Reverse Route Injection

This chapter describes the Reverse Route Injection (RRI) feature supported by the SecGW.

The following topics are covered:

Overview

RRI injects routes in the reverse direction onto the ASR 9000 VSM (IOS-XR blade) so that clear traffic can be routed to the correct interface on the target VSM. The OneP (ConnectedApps [CA]) library provides the necessary API calls to CA clients to communicate to the oneP server (running on IOS-XR).

The RRI feature is used in conjunction with the StarOS SecGW to deal with Site-to-Site (S2S) IPSec SAs and RAS; though the requirement is mainly for S2S. RRI route transaction is initiated is when a tunnel SA is being created.

Interchassis Session Recovery (ICSR) works with RRI to ensure that traffic is correctly routed following an HA switchover.

For additional information, see the sample configurations that appear at the end of this guide.

How It Works

The Connected Apps Linux Process (CALP) receives single or batched route insertion/deletion request, validates the message received is complete, and initiates the update of the route request. A route update API then injects the routes contained in the Routing Information Base (RIB) table of the ASR 9000 Route Processor (RP).

A re-inject (replay) is an asynchronous event message from the ASR 9000 RP asking the StarOS CA client to replay all the route entries in its database from scratch. This message is usually generated in a drastic failure case where the RP has lost all the previously injected RRI routes in its Forwarding Information Base (FIB) table.

Status Handler processes all incoming responses from CALP to batch requests. Each response has a batch_id which will be correlated to the corresponding batch request. Route entries that are not acknowledged are

regrouped and retransmitted. Those that are successful are moved to the route database hash table and removed from this batch. State diagram provided below shows the various states that a RRI route entry can be based on the responses for its batch request.

