the tunnel to CE-1; both SeGW-3 and SeGW-4 use the same local gateway address: 10.10.20.1.

The following configures tunnel group 1 on SeGW-3:

```
# on SeGW-3
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

On SeGW-3, the following router interfaces are configured in the base router. A static route is configured toward CE-1. OSPF is the IGP used between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-3-P-5" {
            port 1/1/1:1000
            ipv4 {
                primary {
                    address 192.168.35.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-3-SeGW-4" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.1
                    prefix-length 30
                }
            }
```

```
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-3-P-5" {
                }
                interface "int-SeGW-3-SeGW-4" {
                }
                interface "system" {
                }
            }
        }
```

The IPSec settings are as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The GRE SDPs are configured as follows:

```
configure {
    service {
        sdp 34 {
            admin-state enable
            description "GRE SDP toward SeGW-4"
            signaling off
            far-end {
                ip-address 192.0.2.4
            }
        }
        sdp 35 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
```

```
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows. In a later step, a VRRP policy will be configured and applied.

```
configure {
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.2
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                    }
                }
            }
```

The private VPRN service "VPRN-2" is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
```

```
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                }
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.2
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.1
                        prefix-length 30
                    }
                }
                spoke-sdp 34:20 {
                    ingress {
                        vc-label 2049
                    }
                    egress {
                        vc-label 2048
                    }
                }
            }
            spoke-sdp 34:2 {
                description "SDP to SeGW-4"
            }
            spoke-sdp 35:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
```

Enable MC-IPSec for tunnel group 1 on SeGW-3

In this section, the following steps are described:

- Create a multi-chassis peer using the system address of SeGW-4.
- Enable MCS for IPSec and tunnel group 1.
- Enable MC-IPSec for the tunnel group with a configured priority 200.
- Bind a central BFD session to MC-IPSec from the system interface.

Multi-chassis peer 192.0.2.4 is configured and MCS and MC-IPSec are enabled for tunnel group 1:

```
# on SeGW-3:
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.4 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 200
                    }
                }
            }
```

BFD is enabled for MC-IPSec in the preceding configuration. BFD is configured on the system interface 192.0.2.3:

```
configure {
    router "Base" {
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.3
                    prefix-length 32
                }
            }
        }
```

Configure MC-IPSec-aware routing on SeGW-3

In this step, a route policy is defined and applied to VPRN "VPRN-2".

Route policy "IPsec-to-MPBGP" exports static route 192.168.1.1/32 in VPRN "VPRN-2" to P-5. This policy sets the local preference of the prefix 192.168.1.1/32 according to the MC-IPSec state:

- for the **ipsec-master-with-peer** state: local preference 200
- for the **ipsec-non-master** state: local preference 100
- for the **ipsec-master-without-peer** state: local preference 200

The state **ipsec-master-without-peer** can be used to attract traffic to the designated primary chassis in case of "dual master" (meaning two chassis lose the MIMP connection in the base routing instance). In this example, SeGW-3 has local preference 200 and SeGW-4 has local preference 100 for **ipsec-master-without-peer**.

The route policy is configured as follows:

```
# on SeGW-3:
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
```

```
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-3 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.4" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
```

The route policy is applied as **vrf-export** in VPRN "VPRN-2":

```
configure {
    service {
        vprn "VPRN-2" {
            bgp-ipvpn {
                mpls {
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
```

Configure MC-IPSec-aware VRRP on SeGW-3

In this section, a VRRP policy is defined that uses the **mc-ipsec-non-forwarding** priority event to lower the in-use VRRP priority upon MC-IPSec switchover, which ensures VRRP and MC-IPSec have the same primary chassis. The VRRP instance needs to be in preempt mode.

This VRRP policy is only configured on the designated VRRP primary chassis SeGW-3, not on the standby chassis. The VRRP policy is applied to the interface "int-SeGW3-S-2" of IES "IES-1".

VRRP policy 1 is configured as follows:

```
# on SeGW-3:
configure {
    vrrp {
        policy 1 {
            priority-event {
                mc-ipsec-non-forwarding 1 {
                    priority {
                        priority-level 50
```

```
                        event-type explicit
                    }
                }
            }
        }
```

VRRP policy 1 is applied in VRRP instance 10 in the IES service:

```
configure {
    service {
        ies "IES-1" {
            interface "int-SeGW-3-S-2" {
                sap 1/1/2:1 {
                    description "SAP to switch S-2"
                }
                ipv4 {
                    primary {
                        address 172.16.1.252
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        priority 200
                        ping-reply true
                        policy 1
                    }
                }
            }
---snip---
```

Configure SeGW-4

The configuration on the standby chassis SeGW-4 is similar, but with different priorities and without the VRRP policy.

The tunnel group is configured in multi-active mode:

```
# on SeGW-4:
configure {
    isa {
        tunnel-group 1 {
            admin-state enable
            isa-scale-mode tunnel-limit-2k
            ipsec-responder-only true
            multi-active {
                isa 1/2 { }
            }
        }
```

The MCS and MC-IPSec configuration is as follows:

```
configure {
    redundancy {
        multi-chassis {
            peer 192.0.2.3 {
                admin-state enable
                sync {
                    admin-state enable
                    ipsec true
                    tunnel-group 1 {
```

```
                        sync-tag "tag-1"
                    }
                }
                mc-ipsec {
                    bfd-liveness true
                    tunnel-group 1 {
                        admin-state enable
                        peer-group 1
                        priority 150
                    }
                }
            }
        }
```

The base router configuration on SeGW-4 includes the following router interfaces and a static route to CE-1. OSPF is used as IGP between SeGW-3, SeGW-4, and P-5.

```
configure {
    router "Base" {
        interface "int-SeGW-4-P-5" {
            port 1/1/2:1000
            ipv4 {
                primary {
                    address 192.168.45.1
                    prefix-length 30
                }
            }
        }
        interface "int-SeGW-4-SeGW-3" {
            port 1/1/3:1000
            ipv4 {
                primary {
                    address 192.168.34.2
                    prefix-length 30
                }
            }
        }
        interface "system" {
            ipv4 {
                bfd {
                    admin-state enable
                }
                primary {
                    address 192.0.2.4
                    prefix-length 32
                }
            }
        }
        static-routes {
            route 10.10.10.0/24 route-type unicast {
                next-hop "172.16.1.100" {
                    admin-state enable
                }
            }
        }
        ospf 0 {
            admin-state enable
            area 0.0.0.0 {
                interface "int-SeGW-4-P-5" {
                }
                interface "int-SeGW-4-SeGW-3" {
                }
                interface "system" {
```

```
                }
            }
        }
```