Figure 1: RRI Requests – State Diagram

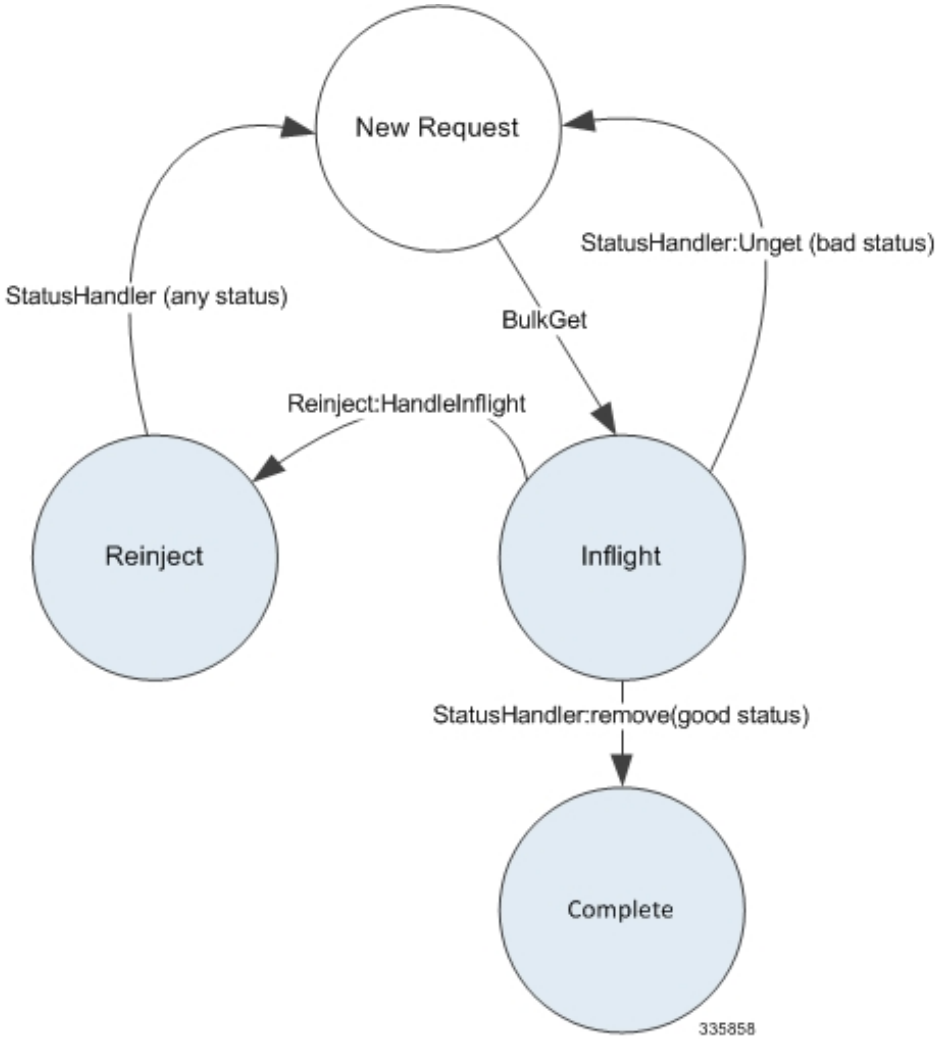

A StarOS proclet (cactrl) manages the creation and maintenance of the session with CALP. This session is the only communication channel between each StarOS VM and the ASR 9000 RSP. This oneP communication session must be established before any form of communication can occur between the two entities. See the *oneP Communication* chapter for detailed information.

High Availability for RRI

Interchassis Session Recovery (ICSR) is implemented for RRI to ensure that the routes are injected correctly on the appropriate VSM to route the traffic to the correct interface after an ICSR switchover.

ICSR can be implemented for:

- Intrachassis or cluster card-level redundancy

- Interchassis L2 card-level redundancy
- Interchassis L3 card-level redundancy

Important RRI is mandatory for S2S StarOS WSG service and optional for RAS.

Intrachassis/Cluster Redundancy

This mode only supports Layer 2, 1:1 redundancy between VPC-VSM instances (StarOS VMs) across two VSMs in the same ASR 9000 chassis. Both instances are located in the same chassis and, therefore, the routes injected by the active VPC-VSM instance to the IOS-XR will still be valid after the failure when the standby card takes over. In this case, the NPU Manager on the standby VSM does not inject the routes to the IOS-XR. The routes only need to be added to the Route DB.

The main requirements for ICSR in this mode are:

- The route DB on the standby VSM must contain only routes that have been successfully injected by the active VPC-VSM instance.
- To prevent IOS-XR from removing the routes, CALP on the standby StarOS VM reconnects to the CA server via the same session ID used prior to the timeout. The session ID is stored in the shared configuration task (SCT) of the CA Controller and a new micro-checkpoint is sent to the standby VPC-VSM instance.

The session manager which programs the IPSec manager and other sessions managers synchronizes the tunnels with the standby VPC-VSM instance via SRP.

Interchassis Redundancy

Overview

This mode supports hot standby redundancy between two VPC-VSM instances in different ASR 9000 chassis. The standby instance is ready to become active once a switchover is triggered. SA re-negotiation is not required and traffic loss is minimal.

The Interchassis Session Recovery (ICSR) model supports both Layer 2 and Layer 3 levels of redundancy. Basic ICSR requirements are:

- The route database on the standby VSM must contain only the routes that were successfully injected by the active VSM.
- L3-based HA SecGW deployment uses the onePK Routing Service Set (RSS) infrastructure to support geo-redundancy. It does this by inserting the necessary routes on the ASR 9000 RSP. The RSP then distributes the relevant routes outwardly such that external traffic would reach the active VSM instead of the standby VSM.
- For Layer 3 redundancy, the routes are injected via IOS-XR as two legs. Only the first leg of the routes is injected to IOS-XR running on the chassis with the standby VSM. The small set of secondary leg routes are reconfigured to point to the newly active VSM after the switchover.

For additional information on StarOS ICSR, see the *VPC-VSM System Administration Guide*.

Mapping of VPC-VSM Instances between VSMs

Because of the asymmetric assignment of VSM resources among StarOS VMs, an operator should configure one-to-one mapping between StarOS VMs across active/standby VSMs in different ASR 9000 chassis. See the table below.

Table 1: Recommended Mapping of Interchassis StarOS VMs

Active VSM	Standby VSM
VM1	VM1
VM2	VM2
VM3	VM3
VM4	VM4

Each VM will be monitored via separate HSRP configurations and connected to separate oneP (CA) sessions so that switchover of one VM will not affect the other VMs.

RRI Configuration Commands

There are several StarOS CLI commands associated with RRI configurations. They are briefly described below. For additional information, see the *Command Line Interface Reference*.

Important You must separately configure RRI on each StarOS VM (VPC-VSM instance).

ip/ipv6 rri Command

This Context Configuration mode CLI command configures Reverse Route Injection egress clear port IP parameters. This command is supported for both Remote Access Service and S2S configurations.

```
configure
      context  context_name
 { ip | ipv6 } rri { ip_address | next-hop nexthop_address } interface
interface_name [ vrf vrf_name ]
```

Notes:

- Use this command for standalone and Interchassis L2-ICSR.
- *ip_address* and *nexthop_address* can be specified in IPv4 dotted-decimal (**ip rri**) or IPv6 colon-separated-hexadecimal (**ipv6 rri**) format.
- The next hop IP address is the SecGW clear interface physical address.
- *interface_name* specifies the egress interface.It should be unique and map the *vrf_name*, under the same context.
- The *vrf_name* is the VRF of clear interface in ASR 9000/RSP (external interface as well as the VSM interface) wherein the clear traffic is forwarded based on the RRI route.

ip/ipv6 rri-route Command

This Context Configuration mode CLI command configures High Availability Routing Parameters for Reverse Route Injection.

```
configure
      context  context_name
{ ip | ipv6 } rri-route network-mode { L2 | L3 } { clear_loopback_ip | rri-ip
virtual_ip_address } { ip_address | next-hop nexthop_address } interface interface_name
[ vrf vrf_name ]
            end
```

Notes:

- Configuring Border Gateway Protocol (BGP) is required when this CLI is used, to support Interchassis L3-ICSR and Intrachassis ICSR. This CLI will add only 1st-leg route. The 2nd-leg routes are added using other routing protocols such as BGP, or OSPF, etc.
- This command is mandatory in the following scenarios:
  - L2 Intrachassis HA (where loopback IP is configured)
  - L3 Interchassis HA (where loopback IP is configured)
- *ip_address*, *virtual_ip_address* and *nexthop_address* can be specified in IPv4 dotted-decimal (**ip rri-route**) or IPv6 colon-separated-hexadecimal (**ipv6 rri-route**) format.
- The next hop IP address is the SecGW clear interface physical address.
- *interface_name* specifies the egress interface. It should be unique and map the *vrf_name*, under the same context.
- The *vrf_name* is the VRF of clear interface in ASR 9000/RSP (external interface as well as the VSM interface) wherein the clear traffic is forwarded based on the RRI route.

ip/ipv6 sri-route Command

Important The **ip/ipv6 sri-route** CLI is deprecated, and not supported in 19.0 and later releases.

This Context Configuration mode command configures L3 High Availability Service Route Injection parameters:

```
configure
      context  context_name
{ ip | ipv6 } sri-route sri-ip network_address next hop nexthop_address interface
interface_name  [ vrf vrf_name ]
            end
```

Notes:

- *network_address* and *nexthop_address* are specified in IPv4 dotted-decimal (**ip sri-route**) or IPv6 colon-separated hexadecimal (**ipv6 sri-route**) notation.
- *interface_name* specifies the egress interface.

rri-mode Command

This ConnectedApps Configuration mode CLI command configures the supported RRI mode.

```
configure
      connectedapps
            rri-mode { both | none | ras | s2s }
            end
```

Notes:

- This command configures the anchor-route for an L3-L3 interchassis HA scenario.
  - **both** = enabled for RAS and S2S
  - **none** = disabled for all flow types
  - **ras** = Remote Access Service only
  - **s2s** = site-to-site only