The IPSec configuration is as follows:

```
configure {
    ipsec {
        ike-policy 1 {
            ipsec-lifetime 7200
            ike-transform [1]
            ike-version-2 {
            }
        }
        ike-transform 1 {
            isakmp-lifetime 172800
        }
        ipsec-transform 1 {
        }
```

The following route policy is configured on SeGW-4, The local preference is lower for the **ipsec-master-without-peer** state.

```
configure {
    policy-options {
        community "vprn2" {
            member "target:64496:2" { }
        }
        prefix-list "CE-1-Internal" {
            prefix 192.168.1.1/32 type exact {
            }
        }
        policy-statement "IPsec-to-MPBGP" {
            entry 10 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-master-with-peer
                }
                action {
                    action-type accept
                    local-preference 200
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 20 {
                from {
                    prefix-list ["CE-1-Internal"]
                    state ipsec-non-master
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            entry 30 {
                from {
                    prefix-list ["CE-1-Internal"]
```

```
                    state ipsec-master-without-peer
                }
                action {
                    action-type accept
                    local-preference 100
                    community {
                        add ["vprn2"]
                    }
                }
            }
            default-action {
                action-type accept
                community {
                    add ["vprn2"]
                }
            }
        }
```

The BGP configuration on SeGW-4 is as follows:

```
configure {
    router "Base" {
        autonomous-system 64496
        bgp {
            group "MPBGP" {
                type internal
                family {
                    vpn-ipv4 true
                }
            }
            neighbor "192.0.2.3" {
                group "MPBGP"
            }
            neighbor "192.0.2.5" {
                group "MPBGP"
            }
        }
```

The following GRE SDPs are configured:

```
configure {
    service {
        sdp 43 {
            admin-state enable
            description "GRE SDP toward SeGW-3"
            signaling off
            far-end {
                ip-address 192.0.2.3
            }
        }
        sdp 45 {
            admin-state enable
            description "GRE SDP toward P-5"
            signaling off
            far-end {
                ip-address 192.0.2.5
            }
        }
```

The public IES service is configured as follows:

```
configure {
```

```
    service {
        ies "IES-1" {
            admin-state enable
            service-id 1
            customer "1"
            interface "int-IPsec-Public-1" {
                static-tunnel-redundant-nexthop 192.168.34.1
                sap tunnel-1.public:1 {
                }
                ipv4 {
                    primary {
                        address 10.10.20.254
                        prefix-length 24
                    }
                }
            }
            interface "int-SeGW-4-S-2" {
                sap 1/1/1:1 {
                }
                ipv4 {
                    primary {
                        address 172.16.1.253
                        prefix-length 24
                    }
                    vrrp 10 {
                        backup [172.16.1.254]
                        ping-reply true
                    }
                }
            }
        }
```

The private VPRN service is configured as follows:

```
configure {
    service {
        vprn "VPRN-2" {
            admin-state enable
            service-id 2
            customer "1"
            ipsec {
                security-policy 1 {
                    entry 10 {
                        local-ip {
                            address 192.168.1.5/32
                        }
                        remote-ip {
                            address 192.168.1.1/32
                        }
                    }
                }
            }
            bgp-ipvpn {
                mpls {
                    admin-state enable
                    route-distinguisher "64496:2"
                    vrf-target {
                        community "target:64496:2"
                    }
                    vrf-export {
                        policy ["IPsec-to-MPBGP"]
                    }
                }
```

```
            }
            interface "int-IPsec-Private-1" {
                tunnel true
                static-tunnel-redundant-nexthop 192.168.20.1
                sap tunnel-1.private:1 {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                        key-exchange {
                            dynamic {
                                ike-policy 1
                                ipsec-transform [1]
                                pre-shared-key "pass"
                            }
                        }
                        tunnel-endpoint {
                            local-gateway-address 10.10.20.1
                            remote-ip-address 10.10.10.1
                            delivery-service "IES-1"
                        }
                        security-policy {
                            id 1
                        }
                    }
                }
            }
            interface "int-Redundant-1" {
                ipv4 {
                    primary {
                        address 192.168.20.2
                        prefix-length 30
                    }
                }
                spoke-sdp 43:20 {
                    ingress {
                        vc-label 2048
                    }
                    egress {
                        vc-label 2049
                    }
                }
            }
            spoke-sdp 43:2 {
                description "SDP to SeGW-3"
            }
            spoke-sdp 45:2 {
                description "SDP to P-5"
            }
            static-routes {
                route 192.168.1.1/32 route-type unicast {
                    ipsec-tunnel "tunnel-1" {
                        admin-state enable
                    }
                }
            }
        }
```

Verification

The following will be verified in this section:

- the MC-IPSec status and VRRP status on SeGW-3 and SeGW-4

- the status of the IPSec tunnel on CE-1
- the status of the IPSec tunnel on the SeGWs

Verify the MC-IPSec status on SeGW-3 and SeGW-4

The following is verified:

- SeGW-3 is the primary chassis (**master**) and SeGW-4 is the standby for tunnel group 1 because SeGW-3 has the higher priority 200.
- SeGW-3 is the primary node for VRRP instance 10 and SeGW-4 is the backup.