Sample StarOS RRI HA Configurations

ConnectedApps (oneP) Configuration

```
config
     context local
          interface CA
               ip address 192.168.122.10 255.255.255.0
               exit
               subscriber default
               exit
               aaa group default
               exit
               no gtpp trigger direct-tunnel
               ip route 0.0.0.0 0.0.0.0 192.168.122.110 CA
          exit
          port ethernet 1/1
               no shutdown
               bind interface CA local
          exit
```

Intrachassis/Cluster Redundancy

```
config
     connectedapps
          sess-userid cisco
          sess-passwd cisco
          sess-name secgw
          sess-ip-address 172.29.98.14
          rri-mode ras
          ha-chassis-mode intra
          ha-network-mode L2
          activate
     exit
```

Figure 2: Intra-chassis/Cluster Redundancy

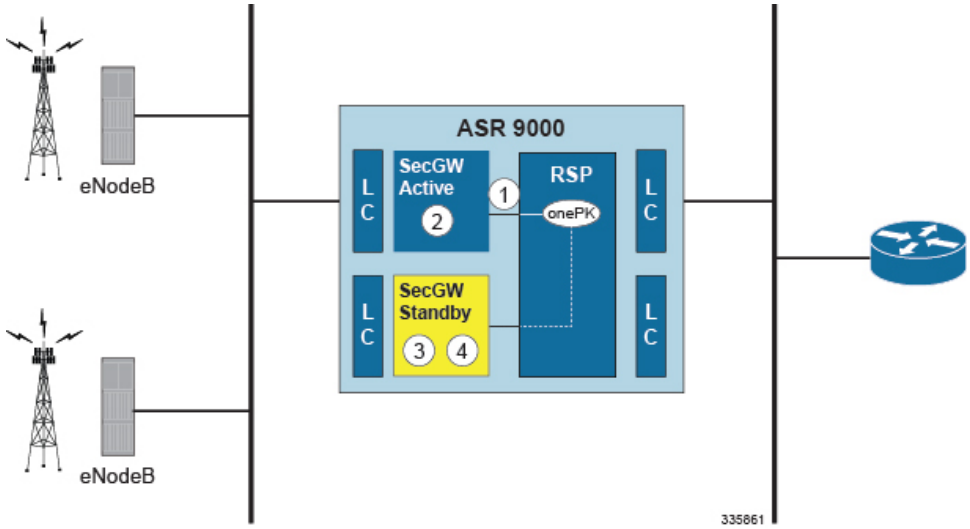

Item	Description
1	Common oneP session is used only by the active SecGW.
2	Only the active SecGW injects routes on tunnel setup.
3	Upon failover the currently active SecGW gives up its oneP session and the newly active SecGW takes over the session.
4	Upon failover the newly active SecGW injects routes for new tunnels.

L2 Interchassis Redundancy

```
config
     connectedapps
          sess-userid cisco
          sess-passwd cisco
          sess-name secgw
          sess-ip-address 172.29.98.14
          rri-mode ras
          ha-chassis-mode inter
          ha-network-mode L2
          activate
     exit
```