SeGW-3 is the primary chassis in tunnel group 1 with priority 200:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is the standby chassis in tunnel group 1 with priority 150:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
```

```
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is backup for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Verify the IPSec tunnel on CE-1

The following is verified in this section:

- the connectivity between CE-1 and P-5
- the IPSec tunnel information

A ping command is launched from the loopback interface in VPRN "VPRN-2" on CE-1 to the loopback interface in VPRN "VPRN-2" on P-5:

```
[/]
A:admin@CE-1# ping 192.168.1.5 router-instance "VPRN-2"
PING 192.168.1.5 56 data bytes
64 bytes from 192.168.1.5: icmp_seq=1 ttl=63 time=2.44ms.
64 bytes from 192.168.1.5: icmp_seq=2 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=3 ttl=63 time=2.38ms.
64 bytes from 192.168.1.5: icmp_seq=4 ttl=63 time=2.51ms.
64 bytes from 192.168.1.5: icmp_seq=5 ttl=63 time=2.50ms.
```

```
---- 192.168.1.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 2.38ms, avg = 2.44ms, max = 2.51ms, stddev = 0.053ms
```

The following command shows the IPSec tunnel information.

```
[/]
A:admin@CE-1# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.10.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.20.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

Verify the IPSec tunnel on the SeGWs

In this section, the following is verified:

- the MCS database is in-sync, so the tunnel status is up on both chassis.
- P-5 receives two VPN-IPv4 routes for prefix 192.168.1.1/32: the route from SeGW-3 has local preference 200; the route from SeGW-4 has local preference 100.

On both SeGWs, the IPSec tunnel with local address 10.10.20.1 and remote address 10.10.10.1 is up:

```
[/]
A:admin@SeGW-3# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
tunnel-1                         10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.10.1        IES-1        Up     1
--------------------------------------------------------------------------------
IPsec Tunnels: 1
================================================================================
```

```
[/]
A:admin@SeGW-4# show ipsec tunnel

================================================================================
IPsec Tunnels
================================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
--------------------------------------------------------------------------------
```

```
tunnel-1                          10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1                10.10.10.1        IES-1        Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

MCS is in sync on both SeGWs:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.4
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.3
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.3
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.4
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
```

```
Sub-mgmt options        :
  DHCP lease threshold  : Inactive
    Local / Remote      : -- / --
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
===============================================================================
```

The following command shows that P-5 received two VPN-IPv4 routes for prefix 192.168.1.1/32: one from SeGW-3 with local preference 200 and one from SeGW-4 with local preference 100:

```
[/]
A:admin@P-5# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.5        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     IGP Cost
      As-Path                                                       Label
-------------------------------------------------------------------------------
u*>i  64496:2:192.168.1.1/32                            200         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*i    64496:2:192.168.1.1/32                            100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.0/30                           100         None
      192.0.2.3                                         None        10
      No As-Path                                                    524286
*>i   64496:2:192.168.20.0/30                           100         None
      192.0.2.4                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.1/32                           100         0
      192.0.2.3                                         None        10
      No As-Path                                                    524286
u*>i  64496:2:192.168.20.2/32                           100         0
      192.0.2.4                                         None        10
      No As-Path                                                    524286
-------------------------------------------------------------------------------
```

```
Routes : 6
===============================================================================
```

MC-IPSec failover scenarios

Two MC-IPSec failover scenarios are described in this section:

- MC-IPSec failover when MS-ISA is disabled
- MC-IPSec failover when the primary chassis SeGW-3 reboots

Failover when MS-ISA is disabled

Initially, MS-ISA is enabled, so SeGW-3 is the primary chassis and SeGW-4 is the standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Master       200       1
                                 IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3
```

```
===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Backup       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The following command disables the MS-ISA on the primary chassis SeGW-3, which will trigger an MC-IPSec failover.

```
configure {
    card 1 {
        mda 2 {
            admin-state disable
```

With MS-ISA disabled, the MC-IPSec state of tunnel group 1 on SeGW-3 becomes **notEligible**, which means that the tunnel group is down, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for details description of MIMP states.:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10
```

```
===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             notEligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-3 is backup for VRRP instance 10 with in-use priority 50, as per the VRRP policy 1:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is now the primary chassis in tunnel group 1. This is triggered by MC-IPSec failover, as per the **mc-ipsec-non-forwarding** event in VRRP policy 1.

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

SeGW-4 is primary for VRRP instance 10;

```
[/]
A:admin@SeGW-4# show router vrrp instance
```

```
===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The situation is restored by enabling MS-ISA on SeGW-3:

```
configure {
    card 1 {
        mda 2 {
            admin-state enable
```

MC-IPSec failover when primary chassis reboots

The following **tools** command on SeGW-3 triggers an MC-IPSec switchover:

```
tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1

[/]
A:admin@SeGW-3# tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1
WARNING! Forcing a mastership switchover may significantly impact traffic. Are you sure (y/n)?
 y
```

Before the failure condition takes place, SeGW-3 is the primary chassis for tunnel group 1:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:09:10

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```

SeGW-3 is primary for VRRP instance 10:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                  10    No  Up   Master       200       1
                                IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is the standby chassis for tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===========================================================================
ID             Peer Group     Priority  Admin State    Mastership
---------------------------------------------------------------------------
1              1              150       Up             standby
---------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===========================================================================
===============================================================================
```

The VRRP state on SeGW-4 is backup:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

The following command reboots the primary chassis SeGW-3:

```
[/]
A:admin@SeGW-3# admin reboot card active now
```

While SeGW-3 reboots, the IPSec state of SeGW-4 becomes **eligible**:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-4 is primary (**master**):

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

When SeGW-3 comes up, the IPSec state of tunnel group 1 is **discovery**, which means that the system has not established the MIMP session with its peer yet.

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
```

```
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             discovery
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

After a while, the preceding **show** command is repeated and the IPSec state for tunnel 1 on SeGW-3 is standby:

```
[/]
A:admin@SeGW-3# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.4

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.4
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:24:41

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
```

The VRRP state on SeGW-3 is backup:

```
[/]
A:admin@SeGW-3# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-3-S-2                   10    No  Up   Backup       200       1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

SeGW-4 is the primary chassis in MC-IPSec tunnel group 1:

```
[/]
A:admin@SeGW-4# show redundancy multi-chassis mc-ipsec peer ip-address 192.0.2.3

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.3
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 02/16/2023 10:10:22

===================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------
1              1              150       Up             master
-------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===================================================================
===============================================================================
```

SeGW-4 is the primary node for VRRP instance 10:

```
[/]
A:admin@SeGW-4# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-4-S-2                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

Configuration guidelines

The following is a list of guidelines for configuring MC-IPSec:

- To avoid high CPU load and issues in some complex cases, the following are suggestions for configuring the IKEv2 lifetime:
  - Both IKE_SA and CHILD_SA lifetime on MC-IPSec chassis (SeGW-3 and SeGW-4) should be around three times larger than on the IPSec peer CE-1.
  - With the first rule, the lifetime of the side with smaller lifetime (IPSec peer CE-1) should not be too small (these being the default values):
    - IKE_SA: >= 86400 seconds
    - CHILD_SA: >= 3600 seconds

  - With the first rule, on the side with smaller lifetime (IPSec peer CE-1), the IKE_SA lifetime must be at least 3 times larger than CHILD_SA lifetime.
- The IKE protocol is the control plane of IPSec, so IKE packets must be treated as high QoS priority in the end-to-end path of the public service. On the public interface, a SAP ingress QoS policy must be configured to ensure that IKE packets get high QoS priority.
- Configure **ipsec-responder-only true** under **tunnel-group** for static LAN-to-LAN tunnels.
- Enable dead peer detection (DPD) on the IPSec peer side (CE-1); disable DPD (default) on the MC-IPSec chassis side.
- The direct and redundant physical link between MC-IPSec chassis must be configured with sufficient bandwidth for MCS and shunting traffic, and proper QoS configuration to make sure the MIMP and MCS packets are treated as high priority traffic.
- The system time must be same on both MC-IPSec chassis.
- Make sure the protection status is **nominal** on both chassis before provoking a controlled switchover. The protection status can be displayed with the **show redundancy multi-chassis mc-ipsec peer ip-address <addr>** command.
- Wait at least five minutes between two consecutive switchovers if possible, to prevent a second switchover happening before the standby is ready to become the primary chassis.

Conclusion

MC-IPSec provides a stateful multi-chassis IPSec redundancy solution. This is very important in a carrier grade network, especially in applications such as mobile backhaul where high value mobile services run over IPSec tunnels.

N:M MC-IPsec Redundancy

This chapter describes N:M MC-IPsec redundancy.

Topics in this chapter include:

Applicability

The information and MD-CLI configuration in this chapter are based on SR OS Release 22.10.R1.

The IPsec tunnel termination configuration described in this chapter requires an MS-ISA2 or an ESA server configured with a virtual machine. Configuration and setup for ISA2 or ESA are beyond the scope of this chapter; see the Multi-Chassis IPSec Redundancy chapter.

Overview

The N:M MC-IPsec redundancy model is a feature of the multi-chassis (MC) capabilities of SR OS when the router is deployed as Security Gateway (SeGW). N:M aims at enhancing the existing 1:1 redundancy model for IPsec tunnels. For the definition of N:M terminology and a description of its benefits, see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide*.

The figure Figure 3: 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 3: Three-node redundancy domain with a 2 DA + 1 DS model

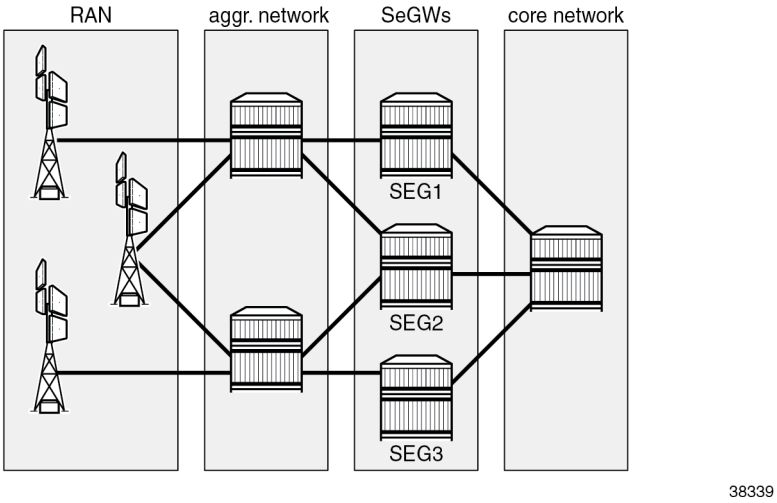

Radio access network (RAN) elements are opening IPsec tunnels toward SeGW cluster tunnel endpoint IP addresses. The RAN, aggregation network, and core network are emulated with standard routing nodes. For this deployment, assume that connectivity between elements is established using routing protocols and, as for a classic SeGW router, the public side where traffic is encrypted is built on top of a public-side VPRN, while private side (clear-text traffic) is associated with another VPRN. ISA2 or ESA resources manage encryption and decryption operations across the VPRN boundary.

This chapter describes configuration of SeGW elements, as well as MD-CLI commands for tracking the functionality of N:M nodes in the same redundancy domain (RD).

Configuration

Assume that IP connectivity is established across the IP network elements in the architecture. It is beyond the scope of this chapter to describe how traffic is carried from the RAN to the SeGW or from SeGW to the mobile packet core. Among the protocols and techniques that are required to speed up convergence of routing, the bidirectional forwarding detection (BFD) protocol is especially useful to keep network convergence time in a range compatible with mobile traffic use case.

ISA2 or ESA setup for N:M

The nodes participating in the IPsec domain have a standard setup for ISA2 or ESA resources.

SEG1 and SEG 2 can each be configured like a classic SeGW, as follows:

```
[gl:/configure isa]
A:admin@SEG1# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            isa 1/2 { }
```