Figure 3: L2 Interchassis Redundancy

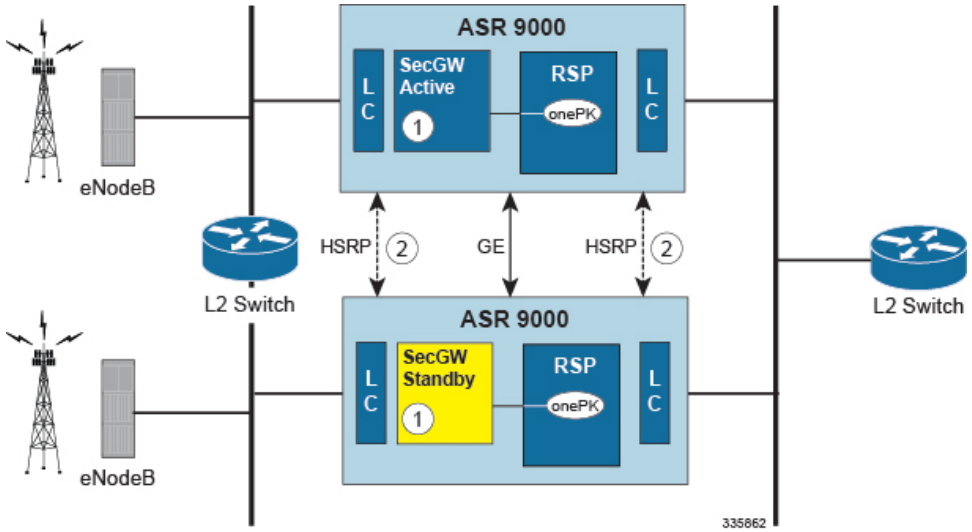

Item	Description
1	Both the active and standby SecGWs insert routes into local chassis only.
2	ICSR is configured to track RSP HSRP groups. HSRP also tracks SecGW using an SLA (Service Level Agreement).

L3 Interchassis Redundancy

```
config
     connectedapps
          sess-userid cisco
          sess-passwd cisco
          sess-name secgw
          sess-ip-address 172.29.98.14
          rri-mode ras
          ha-chassis-mode inter
          ha-network-mode L3
          activate
     exit
```

Figure 4: L3 Interchassis (Geo Redundancy) Mode

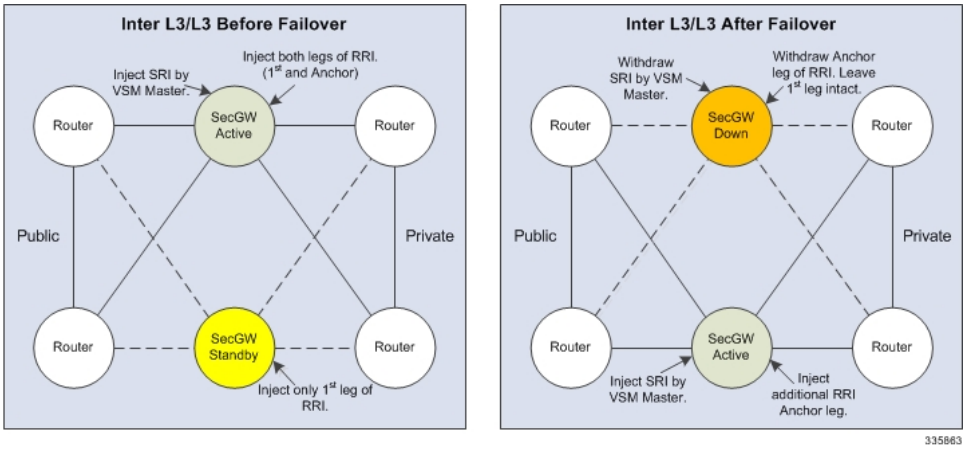

HSRP

Overview

Hot Standby Router Protocol (HSRP) is a Cisco proprietary redundancy protocol for establishing a fault-tolerant default gateway (RFC 2281). The protocol establishes a framework between network routers in order to achieve default gateway failover if the primary gateway becomes inaccessible.

Chassis-to-chassis redundancy employs HSRP to detect failure in the system and notify other elements of the need to change their HA State. Each VSM receives these notifications via oneP (Connected Apps) communication.

An external HSRP-aware entity switches traffic from the primary to the backup chassis. All application instances must failover to the backup chassis.

For additional information on HSRP, see the *ASR 9000 Series Aggregation Services Router IP Addresses and Services Configuration Guide.*

Figure 5: HSRP Notification

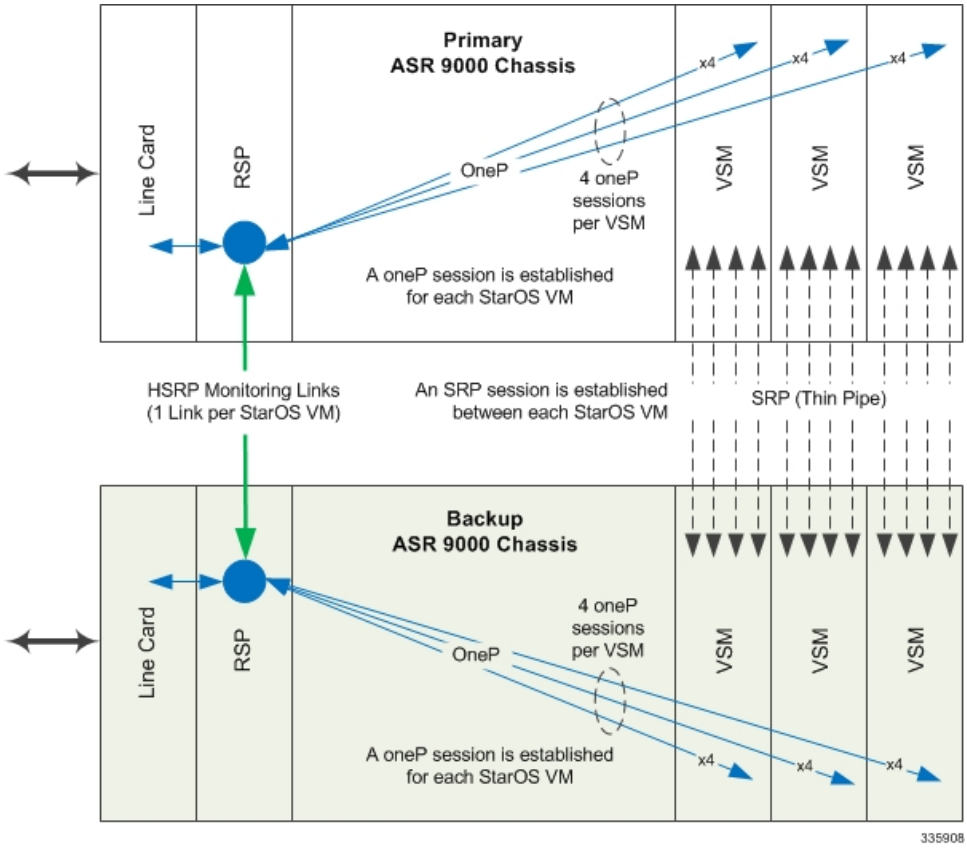

Each StarOS VM requires a separate oneP connection to the RSP (four oneP connections per VSM).

Each StarOS VM is monitored by a separate HSRP link that is established using sub-interfaces.

HSRP Configuration

Parameters

HSRP configuration parameters include:

- Interface name
- Address Family Identifier (AFI) type (IPv4 or IPv6)
- HSRP group number

Important The above parameters must match those of the HSRP configuration in the ASR 9000 RSP.

The following limits also apply to the HSRP configuration

- A maximum of one HSRP monitor is supported per VPC-VSM instance.
- The **monitor hsrp** command is associated with the SRP context.

monitor hsrp Command

The syntax for the **monitor hsrp** command is as follows:

```
config
 context srp_context
 monitor hsrp interface ifname afi-type type group hsrp_group
```

StarOS Configuration

HSRP monitoring must be enabled in the SRP configuration. A sample configuration is provided below.

Important You must configure HSRP for each VPC-VSM instance (StarOS VM) on the active and standby VSMs.

```
configure
     context srp
          service-redundancy-protocol
               checkpoint session duration 30
               route-modifier threshold 10
               priority 10
              monitor hsrp interface GigabitEthernet0/0/1/1 afi-type ipv4
 hsrp-group 4
               peer-ip-address 88.88.88.36
               bind address 88.88.88.37
               exit
```

Important HSRP monitoring is done via the ConnectedApps (oneP) interface in StarOS. A oneP session is established to all VPC-VSM instances on each VSM.

ASR 9000 RSP Configuration

HSRP must be configured on both the primary and backup ASR 9000 chassis. Sample IOS-XR configurations are provided below.

Primary ASR 9000 Chassis

```
router hsrp
   interface GigabitEthernet0/1/0/3
      address-family ipv4
         hsrp 2
            priority 110
            address 10.10.10.100
           |
          |
        |
       |
```

Backup ASR 9000 Chassis

```
router hsrp
   interface GigabitEthernet0/2/0/2
      address-family ipv4
         hsrp 2
            priority 100
            address 10.10.10.100
           |
          |
         |
        |
```