```
            active-isa-number 1
        }
        reassembly {
            max-wait-time 1200
        }
        stats-collection {
            isa-dp-cpu-usage true
        }
    }
```

The **active-isa-number** command specifies the number of active encryption and decryption elements. Nokia recommends implementing the same number of ISA2 and ESA resources among the nodes participating in the RD, which allows for the DS node to activate the same number of ISA2 or ESA resources when failover occurs. However, a failover can occur even if the DS node has a lower number of ISA2 or ESA resources available in its local pool. This allows operators to save costs, but if the ISA2 ore ESA resources on the initial DA nodes were fully loaded, the DS node cannot host all tunnels and the protection is only partial.

N:M redundancy allows DS nodes to cover multiple TGs, and therefore, multiple RDs. DS nodes may have more ISA2 or ESA resources than the DA nodes, because the DS nodes should be able to cover one or more DA node failures, with a maximum of 16.

The output from SEG2 is the same as for SEG1.

SEG3 is configured as the DS node of the domain, where the configuration contains the **tunnel-member-pool** command:

```
[gl:/configure isa]
A:admin@SEG3# info
    tunnel-group 1 {
        admin-state enable
        isa-scale-mode tunnel-limit-2k
        ipsec-responder-only true
        multi-active {
            member-pool "MP1"
        }
        reassembly {
            max-wait-time 1200
        }
    }
    tunnel-member-pool "MP1" {
        isa 1/2 { }
    }
```

The **tunnel-member-pool** option defines the set of ISA2 or ESA resources used by the DS node during failures on active nodes. It is referenced in the tunnel group (TG) configuration, because multiple TGs can use the same tunnel member pool using the same set of ISA2 or ESA resources.

The output of the **show isa tunnel-member-pool** command lists ISA (ISA2 or ESA) members and their states. Under normal conditions, the ISA2 or ESA resource is not active on SEG3.

```
[gl:/configure isa]
A:admin@SEG3# /show isa tunnel-member-pool "MP1" detail

===============================================================================
ISA Tunnel Member Pool : MP1
Description            : (Not Specified)
Associated Tunnel Grps : 1
===============================================================================
Isa Members                   Active In Group     Last Configuration Change
-------------------------------------------------------------------------------
```

```
1/2                                                    11/25/2022 12:10:14
-------------------------------------------------------------------------------
Number of Configured Entries: 1
Number of Active Entries: 0
===============================================================================
```

Redundancy domain configuration

The configuration of MC-IPsec as N:M starts by defining node roles and behavior. The configuration on SEG1 (with system IP address 192.0.2.1) is as follows:

```
[gl:/configure redundancy]
A:admin@SEG1# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
            priority 250
            tunnel-group 1
        }
        peer 192.0.2.2  {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The preceding configuration example shows a multi-chassis IPsec domain, where the following domain characteristics have been specified:

- domain number – must be shared across all the nodes joining the redundancy domain (RD)
- designated role – DA or DS
- priority – required by the multi-chassis IPsec mastership protocol (MIMPv2) when an operationally active (OA) node must be elected. Setting a higher priority for an SeGW increases the likelihood of it being elected as the OA. In this case, SEG1 has the highest priority and DA role, so it is elected OA for RD 1.

- tunnel group – must be defined as per the ISA2 or ESA setup. The TG is always mapped to the RD in a 1:1 relationship
- peers – up to three peers can be added. While full-mesh peering between them is required, Nokia also recommends deploying highly redundant network paths between these peers.

  Each peer has its own CLI tree where the following characteristics must be defined:

  – the domain or domains the peer belongs to

  – the synchronization state for IPsec

  – whether BFD is applied to check peer liveliness.
- (optional) other parameters for keepalives, hold-time, and discovery-interval are configured with default values. Do not change these values unless a different setup is required under specific network conditions.

The configuration for the redundancy domain on SEG2 is the same as on SEG1, but with different IP addresses for peers and different priority:

```
A:admin@SEG2# info
    multi-chassis {
        ipsec-domain 1 {
            admin-state enable
            designated-role active
            priority 240
            tunnel-group 1
        }
        peer 192.0.2.1 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
        peer 192.0.2.3 {
            admin-state enable
            sync {
                admin-state enable
                ipsec true
            }
            mc-ipsec {
                bfd-liveness true
                domain 1 {
                    admin-state enable
                }
            }
        }
    }
```

The designated role of SEG2 is **active**, which means SEG2 behaves similarly to the 1:1 model where tunnel states are synchronized with SEG1 and immediately pushed to ISA2 or ESA resources. This behavior allows for a very quick failover when SEG1 experiences a failure.

The priority is 240, which is lower than for SEG1. As a result, SEG1 receives node role DA and is operationally active (OA) while SEG2 receives node role DA and is operationally standby (OS).

The RD configuration for DS SEG3 is as follows:

```
[gl:/configure redundancy multi-chassis]
A:admin@SEG3# info
    ipsec-domain 1 {
        admin-state enable
        designated-role standby
        priority 230
        tunnel-group 1
    }
    peer 192.0.2.1 {
        admin-state enable
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
    peer 192.0.2.2 {
        admin-state enable
        sync {
            admin-state enable
            ipsec true
        }
        mc-ipsec {
            bfd-liveness true
            domain 1 {
                admin-state enable
            }
        }
    }
```

The peer configuration is similar to those of other nodes where BFD liveliness is enabled.

The designated role is standby (DS). The default value in the configuration is not shown from the **info** command.

The priority is 230 but the node role is DS. The DS node will not become OA because the DA role of SEG1 and SEG2 always prevails when electing the OA, regardless of priority value. Therefore, a DS node can become OA only if there are no DA nodes available in the domain.

After the setup of MC IPsec RD is completed across all the nodes, **show** commands can be used to track RD behavior and state:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 250
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status  : nominal
Router Id          : 192.0.2.1         Current Active     : 192.0.2.1
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
```

```
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                                 Up      standby      active
  192.0.2.2
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

The output shows important information about the redundancy domain:

- the designated role of the node – active or standby
- the activity state based on fault conditions – active or standby
- the protection status – "nominal" means that the nodes are synchronized.
- the domain adjacencies – list of peers and their activity state and designated role
- the tunnel statistics – in this case, seven dynamic tunnels are established

The same **show** command executed on SEG2 provides similar output, with differences for the priority and the designated role. The seven tunnels are shown in the "Installed" state because SEG2 is a DA node.

The same **show** command on DS SEG3 shows the following:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby           Priority          : 230
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : nominal
Router Id          : 192.0.2.3         Current Active    : 192.0.2.1
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                                 Up      active       active
  192.0.2.1
192.0.2.2                                 Up      standby      active
  192.0.2.2
```

```
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                              Static              Dynamic
-------------------------------------------------------------------------------
Installed                     0                   0
Installing                    0                   0
Standby Dormant               0                   7
Awaiting Config               0                   0
Failed                        0                   0
===============================================================================
```

Relevant information from the SEG3 CLI output, apart from the activity state, the designated role, and the peer's state, is the tunnel state, which is now marked as "Standby Dormant".

Tunnels on SEG3 are not installed on the ISA2 or ESA; rather, they are stored in the router CPM and are kept ready to be offloaded on the ISA2 or ESA resources connected to the router. These tunnels are offloaded as soon as SEG3 becomes OA, following a node reboot, failure, or manual switchover.

Services configuration

The tunnels opened by RAN elements are terminated in a public-side VPRN IP address called TEIP (the public side can also be made on a IES service). Assume that the RAN nodes are using a single tunnel setup with a single IKE_SA, whereas the Child_SA's number is specific to the deployment. The configuration of this public side VPRN is the same for all three nodes and follows the standard SeGW setup:

```
[gl:/configure service vprn "100"]
A:admin@SEG1# info
    vprn "100" {
        admin-state enable
        description "public side"
        customer "1"
        ipsec {
            multi-chassis-shunt-interface "to_SEG2_Shunt" {
                next-hop {
                    address 10.1.12.2
                }
            }
            multi-chassis-shunt-interface "to_SEG3_Shunt" {
                next-hop {
                    address 10.1.13.2
                }
            }
            multi-chassis-shunting-profile "MCSPROF1" {
                peer 192.0.2.2 {
                    multi-chassis-shunt-interface "to_SEG2_Shunt"
                }
                peer 192.0.2.3 {
                    multi-chassis-shunt-interface "to_SEG3_Shunt"
                }
            }
        }
        interface "PUBLIC1" {
            multi-chassis-shunting-profile "MCSPROF1"
            sap tunnel-1.public:100 {
```

```
                ipsec-gateway "IPSECGW1" {
                    admin-state enable
                    default-tunnel-template 1
                    ike-policy 1
                    pre-shared-key "uCLxzS3PxoW0foPjmAKJ/Wv41hy6O3H76tg=" hash2
                    default-secure-service {
                        service-name "200"
                        interface "PRIVATE1"
                    }
                    local {
                        gateway-address 10.51.100.1
                    }
                }
            }
            ipv4 {
                primary {
                    address 198.51.100.2
                    prefix-length 24
                }
            }
        }
        interface "to_SEG2_Shunt" {
            spoke-sdp 2000:1 {
            }
            ipv4 {
                primary {
                    address 10.1.12.1
                    prefix-length 30
                }
            }
        }
        interface "to_SEG3_Shunt" {
            spoke-sdp 3000:1 {
            }
            ipv4 {
                primary {
                    address 10.1.13.1
                    prefix-length 30
                }
            }
        }
        ospf 0 {
            export-policy ["EXPORT_OSPF"]
        }
```

The parts of the configuration that are exclusive of N:M are those related to shunt-link setup.

The **multi-chassis-shunting-profile** command can be found under the **ipsec** configuration for the IES or VPRN service, where the multi-chassis shunting (MCS) profile is required to map each peer to a dedicated shunt interface. The MCS profile is referenced under the interface where the IPsec gateway is configured. In this scenario, peer 192.0.2.2 is reached through the to_SEG2_Shunt interface, which is defined under the same VPRN as an interface built on top of sdp:2000:1.

A full mesh of shunt interfaces is made across the RD, for both public and private side services.

```
A:admin@SEG1# show ipsec multi-chassis-shunt-interface service "100"

===============================================================================
IPsec Multi-Chassis Shunt Interfaces
===============================================================================
Service Id  MC Shunt Interface Name          Next Hop                Resolved
-------------------------------------------------------------------------------
100         to_SEG2_Shunt                    10.1.12.2               Yes
```

```
100         to_SEG3_Shunt                   10.1.13.2               Yes
-------------------------------------------------------------------------------
No. of IPsec MC Shunt Interfaces: 2
===============================================================================
```

The **show ipsec multi-chassis-shunt-interface service** command shows the liveliness of shunt interfaces and information on the next-hop resolution, whereas the **show ipsec multi-chassis-shunting-profile service** command provides a summary of the MCS profile and associated peers:

```
A:admin@SEG1# show ipsec multi-chassis-shunting-profile service "100"

===============================================================================
Multi-Chassis Shunting Profile Params Entries
===============================================================================
Service Id  MC Shunting Profile Name                MC Shunt Interface Name
            Peer
-------------------------------------------------------------------------------
100         MCSPROF1                                to_SEG2_Shunt
            192.0.2.2
100         MCSPROF1                                to_SEG3_Shunt
            192.0.2.3
-------------------------------------------------------------------------------
No. of IPsec MC Shunting Profile Params Entries: 2
===============================================================================
```

The SDP full mesh must be configured on both sides, as shown in the figure Figure 4: SDP full mesh.

Figure 4: SDP full mesh

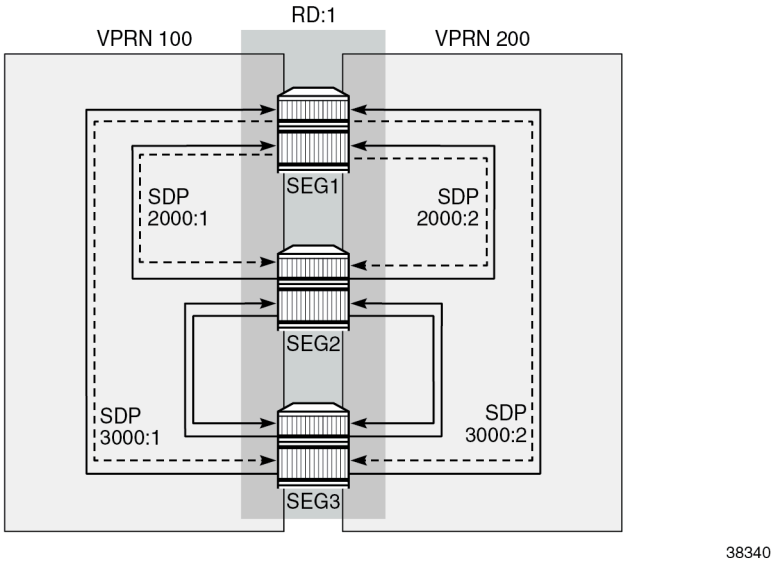

Note: Only the SDPs from SEG1 are shown with IDs.

The shunt link can be built from a standard spoke SDP or from a port-based interface. In this example, the following spoke SDPs are used in the public-side VPRN 100:

```
A:admin@SEG1# show service id "100" sdp
```

```
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:1           Spok     192.0.2.2       Up      Up        524285    524285
3000:1           Spok     192.0.2.3       Up      Up        524283    524285
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

The **show** output for the private-side VPRN 200 looks similar to that for the public-side VPRN, except for the SDP IDs and label values:

```
A:admin@SEG1# show service id "200" sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId            Type     Far End addr    Adm     Opr       I.Lbl     E.Lbl
-------------------------------------------------------------------------------
2000:2           Spok     192.0.2.2       Up      Up        524284    524284
3000:2           Spok     192.0.2.3       Up      Up        524282    524284
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
```

There are no routing policy changes from the 1:1 MC-IPsec cluster, although this example could have a more complex routing setup, considering that the number of routers in a domain is higher than in the 1:1 model. The following configuration shows the SEG1-2-3 export policy used on the public side where the OSPF protocol is used under VPRN 100:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG1# info
    description "EXPORT TEIP OSPF - PUBLIC SIDE"
    entry 10 {
        from {
            state ipsec-master-with-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 30
            }
        }
    }
    entry 20 {
        from {
            state ipsec-non-master
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
```

```
            tag 100
            metric {
                set 190
            }
        }
    }
    entry 30 {
        from {
            state ipsec-master-without-peer
            protocol {
                name [ipsec]
            }
        }
        action {
            action-type accept
            tag 100
            metric {
                set 40
            }
        }
    }
    default-action {
        action-type reject
    }
```

On SEG2, only the metrics are different and are aligned with DA priorities:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG2# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 60
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 190
                }
            }
        }
        entry 30 {
            from {
```

```
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 200
                metric {
                    set 50
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

On SEG3, the export policy is as follows:

```
[gl:/configure policy-options policy-statement "EXPORT_OSPF"]
A:admin@SEG3# info
    policy-statement "EXPORT_OSPF" {
        description "EXPORT TEIP OSPF - PUBLIC SIDE"
        entry 10 {
            from {
                state ipsec-master-with-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 90
                }
            }
        }
        entry 20 {
            from {
                state ipsec-non-master
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
                tag 300
                metric {
                    set 195
                }
            }
        }
        entry 30 {
            from {
                state ipsec-master-without-peer
                protocol {
                    name [ipsec]
                }
            }
            action {
                action-type accept
```

```
                tag 300
                metric {
                    set 60
                }
            }
        }
        default-action {
            action-type reject
        }
    }
```

The export policy on the private-side VPRN is made with the same concept as the public side, but is not shown here.

Note: Parts of the configuration where the parameters remain the same as those in classic SeGW deployments (either stand-alone or 1:1) have not been added to this chapter. This information is described in the Multi-Chassis IPSec Redundancy chapter.

On the private side of SeGWs, a different VPRN is required, as per standard IPsec configuration. The private-side VPRN configuration on SEG1 is as follows:

```
[gl:/configure service vprn "200"]
A:admin@SEG1# info
    admin-state enable
    description "private segw testing"
    customer "1"
    ipsec {
        multi-chassis-shunt-interface "to_SEG2_Shunt" {
            next-hop {
                address 10.2.12.2
            }
        }
        multi-chassis-shunt-interface "to_SEG3_Shunt" {
            next-hop {
                address 10.2.13.2
            }
        }
        multi-chassis-shunting-profile "MCSPROF1" {
            peer 192.0.2.2 {
                multi-chassis-shunt-interface "to_SEG2_Shunt"
            }
            peer 192.0.2.3 {
                multi-chassis-shunt-interface "to_SEG3_Shunt"
            }
        }
    }
    bgp-ipvpn {
        mpls {
            admin-state enable
            route-distinguisher "300:4"
        }
    }
    interface "PRIVATE1" {
        tunnel true
        multi-chassis-shunting-profile "MCSPROF1"
        sap tunnel-1.private:100 {
        }
    }
    interface "to_SEG2_Shunt" {
        ipv4 {
            primary {
                address 10.2.12.1
```

```
                prefix-length 30
            }
        }
        spoke-sdp 2000:2 {
        }
    }
    interface "to_SEG3_Shunt" {
        ipv4 {
            primary {
                address 10.2.13.1
                prefix-length 30
            }
        }
        spoke-sdp 3000:2 {
        }
    }
```

As the configuration shows, the same setup of shunt links is required on the private side to allow path resiliency in case of faults for the traffic going downstream from core toward the RAN.

Failure scenario – active node experiences a power failure

N:M can be triggered by different fault conditions, such as a complete node failure, an ISA2 or ESA failure, or a manual switchover executed with the **tools** command. In this scenario, complete node failures are simulated. When SEG1 experiences a node failure, SEG2 takes over. When SEG2 fails too, SEG3 takes over and remains the only node with active tunnels.

The initial scenario has SEG1 and SEG2 configured as DA nodes, while SEG3 is the DS node for the domain configured as **ipsec-domain 1**. The state can be verified with the **show redundancy multi-chassis ipsec-domain 1** command (as shown above in the Redundancy domain configurationsection).

As soon as SEG1 experiences a node failure, SEG2 takes over:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 240
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.2         Current Active    : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote      Remote
  Router-Id                             State   Activity    Designated
                                                State       Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown     unknown
  0.0.0.0
192.0.2.3                               Up      standby     standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================
```

```
===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

Although the protection status, as seen from SEG2 and SEG3, is initially "notReady", it changes to "nominal" after few minutes. From the SEG2 and SEG3 point of view, SEG1 is unreachable, and its activity state remains unknown. Log 99 also records the failure event:

```
A:admin@SEG2# show log log-id 99

===============================================================================
Event Log 99 log-name 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=187  (not wrapped)]

186 2022/12/13 14:05:32.534 UTC WARNING: MC_REDUNDANCY #2047 Base MC-IPSEC-DOMAIN 1
"Protection status for the multi-chassis ipsec domain 1 changed to nominal"

185 2022/12/13 14:02:19.611 UTC MINOR: VRTR #2061 Base 192.0.2.1
"BFD: Local Discriminator 1 BFD session on node 192.0.2.1 is down due to noHeartBeat "

---snip---

179 2022/12/13 14:02:19.124 UTC WARNING: MC_REDUNDANCY #2004 Base
"The Sync status of peer 192.0.2.1 changed to outOfSync"

178 2022/12/13 14:02:18.746 UTC WARNING: MC_REDUNDANCY #2046 Base MC-IPSEC-DOMAIN 1
"Multi-chassis ipsec domain 1 local activity state changed from standby to active because an
 inter-chassis link went down. The active router in the domain is 192.0.2.2."
```

Next, SEG2 also experiences a full node failure, and SEG3 takes over:

```
A:admin@SEG3# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : standby           Priority          : 230
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.3         Current Active    : 192.0.2.3
Activity State     : eligible
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                               Down    unknown      unknown
  0.0.0.0
```

```
192.0.2.2                                   Down    unknown      unknown
  0.0.0.0
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                             Static              Dynamic
-------------------------------------------------------------------------------
Installed                    0                   7
Installing                   0                   0
Standby Dormant              0                   0
Awaiting Config              0                   0
Failed                       0                   0
===============================================================================
```

Both SEG1 and SEG2 are seen as operationally down with an unknown activity state. On SEG3, the tunnel states have been copied from the CPM to the ISA2 or ESA entities and are now shown as "Installed", rather than "Standby Dormant". As soon as SEG1 or SEG2 are back up, the **revertive** flag configured under the **ipsec-domain** command determines if the tunnels are kept on the current active DS node or if they are moved back to SEG1 ownership.

Failure scenario – using the tools command line

A planned failure condition is commonly seen when executing software upgrades or hardware maintenance on SeGW nodes, which leverages the **tools** command line utility to move tunnels toward other peering nodes.

The initial state is the same as for the previous example where SEG1 is initially the operationally active DA.

The following tools command triggers a switchover and therefore causes all the tunnels installed on the operationally active DA node to move on another node in the domain, selected by the **auto** flag in this case.

```
A:admin@SEG1# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 auto
 now
```

To specify a peer IP address among those available in the domain, the **to <peer_ip>** option could be used instead of **auto**.

The following output shows the domain state as seen from SEG1 after the execution of the tools command:

```
A:admin@SEG1# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active            Priority          : 250
Tunnel Group       : 1                 Revertive         : false
Admin State        : Up                Protection Status : notReady
Router Id          : 192.0.2.1         Current Active    : 192.0.2.2
Activity State     : standby
===============================================================================

===============================================================================
Domain 1 Adjacencies
```

```
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.2                               Up      active       active
  192.0.2.2
192.0.2.3                               Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
Installing                  0                   0
Standby Dormant             0                   0
Awaiting Config             0                   0
Failed                      0                   0
===============================================================================
```

As shown in the output, the current active node is SEG2 (192.0.2.2). The **auto** flag forced all the traffic to move across the second preferred active node in the domain, which is SEG2.

The protection status, as seen from SEG2, changes to "nominal" after a few minutes:

```
A:admin@SEG2# show redundancy multi-chassis ipsec-domain 1

===============================================================================
Multi-Chassis IPsec Domain: 1
===============================================================================
Designated Role    : active             Priority           : 240
Tunnel Group       : 1                  Revertive          : false
Admin State        : Up                 Protection Status  : nominal
Router Id          : 192.0.2.2          Current Active     : 192.0.2.2
Activity State     : active
===============================================================================

===============================================================================
Domain 1 Adjacencies
===============================================================================
Peer                                    Oper    Remote       Remote
  Router-Id                             State   Activity     Designated
                                                State        Role
-------------------------------------------------------------------------------
192.0.2.1                                 Up      standby      active
  192.0.2.1
192.0.2.3                                 Up      standby      standby
  192.0.2.3
-------------------------------------------------------------------------------
Domain Adjacency Entries found: 2
===============================================================================

===============================================================================
Multi-Chassis Tunnel Statistics
===============================================================================
                            Static              Dynamic
-------------------------------------------------------------------------------
Installed                   0                   7
```

```
Installing                    0                    0
Standby Dormant               0                    0
Awaiting Config               0                    0
Failed                        0                    0
===============================================================================
```

After maintenance operations on SEG1 have been completed and the node is operational (which can be verified using the **show** commands described in this chapter), the operator reverts services and traffic back to SEG1. For this purpose and in this specific example, the same **tools** command can be used. The **auto** flag selects SEG1, according to its highest priority in the domain. If more predictability is required in the selection choice, the **to <peer_ip>** flag can be used, as in this example:

```
A:admin@SEG2# tools perform redundancy multi-chassis mc-ipsec force-switchover domain 1 to
 192.0.2.1 now
```

Conclusion

N:M adds a level of redundancy to an already efficient redundancy model; it ensures that RAN elements stay connected to the core network under a wide range of failure conditions. SR OS uses a full set of commands to implement this feature, available for both classic and MD-CLI. N:M also gives network engineers and architects the capability to deploy SeGW services with greater flexibility; for example, to deploy super-resilient SeGW clusters to serve high-density RAN areas, or to introduce cost-optimized solutions with an acceptable level of automated fault recovery.

Customer document and product support

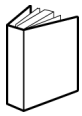

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback