



BNG Geo Redundancy

This chapter provides information about support of geographical redundancy through subscriber redundancy groups (SRGs) and session redundancy groups (SERGs).

Table 1: Feature History for Establishing Geo Redundancy

Release	Modification
Release 5.2.2	Introduced BNG geo redundancy.
Release 5.3.1	Geo redundancy support for PPPoE sessions was added.
Release 5.3.3	Peer route disable functionality was added.
Release 6.1.2	These geo redundancy enhancements were added: <ul style="list-style-type: none"> • Active-active session support • State Control Route • Subscriber Redundancy Group Revertive Timer • Subscriber Redundancy Group-aware IPv6 Neighbor Discovery • Peer-to-peer Traffic Flow • Accounting Trigger Cause
Release 6.2.1	Session Redundancy Groups (SERGs) were introduced for DHCPv6 and IPv6 ND clients.
Release 6.2.2	Added the support for BNG Geo Redundancy over Cisco NCS 5000 Series nV satellite.
Release 6.3.1	Added Multiple State Control Routes for Each SRG feature.
Release 6.3.1	Added SRG Support for BNG SLAAC Sessions.

Release	Modification
Release 6.5.1	These new features were introduced: <ul style="list-style-type: none"> • Address pool usage synchronisation in BNG geo redundant active-active nodes • SRG support for static sessions • SRG for line card subscribers (12-connected IPoE only on QinQ based access interfaces)
Release 7.1.1	Added SRG support for BNG sessions on PWHE with DHCPv4 and DHCPv6 as server mode

This chapter covers these topics:

- [Geo Redundancy Overview, on page 2](#)
- [Supported Features in BNG Geo Redundancy, on page 7](#)
- [BNG Geo Redundancy Configuration Guidelines, on page 9](#)
- [Setting up BNG Subscriber Redundancy Group, on page 10](#)
- [Geographical Redundancy By Using a Session Redundancy Group \(SERG\), on page 15](#)
- [Geo Redundancy for PPPoE Sessions, on page 32](#)
- [BNG Geo Redundancy with Satellite, on page 36](#)
- [Geo Redundancy with Multihoming Ethernet Virtual Private Networks, on page 39](#)
- [Geo Redundancy Features, on page 48](#)
- [Deployment Models for BNG Geo Redundancy, on page 70](#)

Geo Redundancy Overview

To provide redundancy for the subscriber sessions, BNG supports Geographical Redundancy across multiple BNGs, without having any L1 or L2 connectivity between them. The BNG routers may be located in multiple geographical locations, and they have L3 connectivity over a shared core network through IP or MPLS routing.


Geo redundancy feature is supported for IPoE DHCP-triggered (IPv4, IPv6 and dual-stack) sessions and PPPoE (PTA and LAC) sessions.



Note PPPOE LAC geo redundancy is supported only with Multi-chassis Link Aggregation (MC-LAG) based access networks (active-standby mode) and RFC4951 compliant L2TP Network Server (LNS).

This figure depicts a BNG geo redundancy deployment network model:

Figure 1: BNG Geo Redundancy Deployment Network Model



The redundancy pairing between BNG routers work by synchronizing the state from the primary (active) to the subordinate (backup).

Geo redundancy works in conjunction with any of the access technologies. The CPEs are agnostic to redundancy; they see only one BNG or gateway. The access nodes are dual or multi-homed for redundancy using a variety of technologies based on the service provider network design and choices. Multi-chassis Link Aggregation (MC-LAG), dual-homed (Multiple Spanning Tree - Access Gateway or MST-AG), Ring (MST-AG or G.8032), xSTP and seamless MPLS (pseudowires) are a few such access networks.

Subscriber Redundancy Group (SRG)

Geo redundancy for subscribers is delivered by transferring the relevant session state from primary BNG to subordinate BNG which can then help in failover (FO) or planned switchover (SO) of sessions from one BNG to another. Subscriber Redundancy Group (SRG) which is a set of access-interface (or a single access-interface) is introduced in BNG, and all subscribers in an SRG would FO or SO as a group.

The SRG has two modes of operation:

- Hot-standby
- Warm-standby

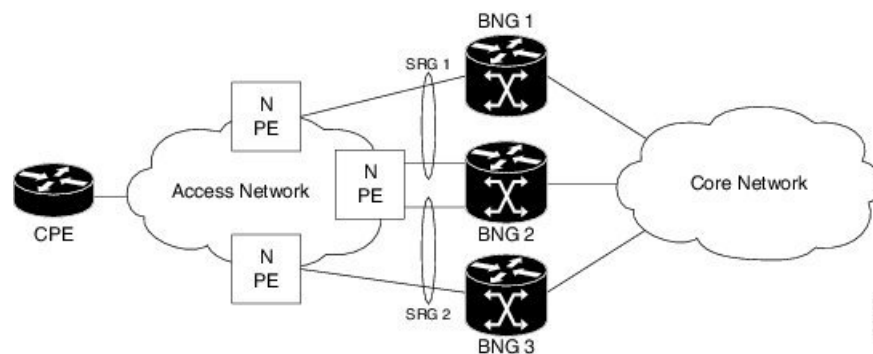
Currently BNG geo redundancy supports only the hot-standby subordinate mode. This is achieved by a 1:1 mirroring of subscriber session state from the primary to the subordinate where the entire provisioning is done before the FO or SO. The sessions provisioned on subordinate is in sync with the set up on the primary. Because the data plane is already set up for sub-second traffic impact, there is minimal action on switchover in the case of hot-standby mode and therefore, it is suitable for subscribers requiring high service level agreement (SLA). With appropriate capacity planning, the sessions can also be distributed across multiple BNGs to achieve an M: N model. The primary-subordinate terminology is always in the context of a specific SRG; not for the BNG device as a whole.



Note Even after the Subscriber Redundancy Group (SRG) configuration is removed from the subordinate node, the CPE continues to receive ARP replies from both the primary node and the subordinate node. This results in the network functioning in an uncertain manner. In order to avoid this uncertainty, shut down the access interface (that which corresponds to the subordinate node from which the configuration is removed) before removing the SRG configuration from the subordinate.

This figure depicts a typical BNG subscriber redundancy group (SRG):

Figure 2: BNG Subscriber Redundancy Group



SRG Virtual MAC

For seamless switchover between two BNGs, the L2-connected CPE devices must not detect change in gateway MAC and IPv4 or IPv6 addresses. The access technology like MC-LAG uses the same MAC address on both BNGs with active-standby roles, providing seamless switchover. Where MAC sharing is not provided by the access technology or protocol (like MST-AG, G.8032), the BNG SRG virtual MAC (vMAC) must be used. vMAC is configured as global MAC prefix or per SRG. This is integrated with BNG's dynamic primary or subordinate role negotiation; additional protocols like VRRP or HSRP is not needed. vMAC (and its derived IPv6 link-local address) is used for control protocol exchanges (for example, ARP, ND, DHCP, PPPOE and so on) and data traffic for subscriber sessions or services only. It allows real port MAC to be used for Ethernet protocols (like E-OAM, xSTP, G.8032 and so on) that are leveraged by the SRG for doing failure detection, recovery and MAC Flush.

Session Distribution Across SRG


The session distribution across SRGs can be in either of these modes:

- Active-standby mode:

In this mode, a dedicated backup BNG can be a subordinate for multiple SRGs from different active BNGs which are primaries for those respective SRGs.

This figure shows an active-standby mode of session distribution across SRGs:

Figure 3: Active-standby Mode of Session Distribution



In figure a:

- Sessions are associated with partitions (VLAN 1, 2, 3 and 4) on BNG1, with each VLAN mapped to separate SRG configured as primary role.
- BNG2 acts as backup for all VLANs.
- Each VLAN has 8K sessions terminated on it.

In figure b:

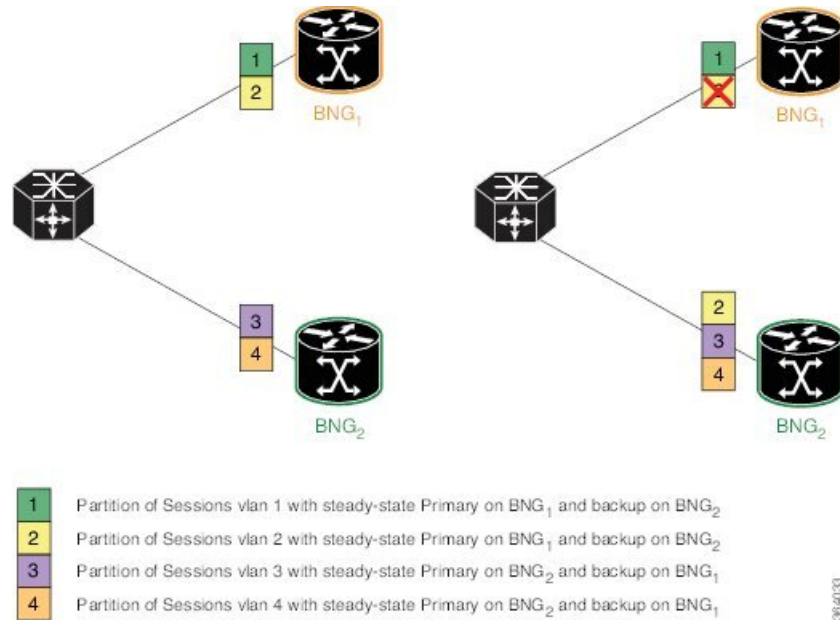
- An interface failure gets detected (using object-tracking of the access-interface) through MC-LAG.
- MC-LAG and SRG for each VLAN on BNG2 gets the primary role.
- All 32K sessions are switched to BNG2.
- BNG2 sees a session termination count of 32K.

- Active-active mode:

In this mode, a BNG can be primary for one SRG and a subordinate for another SRG at the same time.

This figure shows an active-active mode of session distribution across SRGs:

Figure 4: Active-active Mode of Session Distribution



In figure a:

- Sessions are associated with partitions (VLAN 1, 2) on BNG₁, with each VLAN mapped to separate SRG configured as primary role.
- Sessions are associated with partitions (VLAN 3, 4) on BNG₂, with each VLAN mapped to separate SRG configured as primary role.
- Each VLAN has 8K sessions terminated on it.
- Each BNG has 16K sessions terminated on it.

In figure b:

- The interface associated with VLAN 2 on BNG₁ goes down.
- Sessions associated with partitions (VLAN 2) on BNG₁ are switched to BNG₂.
- BNG₁ sees a session termination count of 8K and BNG₂ sees a session termination count of 24K.

Benefits of BNG Geo Redundancy

Major benefits of BNG Geo Redundancy include:

- Supports various redundancy models such as 1:1 (active-active) and M:N, including M:1.
- Provides flexible redundancy pairing on access-link basis.
- Works with multiple access networks such as MC-LAG, dual-home and OLT rings.
- Supports various types of subscribers such as IPv4, IPv6 and dual-stack IPoE sessions.
- Works for RP (bundle and virtual access-links) based subscribers.

- Provides failure protection to access link failures, LC failures, RP failures and chassis failures.
- Performs automatic switchovers during dynamic failures or planned events such as maintenance, upgrades and transitions.
- Co-exists with other high availability (HA) or redundancy mechanisms.
- Does switchover of the impacted session group only; other session groups remain on the same BNG.
- Provides fast convergence and rapid setup of sessions, with minimal subscriber impact during switchover.
- Provides automatic routing convergence towards core and efficient address pool management.
- Provides seamless switchover for subscriber CPE without the need for any signaling.
- Integrates with RADIUS or policy and charging rule function (PCRF) systems.
- Provides minimal to zero incremental load on back end servers and PCRFs during normal operations and switchover.
- Does not impact session scale and call-per-second (CPS) during normal operation.

Supported Features in BNG Geo Redundancy

Supported Features in BNG Geo Redundancy

These access topologies are supported:

- SRG active-active mode without any access protocol.
- MC-LAG topology (recommended only for IPv4 BNG sessions).
- Dual-home bundle interfaces with SRG vMAC using CFM or EFD fault detection and MST-AG for blocking.
- Ring bundle interfaces with SRG vMAC using CFM or EFD fault detection and MST-AG for blocking.
- Other access topologies and design variations may also be used for this feature.

These base geo redundancy features are supported:

- RP subscribers.
- LC subscribers (only for dual-stack IPoE sessions with BNG as DHCPv4 or DHCPv6 proxy), No PPPoE supported.
- Multiple SRG groups to different peer routers.
- Setting up peering statically through IPv4 or IPv6 TCP sessions.
- Hot-standby mode for subordinate (that is, subscribers provisioned in hardware on the subordinate as they are synchronized).
- Dynamic role negotiation between peers.
- Manual SRG switchover through command line interface (CLI).
- Dynamic failure detection using object tracking (link up-down, route and IPSLA tracking).

- Hold timer for dynamic switchover or switchback.
- Protocol bindings alone synchronized to subordinate; whereas AAA authorization for subscriber profile download performed by subordinate.
- Full BNG scale support (that is, half the scale number with redundancy).
- G.8032 (dual-home and ring) access technologies.
- PPPOE LAC geo redundancy only with Multi-chassis Link Aggregation (MC-LAG) based access networks (active-standby mode) and RFC4951 compliant L2TP Network Server (LNS).
- SRG for ambiguous VLAN BNG session is supported only for IPoE subscriber sessions over bundle interface.
- SRG between Cisco IOS XR 64-bit BNG node and 32-bit BNG node is supported.
- Starting Cisco IOS XR Software Release 7.6.1, PWHE SRG support for PPPoE session is available for both ambiguous and non-ambiguous VLAN.
- Starting Cisco IOS XR Software Release 7.9.1, PWHE SRG support for IPoE session is available for non-ambiguous VLAN.
- Starting Cisco IOS XR Software Release 7.10.1, PWHE SRG support for IPoE session is available for ambiguous VLAN.

These DHCP features are supported:

- DHCPv6 IA-NA and IA-PD support for L2 connected sessions.
- DHCPv4 support for L2 connected sessions.
- DHCPv4 or DHCPv6 dual-stack support.
- DHCP proxy mode.
- SRG support for BNG sessions on PWHE with DHCPv4 and DHCPv6 as server mode.



Note This feature is supported only on 64-bit Linux-based IOS XR ASR 9000 operating system.

- Session initiation through DHCPv4 or DHCPv6 protocol.
- Subscriber Redundancy Group (SRG) requires ARP table to be populated and is now compatible with the **subscriber arp scale-mode-enable** configuration. ARP entries maintained for each subscriber interface is required to send GARP during SRG role change from standby to active.

Unsupported Features and Restrictions for BNG Geo Redundancy

This section lists the unsupported features and restrictions for BNG geo redundancy.

These are not supported in BNG geo redundancy:

- IPoE packet-triggered sessions.
- Routed (L3 connected) sessions

- Multicast
- Both RP and LC subscribers do not support enabling fast switchover for subscriber framed-routes.
- PPPoE is not supported for LC subscriber sessions with SRG.
- IPoE and PPPoE sessions do not support `idle-timeout` configuration for both RP and LC-based subscribers.
- On Cisco ASR 9000 series router acting as a SRG primary, LAC sessions are not maintained across RPFO. Sessions are cleared during failover and session reestablishment starts when the next PPPoE discovery packet is received.

These are planned to be fully qualified only in future releases of Cisco IOS XR Software:

- Warm-standby subordinate mode.

BNG Geo Redundancy Configuration Guidelines

While configuring BNG geo redundancy, certain guidelines must be followed in these areas:

- BNG Configuration Consistency
- Access-link Integration
- Core Routing Integration
- RADIUS-PCRF Integration

BNG Configuration Consistency

- Geo redundancy feature infrastructure synchronizes individual subscriber session state from primary to subordinate. But, it does not synchronize the BNG related configurations (namely dynamic-template, DHCP profiles, policy-maps, access-interface configurations, external RADIUS or DHCP server and so on).
- For successful synchronization and setup of subscriber sessions between the two BNGs, it is mandatory that the relevant BNG configurations must be identical on the two routers and on the access-interfaces pairs in the SRG.
- While the access-interfaces or their types (or both) may vary between the paired BNGs, their outer-VLAN tag (that is, S-VLAN imposed by the access or aggregation devices) must be identical.
- Inconsistencies in base BNG or SRG configurations may result in synchronization failure and improper setup of sessions on the subordinate.

Access-link Integration

- You must use only those dual-homing techniques where one side is up or active, and the other side is down or standby. Both sides must not be up and forwarding traffic at the same time.
- You must use access-tracking mechanism under the SRG to ensure that its BNG role is always in synchronization with its access-link. Without this, the data or control traffic may get dropped.

- The access-tracking object used by the SRG must be same as the one used in the routing configuration for conditional advertisement of the subscriber summary route(s) corresponding to that SRG's subscriber address or subnet pool(s).
- Including multiple access-links (which do not fail or switchover their roles) together into a single SRG may be challenging, unless mechanisms are implemented to ensure that all these links change state even when one of them fails.
- Synchronisation of the framed IPv6 prefix addresses in SRG or SERG scenario is not supported on satellite bundle access interfaces in dual-homed satellite topology.

IPv6 Neighbour Discovery process maintains the synchronisation of a framed-prefix-pool. For the Neighbour Discovery process to work, the interface must be up. Bundle state is up only on the active host in a dual-homed satellite scenario.

Core Routing Integration

- Redistribution of individual subscriber routes into the routing protocol is not recommended because it slows convergence in failure or switchover events.
- Recommended design option is to conditionally advertise the summary static route for the subscriber address/subnet pool(s) of the SRG into the core routing protocol, through access-tracking.
- You can also advertise from both routers with different preferences and use various fast-reroute techniques.
- To avoid core routing changes in certain failure conditions, there are options to re-route the traffic from the subordinate to the primary (for example, a tunnel or inter-chassis link) for transient or prolonged intervals.
- Routing convergence and its correlation with access failures or convergence is a key to overall end-to-end service impact for subscribers. Multiple options exist to achieve sub-second intervals.

RADIUS-PCRF Integration

The backend policy and charging rule function (PCRF) system must send the CoA message to both primary and subordinate nodes. The message can be sent to the subordinate either at the same time as it is sent to primary, or it can be sent after the subordinate takes over the primary role and sends the Accounting START message.

From Cisco IOS XR Software Release R5.3.1 and later, the backend PCRF system need to send the CoA message only to the primary node.

Session Sync

Once the session is up on the primary node, the entire session information gets synced to the subordinate node. This includes dynamic synchronization of updates such as CoA or service logon. This is applicable from Cisco IOS XR Software Release R5.3.1 and later.

Setting up BNG Subscriber Redundancy Group

Guidelines in setting up SRG

Setting up SRG is subjected to these guidelines:

- The configurations and subscriber policies applied on the two routers (where the SRG access-interfaces are dual homing) must be identical to ensure seamless session mirroring and switchover.
 - SRG IDs (group IDs) must be same across BNGs.
 - Access-interface names or types need not be the same across routers.
 - Interface mapping-IDs must be same for the acces-interfaces across BNGs.
 - Server configurations (namely, RADIUS and DHCP configurations), IP pools, subscriber policies and templates must be identical across routers.
- The database of SRGs is scoped to a particular control plane instance (that is, at RP or LC node level). Therefore, you cannot form a single SRG with member links across LCs or with a mix of virtual interfaces (for example, bundles) and physical ports.
- The **global** BBA-Group is not valid for SRG, and hence the **pppoe bba-group global** command must not be used in BNG geo redundancy scenarios. Because **global** is a reserved keyword for IOS XR PPPoE call flow, you must use a different keyword for SRG.



Note You should configure a loopback interface under dynamic template only if it's different from the loopback interface associated with the access interface. Otherwise, after the SRG switchover, the new active BNG sends two GARP packets, one from the access interface with physical MAC and another one for the subscriber interface with VMAC for same IP, causing ARP conflict on the CPE.

Setting up a BNG subscriber redundancy group (SRG) involves these steps:

- Enable BNG Geo-Redundancy:

```
subscriber redundancy
source-interface loopback1
```

- Setup SRG and specify peer IPv4 or IPv6 address:

```
subscriber redundancy
group 1
peer 1.1.1.2
```

- Specify access-interfaces or VLANs, and mapping IDs:

```
subscriber redundancy
group 1
interface-list
interface Bundle-Ether1.10 id 210
```

- Set up access object tracking for SRG and summary subscriber route:

```
track mc-lag-bel
type line-protocol state
interface bundle-ether1

subscriber redundancy
```

```
group 1
  access-tracking mc-lag-bel

router static
  address-family ipv4 unicast
    200.0.0.0/16 Null0 track mc-lag-bel
```

Some optional configurations such as **preferred-role**, **slave-mode** and **hold-timer** also exist for SRG.



Note

Subscriber sessions should always be in sync between the SRG primary and subordinate. Don't clear all the SRG subscriber sessions manually. In an exceptional situation, follow the recommended steps for safe clearing of all SRG subscriber sessions.

To clear the subscriber sessions manually, perform the following:

To Clear SRG subordinate sessions:

1. Shut down the subordinate access interface.

For example,

```
RP/0/RSP0/CPU0# configure
RP/0/RSP0/CPU0(config)# interface Hu0/0/0/1
RP/0/RSP0/CPU0(config-if)# shutdown
```

2. Disable SRG (in global SRG configuration or SRG group level configuration).

For example,

```
Disable SRG in global configuration
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# subscriber redundancy disable

Disable SRG in group level configuration
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# subscriber redundancy group 1 disable
```

3. Clear subscriber sessions using the **clear subscriber srg slave session all** command.

For example,

```
RP/0/RSP0/CPU0:router#clear subscriber srg slave session all
```

4. Re-enable SRG (in global SRG configuration or SRG group level configuration).

For example,

```
Re-enable SRG in the global configuration:
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# session redundancy

Re-enable SRG in the group configuration:
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
```

5. Bring up the subordinate access interface.

For example,

```
RP/0/RSP0/CPU0# configure
RP/0/RSP0/CPU0(config)# interface Hu0/0/0/1
RP/0/RSP0/CPU0(config-if)# no shutdown
```

Clearing SRG primary sessions:

1. Perform SRG switchover using the **subscriber redundancy switchover [group <group>]** command.

For example,

```
RP/0/RSP0/CPU0# configure
RP/0/RSP0/CPU0:router(config)# subscriber redundancy switchover group 1
```

2. Shut down the primary access interface.

For example,

```
RP/0/RSP0/CPU0# configure
RP/0/RSP0/CPU0(config)# interface Hu0/0/0/1
RP/0/RSP0/CPU0(config-if)# shutdown
```

3. Disable SRG (in global SRG configuration or SRG group level configuration).

For example,

```
Disable SRG in global configuration
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# subscriber redundancy
RP/0/RSP0/CPU0:router(config-subscr-red)# disable

Disable SRG in group level configuration
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# subscriber redundancy group 1 disable
```

4. Clear subscriber sessions using the **clear subscriber srg master session all** command.

For example,

```
RP/0/RSP0/CPU0:router#clear subscriber srg master session all
```

5. Re-enable SRG (in global SRG configuration or SRG group level configuration).

For example,

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# session redundancy

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
```

6. Bring up the primary access interface.

For example,

```
RP/0/RSP0/CPU0# configure
RP/0/RSP0/CPU0(config)# interface Hu0/0/0/1
RP/0/RSP0/CPU0(config-if)# no shutdown
```


Geographical Redundancy By Using a Session Redundancy Group (SERG)

In large scale network implementations, it becomes essential to have redundancy between routers that share the same core network (IP and MPLS), but are geographically apart. A redundancy thus achieved is known as geographical redundancy, and often consists of a switchover (SO) from the active (primary) router to the standby (subordinate) router.

To achieve geographical redundancy for IPv6 Neighbor Discovery (ND) entries, or for DHCPv6 bindings, we use a Session Redundancy Group (SERG). A SERG comprises of sessions mapped to the access interfaces on the active RP of the router. If a single SERG is configured on the active RPs of the primary and subordinate


routers, then the router hosting the primary SERG serves as the primary, and the router hosting the subordinate SERG serves as the subordinate. This is illustrated in the following figure.

Figure 5: Geo Redundancy with a Single SERG



When multiple SERGs are configured on the active RPs, you could have both primary and subordinate SERGs on a single router. This is illustrated in the following figure.

Figure 6: Geo Redundancy with Multiple SERGs



Each router has an inbuilt redundancy between the RPs. When the active RP fails, the session (s) is transferred to the standby RP. This is known as a failover (FO).

The Session Redundancy Manager (SERM) runs on the active RP of both primary and subordinate routers. The SR clients running on the routers interact with the Session Redundancy Infrastructure (Session Redundancy Agent (SRA) and the Session Redundancy Library (SRL)).

The various components and their functions are briefly described as follows:

- **Session Redundancy Manager (SERM):** The SERM runs as a separate process on the active RP and manages the SERG configuration. The SERM peers with other routers that need to form a redundancy relationship, and establishes a point-to-multipoint communication channel to Session Redundancy Agents (SRAs) on the RP.
- **Session Redundancy Agent (SRA):** One or more SRAs run as a separate process on the active RP and supported line cards. A SRA acts on the SERG configuration, setting up operational context and database tables. The SRA implements the state machine for primary/subordinate selection and role change and orchestrates it using the TCP channel and provided APIs. The SRA receives the session entries on

the primary router and updates its database prior to synchronizing with the database on the subordinate router. The SRA orchestrates the session context setup on the subordinate router during the FO or SO. The SRA maintains a separate session database for each session client configured in the SERG.



Note The SRA works only on specific, defined keys, such as the IPv6 address, DHCPv6 client ID, and so on. Any undefined session data is handled as opaque data by the SRA. The respective session components must provide their access library to the SRA for handling any transformation or data retrieval.

- **Session Redundancy Library (SRL):** The SRL is used by session components for communicating with the SRA. The SRL uses IPC semantics for communicating with the SRA. SERG clients use an asynchronous API for storing and retrieving the session state from the SRL.

You can configure object tracking for one or more access interfaces in the SERG to enable automatic switchovers when an interface goes down. For more information on this configuration, see the *BNG Command Reference Guide for Cisco ASR 9000 Series Routers*.

Limitations for SERG

- If the Address Cached in the primary router: When the device configured as a primary router loses connectivity with the device configured as a subordinate router, the primary router continues to provide addresses assigned to the group. If one or more subscribers restart, the previously assigned IP addresses remain in the cache until communication is restored between both SERG entities. This situation could potentially exhaust the address pool if many subscribers frequently disconnect and reconnect, if a few subscribers continuously flap, or if the **clear subscribers command** is used.
- If a device is configured as a subordinate router: This device does not provide any address if the connectivity to the primary device is lost. In case this device assumes the role of primary, for example, after a session-redundancy switchover group was executed or for any other reason, there could be situations where address duplication occurs.

Guidelines for SERG

- The connectivity between SERG devices should be stable, as communication occurs over a TCP session using port 4001. It is good practice to prioritize this traffic.
- In the event of a prolonged disconnection between devices configured for SERG, it is advisable to configure a peer removal due to the caching mechanism. During the disconnection, addresses provided by the pool are tagged as local. Once the connection is restored, they will be tagged as local and remote.

Configuring and Verifying Session Redundancy for DHCPv6 Clients

Use the following procedure to configure geo-redundancy through session redundancy for DHCPv6 clients.

In this example, we configure Router 1 as Primary and Router 2 as Subordinate.

1. On Routers R1 and R2, enter the global configuration mode and configure session redundancy by specifying Loopback 0 as the source interface.

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# session redundancy
RP/0/RSP0/CPU0:router(config-session-red)# hold-timer 5
RP/0/RSP0/CPU0:router(config-session-red)# source-interface loopback0
```



Note The hold timer values on Routers R1 and R2 must match for them to peer with each other.

- Configure the session redundancy group by specifying the preferred role as Primary for Router R1 using the **master** keyword, and as subordinate for Router R2 using the **slave** keyword.

Router R1:

```
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
RP/0/RSP0/CPU0:router(config-session-red-group)# preferred-role master
RP/0/RSP0/CPU0:router(config-session-red-group)# hold-timer 7
RP/0/RSP0/CPU0:router(config-session-red-group)# peer 2.2.2.2
RP/0/RSP0/CPU0:router(config-session-red-group)# revertive-timer 5 maximum 15
RP/0/RSP0/CPU0:router(config-session-red-group)# interface-list
RP/0/RSP0/CPU0:router(config-session-red-grp-intf)# interface GigabitEthernet0/1/0/0 id
1
```

Router R2:

```
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
RP/0/RSP0/CPU0:router(config-session-red-group)# preferred-role slave
RP/0/RSP0/CPU0:router(config-session-red-group)# hold-timer 7
RP/0/RSP0/CPU0:router(config-session-red-group)# peer 1.1.1.1
RP/0/RSP0/CPU0:router(config-session-red-group)# revertive-timer 5 maximum 15
RP/0/RSP0/CPU0:router(config-session-red-group)# interface-list
RP/0/RSP0/CPU0:router(config-session-red-grp-intf)# interface GigabitEthernet0/1/0/0 id
1
```



Note The hold timer, revertive timer, and interface ID values on Routers R1 and R2 must match for them to peer with each other.

- Exit to the global configuration mode and commit your configuration on Routers R1 and R2.

```
RP/0/RSP0/CPU0:router(config)# commit
```

- Confirm your configuration on Router R1.

```
RP/0/RSP0/CPU0:router# show running-config session-redundancy
...
session-redundancy
source-interface Loopback0
hold-timer 5
group 1
  preferred-role master
  hold-timer 7
  peer 2.2.2.2
  revertive-timer 5 maximum 15
  interface-list
    interface GigabitEthernet0/1/0/0 id 1
  !
!
```

- Confirm your configuration on Router R2.

```
RP/0/RSP0/CPU0:router# show running-config session-redundancy
...
session-redundancy
source-interface Loopback0
hold-timer 5
group 1
  preferred-role slave
  hold-timer 7
  peer 1.1.1.1
  revertive-timer 5 maximum 15
  interface-list
    interface GigabitEthernet0/1/0/0 id 1
  !
!
```

6. Verify the session redundancy group on the routers by running the following show commands.

```
RP/0/RSP0/CPU0:router# show session-redundancy group
```

```
...
Session Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred Slave
           : H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S        : Peer Status
           : I - Initialize, Y - Retry, X - Cleanup, T - Connecting
           : L - Listening, R- Registered, C - Connected, E - Established
I/F Count: Interface Count
SS Count  : Session Count
```

Node Name	Group ID	Role	Flags	Peer Address	P/S	I/F Count
0/1/CPU0	1	Master	EMH-	2.2.2.2	E	
1	0	0				

```
Session Summary Count(Master/Slave/Total): 0/0/0
```

```
RP/0/RSP0/CPU0:router# show session-redundancy group 1
```

```
...
Session Redundancy Group ID: 1
Description      : <<not-configured>>
Status           : Enabled

Init-Role        : Master
Negotiated-Role  : Master
Current-Role     : Master

Hold Time        : 7
Revert Time      : 5

Tracking Status  : Enabled
Core-Tracking    : <<not-configured>>
Status          : n/a
Access-Tracking  : <<not-configured>>
Status          : n/a

Peer:
IP-address       : 2.2.2.2
Status           : Established
Role (Init/Neg/Cur) : Slave/Slave/Slave
Tracking Status  : Up
```

```
Last Neg-Time      : 2017 Mar  2 18:14:42
Last Up-Time       : 2017 Mar  2 18:14:42
Last Down-Time     : 2017 Mar  2 18:14:26
```

Switchover:

```
Last Switchover   : 2017 Mar  2 18:14:42      Reason       : Peer Up
Switchover Count  : 1
Hold Time         : Not-Running
Revert Time       : Not-Running
```

Session Statistics:

```
Count             : 0                      Slave-Upd-Fail : 0
Pending Update    : 0                      Pending Delete  : 0
Client:
  IPv6ND           : 0
  DHCPv6           : 0
```

Interface Count : 1

```
GigabitEthernet0/1/0/0                      Map-ID       : 1
```

```
RP/0/RSP0/CPU0:router# show session-redundancy summary interface
```

```
...
```

```
Session Redundancy Interface Summary
Status: E - Exists, F - Forward Reference
```

Interface Name	Status	Group ID	Map ID	Role
GigabitEthernet0/1/0/0	E	1	1	Master

7. Verify the SRG session information on the routers.

```
RP/0/RSP0/CPU0:router# show session-redundancy group 1 session verbose
```

```
...
```

```
Session Redundancy Agent Group Session
```

```
Flags: M-Master, V-Valid MAC, N-Neg Ack
```

```
Comp: SA-Agent, ND-ipv6nd, D6-dhcpv6
```

```
Comp Flags: U-Update, D-Delete, S-InSync, F-TxListFail, T-Dirty, C-Cleanup
```

```
Err Info: X-xxxx-ec - H/S - Hard/Soft, xxxx - No. of Times, ec - Error Code
```

Parent Interface	Key index	Flags
Comp Flags	Synchronization Error Info	
GigabitEthernet0/1/0/0	00030001ca011bba00000000000000000000000000000000	M-
SA{S} D6{S}	-	

8. Verify the SRA information and statistics.

```
RP/0/RSP0/CPU0:router# show session-redundancy agent interface
```

```
...
```

```
Session Redundancy Agent Interface
```

```
Status : F - Forward Referenced, S - Stale, R - Registered,
```

```
A - CAPS Added, O - Resource Owned, P - EOMS Pending
```

```
C - Pending CAPS Remove, U - Pending Reg Disable
```

```
Err Stats: Enable - Disable - Caps Add - Caps Remove - Attr Updated
```

Interface Name	ID	Group ID	Role	Status	Oper	Err Stats
GigabitEthernet0/1/0/0	1	1	Master	--RA----	-----	0-0-0-0-0

RP/0/RSP0/CPU0:router# show session-redundancy agent statistics

...

Session Redundancy Agent Summary - Node 0/0/CPU0

Process State : Active

Source Interface : Loopback0

VRF Name : default

IPv4 Address : 1.1.1.1

IPv6 Address : 192::2

Restart Client Sync In Progress : No

Client Init Sync TimeStamp : -

Restart Peer Sync In Progress : No

Peer Init Sync TimeStamp : -

Sync in Progress : No

Peer Action Timer : Not-Running

Retry Timer : Not-Running

Interface Status Statistics

Bound to group : 1

Non stale : 0

Pending caps remove : 0

Pending reg disable : 0

Pending other batch oper : 0

Sync in Progress : No

Client Statistics:

Status: U - Connection UP, S - Init-Sync Pending, E - Sync EOD Pending

Comp	Status	Up Timestamp	Down Timestamp	Cleanup Timer
SERGAGT	---	-	-	0
IPv6ND	U--	2017 Mar 2 18:14:25	-	0
DHCPv6	U--	2017 Mar 2 18:14:25	-	0

TxList Statistics:

	Ok	Part-Write	Clean
Marker Encode	4	0	4
Command Encode	0	0	0
Negotiation Encode	0	0	0

Client Statistics:

	Ok	NotOk
Invalid Registration		0
Invalid DeRegistration		0
Connection Up Count	2	
Connection Down Count		0
Message CallBack Count	2	
Message Received	4	0
Command Message Received	0	0
Session Message Received	4	0
Peer Done	2	

Peer Statistics:

	Ok	NotOk
Timer Handler	0	
Invalid Registration		0
Invalid DeRegistration		0

```

Message CallBack Count           :           0           0
Command Connection Up            :           0           0
Command Connection Down         :           0           0
Session Connection Up            :           0           0
Session Connection Down         :           0           0
Peer Done                        :           0           0
-----

```

9. Verify the DHCPv6 SR client information on the routers.

```
RP/0/RSP0/CPU0:router#show session-redundancy agent client dhcpv6
```

```
...
```

Session Redundancy Agent Client Statistics - Node 0/0/CPU0

Component - DHCPv6

```
Statistics:                               Ok           NotOk
```

Sent To Client:

```

Command
  Start of Download - SOD           :           1           0
  End of Download   - EOD           :           1           0
  End of Master Sync - EOMS         :           0           0
  Clear             - All            :           0           0
  Clear             - Selected       :           0           0
  Replay           - All            :           0           0
  Replay           - Selected       :           0           0
Session
  Update           :           0           0
  Delete           :           0           0

```

TxList Operation:

```

Encode - Complete Write           :           0
Encode - Partial Write            :           0
Cleanup CallBack                  :           0
Last Replay Count                 :           0

```

Received From Client:

```

Command
  Start of Download - SOD - All     :           1
  Start of Download - SOD - Selected :           0
  End of Download   - EOD - All     :           1
  End of Download   - EOD - Selected :           0
  End of Master Sync - EOMS         :           0
  Clear             - All            :           0
  Clear             - Selected       :           0
  Replay           - All            :           0
  Replay           - Selected       :           0
Session
  Update           :           0           0
  Delete           :           0           0
  Negative Acknowledgement         :           0           0

```

Client Activity Statistics:

```

Active           :           1           0
Deactive         :           0           0
Registration     :           1           0
DeRegistration   :           0
Connection Down  :           0
Cleanup          :           0
-----

```

Session Redundancy Agent Client Statistics - Node 0/1/CPU0

Component - DHCPv6

```
Statistics:                               Ok           NotOk
```

```

-----
Sent To Client:
Command
  Start of Download - SOD           :           1           0
  End of Download   - EOD           :           1           0
  End of Master Sync - EOMS         :           1           0
  Clear             - All            :           0           0
  Clear             - Selected       :           0           0
  Replay           - All            :           0           0
  Replay           - Selected       :           0           0
  Session          :           0           0
  Update          :           0           0
  Delete          :           0           0

TxList Operation:
  Encode - Complete Write           :           0
  Encode - Partial Write            :           0
  Cleanup CallBack                  :           0
  Last Replay Count                  :           0

Received From Client:
Command
  Start of Download - SOD - All     :           1
  Start of Download - SOD - Selected :           0
  End of Download   - EOD - All     :           1
  End of Download   - EOD - Selected :           0
  End of Master Sync - EOMS         :           0
  Clear             - All            :           0
  Clear             - Selected       :           0
  Replay           - All            :           0
  Replay           - Selected       :           0
  Session          :           0           3
  Update          :           0           2
  Delete          :           0           0
  Negative Acknowledgement          :           0           0

Client Activity Statistics:
  Active           :           1           0
  Deactive         :           0           0
  Registration     :           1           0
  DeRegistration   :           0
  Connection Down  :           0
  Cleanup          :           0
-----

```

Session Redundancy Agent Client Statistics - Node 0/2/CPU0

Component - DHCPv6

```

Statistics:                               Ok           NotOk
-----

```

```

Sent To Client:
Command
  Start of Download - SOD           :           1           0
  End of Download   - EOD           :           1           0
  End of Master Sync - EOMS         :           0           0
  Clear             - All            :           0           0
  Clear             - Selected       :           0           0
  Replay           - All            :           0           0
  Replay           - Selected       :           0           0
  Session          :           0           0
  Update          :           0           0
  Delete          :           0           0

```

```

TxList Operation:
  Encode - Complete Write           :           0

```

```

Encode - Partial Write          :          0
Cleanup CallBack               :          0
Last Replay Count              :          0

Received From Client:
Command
  Start of Download - SOD - All :          1
  Start of Download - SOD - Selected :          0
  End of Download - EOD - All :          1
  End of Download - EOD - Selected :          0
  End of Master Sync - EOMS :          0
  Clear - All :          0
  Clear - Selected :          0
  Replay - All :          0
  Replay - Selected :          0
Session
  Update :          0          0
  Delete :          0          0
  Negative Acknowledgement :          0          0

Client Activity Statistics:
Active :          1          0
Deactive :          0          0
Registration :          1          0
DeRegistration :          0
Connection Down :          0
Cleanup :          0

```

You have successfully configured and verified geo redundancy using session redundancy groups for DHCPv6 clients.

Managing Session Redundancy Groups

After you have configured and verified the session redundancy groups (SERGs), you can use the commands in this section to trigger a manual switchover, trigger a manual synchronization, or clear sessions for all or a specific SERG.

Triggering a Manual Switchover

After you have configured SERGs on the primary and subordinate routers, if you want to remove/replace the primary router, you can trigger a manual switchover from the primary to the subordinate by running the following commands.



Note The following commands can be executed only on the primary router.

- To trigger a redundancy switchover for all SERGs, run the following command.

```
RP/0/RSP0/CPU0:router# session redundancy switchover
```

- To trigger a redundancy switchover for a specific SERG, run the following command.

```
RP/0/RSP0/CPU0:router# session redundancy switchover group 210
```


Triggering Manual Synchronization

If the sessions between the primary and subordinate routers are not getting synchronized, either because of some change in the network topology, or some network latency, you can trigger synchronization manually by running the following commands.



Note The following commands can be executed on either the Primary or the Subordinate router.

- To trigger a redundancy synchronization for all SERGs, run the following command.

```
RP/0/RSP0/CPU0:router# session redundancy synchronize
```

- To trigger a redundancy synchronization for a specific SERG, run the following command.

```
RP/0/RSP0/CPU0:router# session redundancy synchronize group 210
```

Clearing Sessions in a SERG

If you want to clear the existing sessions on the primary and subordinate routers, either because of a switchover, or a change in network topology, you can run the following commands.



Note The following commands can be executed on either the primary or the subordinate router.

When issued on the subordinate, the session context is deleted from the router and a synchronization is requested with the primary. If the router is in hot-standby mode, the sessions are deleted on the subordinate.

When issued on the primary, the session entries are deleted first on the primary and later on the subordinate. The SRA then requests a fresh session from the SR client, which is eventually synchronized with the subordinate.

- To clear sessions for all SERGs, run the following command.

```
RP/0/RSP0/CPU0:router# clear session-redundancy
```

- To clear sessions for a specific SERG, run the following command.

```
RP/0/RSP0/CPU0:router# clear session-redundancy group 1
```

Configuring and Verifying Session Redundancy for IPv6 ND Clients

Use the following procedure to configure geo-redundancy through session redundancy for IPv6 ND clients.

In this example, we configure Router 1 as Primary and Router 2 as Subordinate.

1. On Routers R1 and R2, enter the global configuration mode and configure session redundancy by specifying Loopback 0 as the source interface.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# session redundancy
RP/0/RSP0/CPU0:router(config-session-red)# hold-timer 5
RP/0/RSP0/CPU0:router(config-session-red)# source-interface loopback0
```



Note The hold timer values on Routers R1 and R2 must match for them to peer with each other.

- Configure the session redundancy group by specifying the preferred role as Primary for Router R1 using the **master** keyword, and as subordinate for Router R2 using the **slave** keyword.

Router R1 :

```
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
RP/0/RSP0/CPU0:router(config-session-red-group)# preferred-role master
RP/0/RSP0/CPU0:router(config-session-red-group)# hold-timer 7
RP/0/RSP0/CPU0:router(config-session-red-group)# peer 2.2.2.2
RP/0/RSP0/CPU0:router(config-session-red-group)# revertive-timer 5 maximum 15
RP/0/RSP0/CPU0:router(config-session-red-group)# interface-list
RP/0/RSP0/CPU0:router(config-session-red-grp-intf)# interface GigabitEthernet0/1/0/0 id
1
```

Router R2:

```
RP/0/RSP0/CPU0:router(config)# session redundancy group 1
RP/0/RSP0/CPU0:router(config-session-red-group)# preferred-role slave
RP/0/RSP0/CPU0:router(config-session-red-group)# hold-timer 7
RP/0/RSP0/CPU0:router(config-session-red-group)# peer 1.1.1.1
RP/0/RSP0/CPU0:router(config-session-red-group)# revertive-timer 5 maximum 15
RP/0/RSP0/CPU0:router(config-session-red-group)# interface-list
RP/0/RSP0/CPU0:router(config-session-red-grp-intf)# interface GigabitEthernet0/1/0/0 id
1
```



Note The hold timer, revertive timer, and interface ID values on Routers R1 and R2 must match for them to peer with each other.

- Exit to the global configuration mode and commit your configuration on Routers R1 and R2.

```
RP/0/RSP0/CPU0:router(config)# commit
```

- Confirm your configuration on Router R1 (Primary).

```
RP/0/RSP0/CPU0:router# show running-config session-redundancy
...
session-redundancy
source-interface Loopback0
hold-timer 5
group 1
  preferred-role master
  hold-timer 7
  peer 2.2.2.2
  revertive-timer 5 maximum 15
  interface-list
  interface GigabitEthernet0/1/0/0 id 1
  !
  !
  !
```

5. Confirm your configuration on Router R2.

```
RP/0/RSP0/CPU0:router# show running-config session-redundancy
...
session-redundancy
 source-interface Loopback0
 hold-timer 5
 group 1
   preferred-role slave
   hold-timer 7
   peer 1.1.1.1
   revertive-timer 5 maximum 15
   interface-list
     interface GigabitEthernet0/1/0/0 id 1
   !
 !
 !
```

6. Verify the session redundancy group on the routers by running the following show commands.

```
RP/0/RSP0/CPU0:router# show session-redundancy group
```

```
...
Session Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred Slave
            H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S       : Peer Status
            I - Initialize, Y - Retry, X - Cleanup, T - Connecting
            L - Listening, R- Registered, C - Connected, E - Established
I/F Count: Interface Count
SS Count  : Session Count
```

Node Name	Group ID	Role	Flags	Peer Address	P/S	I/F Count
SS Count	Sync Pending					
0/1/CPU0	1	Master	EMH-	2.2.2.2	E	
1	0	0				

```
Session Summary Count(Master/Slave/Total): 0/0/0
```

```
RP/0/RSP0/CPU0:router# show session-redundancy group 1
```

```
...
Session Redundancy Group ID: 1
Description      : <<not-configured>>
Status           : Enabled

Init-Role        : Master
Negotiated-Role  : Master
Current-Role     : Master

Hold Time        : 7
Revert Time      : 5

Tracking Status  : Enabled
Core-Tracking    : <<not-configured>>
  Status         : n/a
Access-Tracking  : <<not-configured>>
  Status         : n/a

Peer:
IP-address       : 2.2.2.2
Status           : Established
Role (Init/Neg/Cur): Slave/Slave/Slave
```



```

-----
Interface Name          | ID | Group ID | Role | Status | Oper | Err Stats
-----
GigabitEthernet0/1/0/0          1      1  Master  --RA----  -----  0-0-0-0-0
-----
    
```

RP/0/RSP0/CPU0:router# show session-redundancy agent statistics

...

Session Redundancy Agent Summary - Node 0/0/CPU0

```

Process State           : Active

Source Interface       : Loopback0
VRF Name               : default
IPv4 Address           : 1.1.1.1
IPv6 Address           : 192::2
    
```

```

Restart Client Sync In Progress : No
Client Init Sync TimeStamp      : -
Restart Peer Sync In Progress   : No
Peer Init Sync TimeStamp       : -
Sync in Progress                : No
Peer Action Timer               : Not-Running
Retry Timer                     : Not-Running

Interface Status Statistics
Bound to group                  : 1
Non stale                       : 0
Pending caps remove            : 0
Pending reg disable            : 0
Pending other batch oper       : 0
Sync in Progress               : No
    
```

Client Statistics:

Status: U - Connection UP, S - Init-Sync Pending, E - Sync EOD Pending

```

-----
Comp | Status | Up Timestamp | Down Timestamp | Cleanup Timer
-----
SERGAGT    ---  -              -              0
IPv6ND     U--  2017 Mar 2 18:14:25 -              0
DHCpv6     U--  2017 Mar 2 18:14:25 -              0
-----
    
```

TxList Statistics:

```

-----
Ok      Part-Write      Clean
-----
Marker Encode           :           4           0           4
Command Encode          :           0           0           0
Negotiation Encode      :           0           0           0
    
```

Client Statistics:

```

-----
Ok      NotOk
-----
Invalid Registration    :           0
Invalid DeRegistration  :           0
Connection Up Count     :           2
Connection Down Count   :           0
Message CallBack Count  :           2
Message Received        :           4           0
Command Message Received :           0           0
Session Message Received :           4           0
Peer Done               :           2
    
```

Peer Statistics:

```

-----
Ok      NotOk
-----
    
```

```

Timer Handler                :                0
Invalid Registration          :                0
Invalid DeRegistration        :                0
Message CallBack Count       :                0
Command Connection Up        :                0
Command Connection Down      :                0
Session Connection Up        :                0
Session Connection Down      :                0
Peer Done                    :                0

```

...

9. Verify the IPv6 ND SR client information on the routers.

```
RP/0/RSP0/CPU0:router# show session-redundancy client ipv6nd
```

```
Session Redundancy Client Statistics - Node 0/0/CPU0
```

```

Config      : True
Status     : Enabled
Active      : True

```

```

Connection Status      : Connected
Last Connection Up Time : 2017 Mar  7 10:28:03
Last Connection Down Time : 1970 Jan  1 05:30:00

```

TxList Operation:

```

Message CallBack                :                2
Encode - Complete Write         :                7
Encode - Partial Write          :                0
Cleanup CallBack                :                0
Decode Message Error            :                0
Unknown Operation Error         :                0

```

TxList Statistics:

	Ok	Part-Write	Clean
Marker Encode	2	0	2
Command Encode	0	0	0
Negotiation Encode	0	0	0

Statistics:

	Ok	NotOk

Sent To Agent:

```

Command
  Start of Download - SOD          :                1                0
  End of Download   - EOD          :                1                0
  End of Master Sync - EOMS        :                0                0
  Clear             - All           :                0                0
  Clear             - Selected      :                0                0
  Replay           - All           :                0                0
  Replay           - Selected      :                0                0
Session
  Add              :                7                0
  Delete           :                7                0
Negative Acknowledgement
  Synchronous     :                0                0
  Asynchronous    :                0                0

```

Received From Agent:

```

Message                :                9                0
Command
  Start of Download - SOD - All    :                1                0
  Start of Download - SOD - Selected :                0                0
  End of Download   - EOD - All    :                1                0
  End of Download   - EOD - Selected :                0                0

```

```

End of Master Sync - EOMS          :          7
Clear          - All                :          0
Clear          - Selected           :          0
Replay        - All                :          1
Replay        - Selected           :          0
Session                          :          9          0
Update                          :          0          0
Delete                          :          0          0

```

Agent Activity Statistics:

```

Active          :          1          0
Deactive       :          0          0
Connection Up   :          1          0
Connection Down :          0
Peer Done      :          8
...

```

```
RP/0/RSP0/CPU0:router# show ipv6 nd statistics
```

Service	Attrib	Oper	Success#	Failure#	MinTime (usec)	MaxTime (usec)	AvgTime (usec)
AIB	IDB	Init	1	0	84110	84110	84110
AIB	IDB	Reg	1	0	295	295	295
CHKPT	IDB	Init	1	0	73	73	73
CHKPT	IDB	Update	47	0	23	100	27
CHKPT	NBR	Create	4	0	68	110	80
CHKPT	NBR	Del	2	0	150	1475	812
CHKPT	NBR	Init	1	0	232	232	232
CHKPT	NBR	Update	4	0	54	166	120
GSP	BATCH	Sent	4	32	0	2411	68
GSP	IDB	Init	1	0	5197	5197	5197
GSP	IDB	Reg	1	0	1816	1816	1816
GSP	IDB	Sent	24	0	13	104	15
GSP	MNODE	Sent	8	0	14	2482	322
GSP	PING	Recv	4	0	0	0	0
GSP	PING	Sent	4	0	76	1476	475
GSP	Ucast	Recv	5	0	0	0	0
IM	IDB	Create	10	0	0	0	0
IM	IDB	Init	2	0	670	4831	2750
IM	IDB	Reg	1	0	465	465	465
IM	MAC	Recv	8	0	0	0	0
IM	MAC	Reg	1	0	134	134	134
IM	MTU	Recv	8	0	0	0	0
IM	MTU	Reg	1	0	121	121	121
IM	VLAN	Recv	7	0	0	0	0
IM	VLAN	Reg	1	0	144	144	144
IM	MEMBER	Recv	24	0	0	0	0
IM	MEMBER	Reg	1	0	1603	1603	1603
LPTS	IDB	Add	8	0	2	174	25
LPTS	IDB	Init	1	0	4545	4545	4545
LPTS	IDB	Update	10	0	0	1	0
IPV6-MA	BATCH	Recv	3	0	0	0	0
NETIO	IDB	Init	1	0	244787	244787	244787
NETIO	RA	Recv	14	0	0	0	0
NETIO	NA	Recv	4	0	0	0	0
NETIO	NS	Recv	4	0	0	0	0
STATS	IDB	Init	1	0	21410	21410	21410

```

STATS   IDB   Reg      8      0      2      19532      2449

```

```

Session Redundancy Stats
-----

```

Type	Success	Error
serg_init	1	0
serg_shutdown	0	0
serg_activate	1	0
serg_active_txlist_add	0	0
serg_active_txlist_del	0	0
serg_active_txlist_encode	6	0
serg_active_txlist_clean	6	0
serg_active_replay	1	0
serg_active_cleanup	0	0
serg_standby_receive	0	0
serg_standby_sess_update	0	0
serg_standby_sess_delete	0	0
serg_standby_sess_nack	0	0
serg_standby_sess_mark	1	0
serg_standby_sess_sweep	1	0
serg_standby_cleanup	0	0

You have successfully configured and verified geo redundancy using session redundancy groups for IPv6 ND clients.

For information on managing SERGs, see [Managing Session Redundancy Groups, on page 24](#).

Geo Redundancy for PPPoE Sessions

BNG supports geo redundancy for PPPoE-PPP Termination and Aggregation (PPPoE-PTA) and PPPoE-L2TP Access Concentrator (PPPoE-LAC) sessions.


PPPoE-PTA Geo Redundancy

Geo redundancy behavior for the PPPoE-PTA sessions remains the same as for basic geo redundancy set up, except that the keepalives are disabled on the subordinate BNG node. The keepalives are sent only after the subordinate switches its role to primary.

PPPoE-LAC Geo Redundancy

This figure shows a PPPoE-LAC Geo Redundancy set up with BNG

Figure 7: PPPoE-LAC Geo Redundancy Topology



For a PPPoE-LAC geo redundancy setup, the SRG is formed by grouping together the access-links on which LAC sessions are to arrive (co-exists with PTA). To enable SRG level redundancy switchover, tunnels for each SRG for each L2TP network server (LNS) must be setup. L2TP ensures that sessions belonging to different SRGs do not share the same tunnel even if they are going to the same LNS. The tunnel is set up on both primary and subordinate nodes. By default, the tunnel is down on subordinate and it gets activated upon switchover. The BNG sync takes care of both tunnel and session-state sync from the primary to the subordinate. The L2TP tunnel attributes and negotiated parameters are also synchronized through the BNG sync.

You must use this command in subscriber redundancy group configuration mode, to configure the source IP used for L2TP tunnel for subscribers coming under an SRG group:

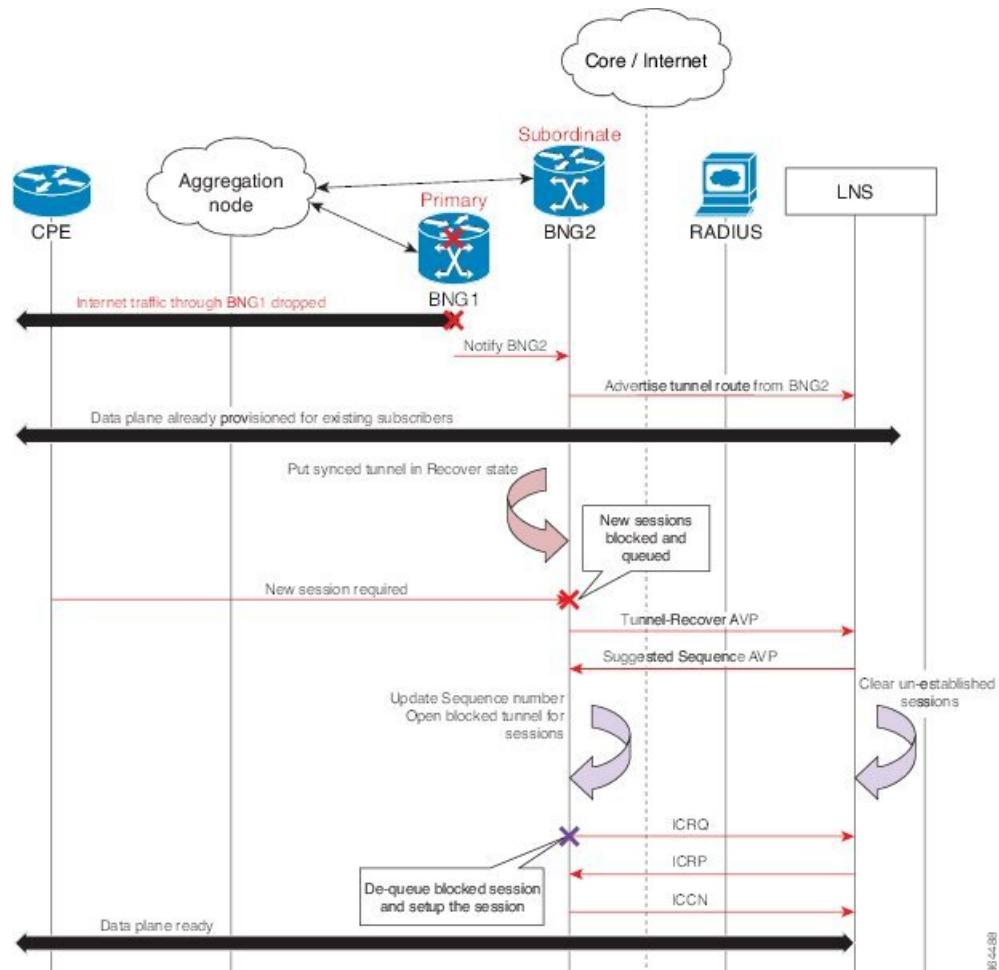
```
l2tp-source-ip ipv4-address
```

This ensures that there is a separate tunnel from each SRG group, in spite of having the same LNS.

PPPoE-LAC Session Switchover

This figure shows the call flow of PPPoE-LAC session switchover.

Figure 8: PPPoE-LAC Session Switchover



During switchover, the tunnel endpoint switches from the primary (BNG1) to subordinate (BNG2) node as soon as the routing converges, and advertises the loopback address of subordinate (BNG2) to the LNS. The sessions and tunnels that are already provisioned on the data path on subordinate (BNG2) then seamlessly take over. The L2TP control plane on subordinate (BNG2) places the tunnel in re-sync state to recover the tunnel sequence number (Ns and Nr) during which only control messages are queued up for further processing. After the tunnel recovery, the LAC gets the sequence number from the LNS. The existing tunnels or sessions are not lost as the subordinate (BNG2) takes over. The signaling for the new session resumes and the queued requests also get processed. The unestablished sessions are then cleared off. For LNS, this switchover appears to be a convergence event where the tunnel has flapped.

Verification of Geo Redundancy for PPPoE Sessions

Listed below are some of the show commands that can be used to verify the Geo Redundancy configuration in BNG. For complete command reference, see the *Subscriber Commands*, *PPPoE Commands* and *PPPoE LAC-Specific Commands*, chapters in the *Cisco ASR 9000 Series Aggregation Services Router Broadband Network Gateway Command Reference*.

- `show subscriber redundancy group 210`

```

Subscriber Redundancy Group ID: 210
  Description : <<not-configured>>

Status           : Enabled
Init-Role          : Master
Negotiated-Role    : Master           Current-Role : Master

Slave-mode         : Hot               Hold Time : 15
- - -
- - -
Peer:
  11::2                               Status : Established
  Role(Init/Neg/Cur): Slave/Slave/Slave
  Tracking Status   : Down
- - -
- - -
Switchover:
  Last Switchover   : 2014 Sep 12 07:12:11      Reason : Object Tracking Status
Change
- - -
- - -
Subscriber Session Statistics:
  Count             : 8000                    Slave-Upd-Fail : 0
  Pending Update    : 0                      Pending Delete  : 0
  Tunnel Count      : 0

Interface Count     : 1
  Bundle-Ether1.10 Map-ID                   : 210

```

• show ppp interfaces

```

Bundle-Ether2.1.pppoe16534 is up, line protocol is up
SRG Role: Slave
LCP: Open
  Keepalives enabled (60 sec, retry count 5)
  Local MRU: 1492 bytes
  Peer MRU: 65531 bytes
Authentication
  Of Peer: PAP (Completed as user1@domain.com)
  Of Us: <None>
IPCP: Open
  Local IPv4 address: 12.16.0.1
  Peer IPv4 address: 12.0.250.23
IPv6CP: Initial
  Local IPv6 address: fe80::
  Peer IPv6 address: fe80::

```

• show pppoe interfaces

```

Bundle-Ether2.1.pppoe16534 is Complete
Session id: 16534
Parent interface: Bundle-Ether2.1
BBA-Group: BBA1
Local MAC address: 0002.0003.0004
Remote MAC address: 0000.6201.0103
Outer VLAN ID: 10
Tags:
  Service name: AGILENT
  Host-Uniq: 4 bytes, (000e0000)

```

```
SRG-state: SRG-Standby
```

- **show vpdn**

```
RP/0/RSP0/CPU0:router# show vpdn session
```

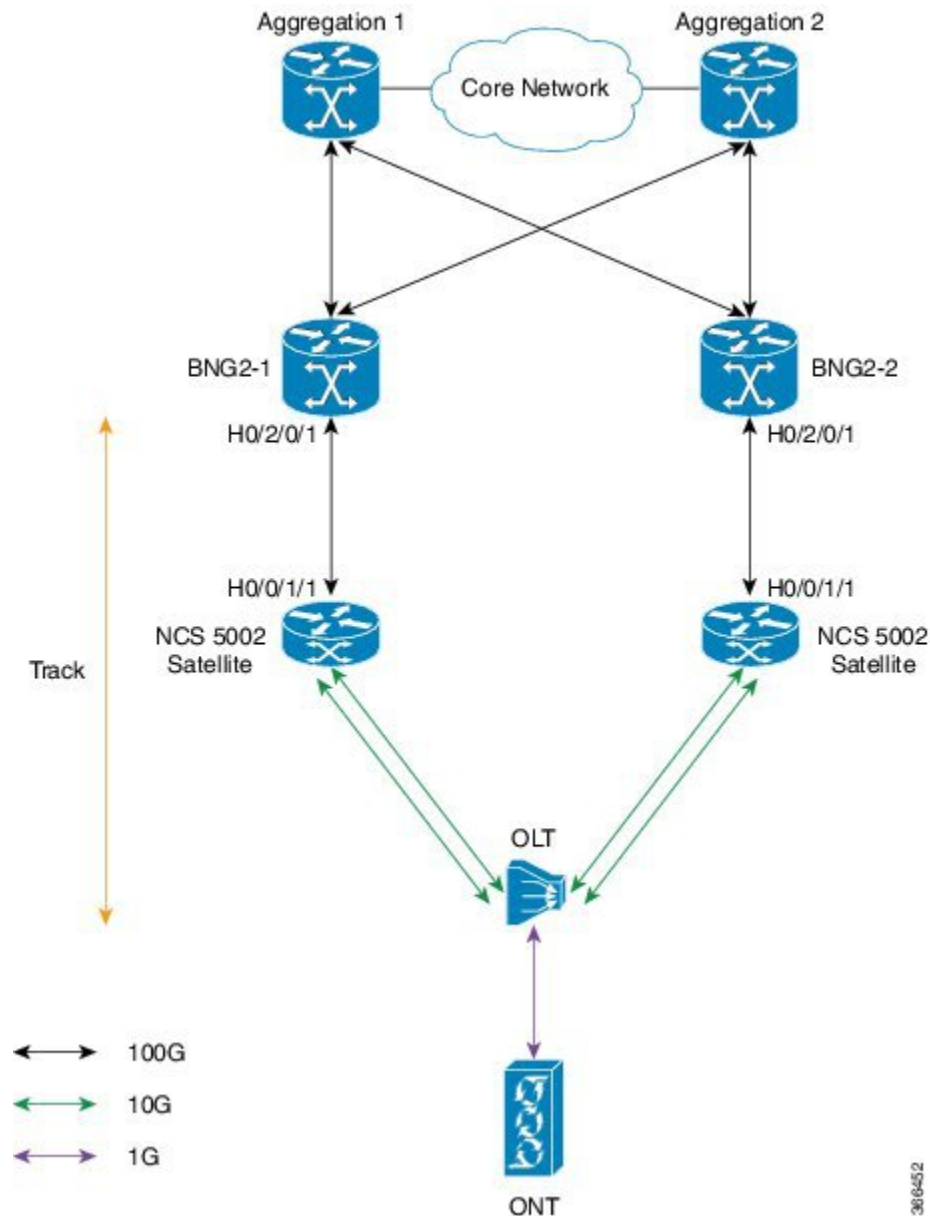
```
SRG Role: Master
Subscriber label: 0x42, interface name: Bundle-Ether1.10.pppoe3
user name: user1@lms2.com
parent interface: Bundle-Ether1.10
state: est last change: 00:01:01
time to setup session: 0:2 (s:msec)
conditional debug flags: 0
L2TP data
local end point: 11.1.1.1 remote end point: 19.9.9.2
call serial number: 1970100002
local tunnel id: 46813 remote tunnel id: 40849
local session id: 36198 remote session id: 33437 remote port: 1701
tunnel assigned id:
tunnel client authentication id: LAC
tunnel server authentication id: LNS
tunnel authentication: disabled
class attribute mask:
Subscriber data
NAS port id: 0/0/1/10
NAS port type: Virtual PPPoE over VLAN
physical channel id: 0
Rx speed: 1000000000, Tx speed: 1000000000
Configuration data
table id: 0xe0000000, VRF id: 0x60000000, VPN id: 0:0
VRF name: default
dsl line info forwarding: disabled, l2tp busy timeout: 60
TOS mode: default
```

BNG Geo Redundancy with Satellite

From Cisco IOS XR Software Release 6.2.2 and later, the BNG geo redundancy feature in Cisco ASR 9000 Series Routers is enhanced to provide a satellite-based solution. The satellite box provides high density 10-Gigabit ports to terminate optical line terminals (OLTs) which works seamlessly with BNG geo redundancy solution for loss-of-signal (LOS) based detection and failover. Currently, this feature is supported only on Cisco IOS XR 32 bit IOS XR operating system, and only with the Cisco NCS 5000 Series nV satellite.

Sample Topology of BNG Geo Redundancy with Cisco NCS 5000 Series nV Satellite

Figure 9: BNG Geo Redundancy with Cisco NCS 5000 Series nV Satellite



The Cisco NCS 5000 Series (NCS 5001 or NCS 5002) nV satellite which is used in this topology, provides the functionality of an extended or virtual line card for BNG, thereby increasing the access ports on BNG.

Configure BNG Geo Redundancy with Cisco NCS 5000 Series nV Satellite

You have to accomplish the following in order to configure BNG Geo Redundancy with Cisco NCS 5000 Series nV satellite:

- Global satellite configuration

- ICL configuration
- Access-tracking configuration

Except for the additional configurations related to satellite, all other configurations on BNG remain unchanged in order to interwork with Cisco NCS 5000 Series (NCS 5002, in this example) nV satellite.

For information on nV satellite configuration, see the *Configuring the Satellite Network Virtualization* chapter in the *nV System Configuration Guide for Cisco ASR 9000 Series Routers* located [here](#).

Configuration Example

The configurations given here for access-tracking in BNG2-1 are to be repeated for BNG2-2 as well.

```
RP/0/RSP0/CPU0:router(config)# interface TenGigE100/0/0/1
RP/0/RSP0/CPU0:router(config-if)# bundle id 100 mode on
RP/0/RSP0/CPU0:router(config-if)# commit

RP/0/RSP0/CPU0:router(config)# track ACCESS100
RP/0/RSP0/CPU0:router(config-track)# type line-protocol state
RP/0/RSP0/CPU0:router(config-track-line-prot)# interface Bundle-Ether100
RP/0/RSP0/CPU0:router(config-satellite)# commit
```

You must also configure BNG on the sub-interfaces of Bundle-Ether100 and add the BNG access-interface under the SRG group.

Running Configuration

```
interface TenGigE100/0/0/1
  bundle id 100 mode on

track ACCESS100
  type line-protocol state
  interface Bundle-Ether100
  !
```

Related Topics

- [Geo Redundancy Overview, on page 2](#)
- [Setting up BNG Subscriber Redundancy Group, on page 10](#)
- [BNG Interoperability](#)

Geo Redundancy with Multihoming Ethernet Virtual Private Networks

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
Geo Redundancy with Ethernet Virtual Private Network	Release 7.11.1	<p>In multi-homing EVPN access networks connected across different geographic locations, you can now deploy BNG Geo redundancy on BNG devices. In the event of a failover, the backup BNG device synchronizes subscriber session states at regular intervals. This guarantees seamless load balancing and failover between devices, and provides uninterrupted connectivity. Previously, BNG Geo redundancy was only available for non-EVPN access networks.</p> <p>This feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> The src-driven keyword is added to the ethernet-segment service-carving command. The vlan-aware keyword is added to the neighbor evpn command. <p>YANG Data Model:</p> <ul style="list-style-type: none"> New Xpaths for <code>Cisco-IOS-XR-l2vpn-cfg.yang</code> <p>(see GitHub, YANG Data Models Navigator)</p>

Geo redundancy with EVPN allows integration of EVPN headend VLAN aware functionality within a BNG system to extend VLANs across the broadband access network. This capability enables service providers to offer advanced Layer 2 services to broadband customers, including seamless integration of customer VLANs and efficient resource utilization. You can enable geo redundancy with EVPN headend VLAN aware using the **src-driven** command.

Enabling SRG-driven configuration allows the distribution of the SRG active role for subinterfaces across BNG nodes (PE3, PE4). EVPN learns and translates each subinterface's active or backup role to the designated forwarder (DF) role used in EVPN signaling.

When SRG-driven configuration is enabled, the SRG role for the subinterfaces is learned and translated to the designated forwarding (DF) role.

In the above topology, the subinterface PW-Ether1002.6 on PE3 is set to the active node by SRG, and this attribute is translated to EVPN. PE3 functions as the primary node for VLAN 2 on PW-Ether1002.6 and advertises the Ethernet Auto-Discovery (EAD) route with the DF role set to primary.

Similarly at PE4, subinterface PW-Ether1002.6 serving as the SRG backup is translated to EVPN and advertises the EAD route with the DF role set to backup.

Traffic in the access to core direction from CE1 will use the primary path for the VLAN 2 through PE3 until an SRG failover occurs and PE4 starts signaling primary. PE1 or PE2 may also fail over to PE4 due to losing reachability to PE3.

The traffic from CE1 passes through either PE1 or PE2, depending on the EVPN ethernet-segment peering mode at PE1 and PE2. However, both PE1 and PE2 receive the same EVPN routes from PE3 and PE4, resulting in the traffic taking the same path regardless of which PE it passes through.

Restrictions for Geo Redundancy with EVPN

- An Ethernet-Segment Identifier (ESI) must be configured on the PWHE main interface.
- When there are multiple remote nexthops available, PWHE does not support ECMP. Instead, it only binds itself to a single remote nexthop without considering the option of using ECMP for load balancing or redundancy purposes.
 - The nexthop selection prioritizes the earliest incoming nexthop. This means that if a new remote nexthop is discovered while the PWHE is already associated with a specific nexthop, it remains bound to the current nexthop as long as it remains valid. es. This approach guarantees that introducing a new nexthop doesn't disrupt the traffic flow for existing services.
 - The PWHE uses the first path received may be overridden by configuration with Preferred-Nexthop Lowest-IP, Highest-IP.
- If there are multiple remote nexthops conflicts during the load phase of the In-Service Software Upgrade (ISSU) process, EVPN halts the process. It is crucial to address this condition during normal operating conditions (steady-state) before ISSU to ensure proper functioning. Before ISSU, existing nexthops were used by EVPN, and conflicting next-hops were rejected. However, during ISSU, there may be a rearrangement of nexthops that could lead to traffic loss or changes in traffic patterns.

Configure Geo Redundancy with EVPN

Let's consider the example where PE3 and PE4 routers running BNG headend configured with VLAN aware.

There are two SRG groups each with one subinterface belonging to the main port (PW-Ether 1002) to show that one subinterface can be DF and the other can be non-DF.

Perform the following tasks on PE3 and PE4:

- Enable EVPN-Geo Redundancy on PE3 and P34 using the **srg-driven** command:

- Configure VLAN aware PWHE on PE3 and PE4 using the **vlan-aware** command.
- Configure SRG on PE3 and PE4.

Configuration Example

Enable EVPN-Geo Redundancy on PE3 and PE4 using the **srg-driven** command:

```
Router#configure
Router(config)#i
Router(config-evpn)#interface PW-Ether1002
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)#identifier type 0 00.10.02.00.00.00.00.10.02
Router(config-evpn-ac-es)#service-carving preference-based
Router(config-evpn-ac-es-sc-pref)# srg-driven
Router(config-evpn-ac-es-sc-pref)#commit
Router(config-evpn-ac-es-sc-pref)#root
```

Configure VLAN aware PWHE on PE3 and PE4:

```
Router(config)#l2vpn
Router(config-l2vpn)#xconnect group evpn-headend
Router(config-l2vpn-xc)#p2p headend-va-1002
Router(config-l2vpn-xc-p2p)#interface PW-Ether 1002
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1002 service vlan-aware
Router(config-l2vpn-xc-p2p)#root
```

Configure SRG on PE3 and PE4:

```
Router(config)#subscriber redundancy
Router(config-subscr-red)#source-interface GigabitEthernet0/0/0/0
Router(config-subscr-red)#group 6
Router(config-subscr-red-group)#peer 10.10.10.5
Router(config-subscr-red-group)# interface-list
Router(config-subscr-red-grp-intf)#interface PW-Ether1002.6 id 206
Router(config-subscr-red-grp-intf)#exit
Router(config-subscr-red-group)#group 207
Router(config-subscr-red-group)#peer 10.10.10.5
Router(config-subscr-red-group)#interface-list
Router(config-subscr-red-grp-intf)#interface PW-Ether1002.7 id 207
Router(config-subscr-red-grp-intf)#commit
Router(config-subscr-red-grp-intf)#root
```

Running Configuration

Here's the running configuration on PE3 and PE4:

```
evpn
 interface PW-Ether1002
   ethernet-segment
   identifier type 0 00.10.02.00.00.00.00.10.02
   service-carving preference-based
   srg-driven
   !
 !
l2vpn
 xconnect group evpn-headend
 p2p headend-va-1002
```

```

        interface PW-Ether1002
        neighbor evpn evi 1002 service vlan-aware
        !
        !
        !
subscriber
  redundancy
    source-interface GigabitEthernet0/0/0/0
    group 6
      peer 10.10.10.5
      interface-list
        interface PW-Ether1002.6 id 206
      !
    !
    group 207
      peer 10.10.10.5
      interface-list
        interface PW-Ether1002.7 id 207
      !
    !
  !
!
!
!

```

Verification

On PE3, verify the SRG role.

```
RP/0/0/CPU0:PE3#show subscriber redundancy group 6
```

```
Subscriber Redundancy Group ID: 6
Description : <<not-configured>>

Status           : Enabled
Init-Role        : None
Negotiated-Role  : Master          Current-Role     : Master
...

```

```
RP/0/0/CPU0:PE3#show subscriber redundancy group 207
```

```
Subscriber Redundancy Group ID: 207
Description : <<not-configured>>

Status           : Enabled
Init-Role        : None
Negotiated-Role  : Slave          Current-Role     : Slave
...

```

Verify the EVPN state and the DF role in the output. This example shows that the the subinterface PW-Ether1002.6 acts as primary and PW-Ether1002.7 acts a backup.

```
Routerrouter#show evpn ethernet-segment interface PW-ether 1002 private
Tue Oct 24 15:13:20.728 EDT
```

Legend:

```

B - No Forwarders EVPN-enabled,
C - MAC missing (Backbone S-MAC PBB-EVPN / Grouping ES-MAC vES),
RT - ES-Import Route Target missing,
E - ESI missing,
H - Interface handle missing,
I - Name (Interface or Virtual Access) missing,
M - Interface in Down state,
O - BGP End of Download missing,
P - Interface already Access Protected,
Pf - Interface forced single-homed,
R - BGP RID not received,
S - Interface in redundancy standby state,
X - ESI-extracted MAC Conflict

```

SHG - No local split-horizon-group label allocated
 Hp - Interface blocked on peering complete during HA event
 Rc - Recovery timer running during peering sequence

Ethernet Segment Id	Interface	Nexthops
0000.1002.0000.0000.1002	PE1002	192.168.0.4 192.168.0.5

```

ES to BGP Gates : Ready
ES to L2FIB Gates : Ready
Main port :
  Interface name : PW-Ether1002
  Interface MAC : 021b.37c5.0807
  IfHandle : 0x00000130
  State : Up
  Redundancy : Not Defined
ESI ID : 3
ESI type : 0
  Value : 0000.1002.0000.0000.1002
ES Import RT : 0010.0200.0000 (from ESI)
Source MAC : 0000.0000.0000 (N/A)
Topology :
  Operational : MH, Anycast mode
  Configured : Anycast Single-active (default)
Service Carving : Preferential
  Config Weight : 32767
  Oper Weight : 32767
  Non-Revertive : Disabled, Inactive
  Access Driven : Disabled
  SRG Driven : Enabled
  Multicast : Disabled
Convergence :
Peering Details : 2 Nexthops
  192.168.0.4 [PREF:P:7fff:T][1]
  192.168.0.5 [PREF:P:7fff:T][2]
Service Carving Synchronization:
  Mode : NONE
  Peer Updates :
    192.168.0.4 [SCT: N/A]
    192.168.0.5 [SCT: N/A]
Service Carving Results:
  Forwarders : 1
  Elected : 0
  Not Elected : 0
EVPN-VPWS Service Carving Results:
  Primary : 1
  Backup : 0
  Non-DF : 0
MAC Flushing mode : STP-TCN
Peering timer : 3 sec [not running]
Recovery timer : 30 sec [not running]
Carving timer : 0 sec [not running]
Revert timer : 0 sec [not running]
HRW Reset timer : 5 sec [not running]
AC Debounce timer : 3000 msec [not running]
Local SHG label : 400004
Remote SHG labels : 1
  500004 : nexthop 192.168.0.5
Access signal mode: Unsupported

Object: EVPN ES
Base info: version=0xdbdb0007, flags=0x0, type=7, reserved=0
EVPN ES event history [Num events: 49]

```

```

Time          Event          Flags          Flags
====          =====          =====          =====
Oct 24 15:10:33.600 ES State Change 00000000 00000003 - -
Oct 24 15:10:33.600 ES DB Bind      00000000 00000001 - -
Oct 24 15:10:33.600 Create         00000000 00000001 - -
Oct 24 15:10:33.600 API Config Ifname Add 00000000 00000001 - -
Oct 24 15:10:33.600 ES DB Bind      00000000 00010001 - -
    
```

EVPN ES state chart history [Total events: 5]

```

Time          Event          Value(s)          Value(s)
====          =====          =====          =====
Oct 24 15:10:37.184 ES Weight Change [many] 0x00000000 0x00000000
Oct 24 15:12:22.656 ES Weight Change 0x00020000 0x04000000
Oct 24 15:12:22.656 ES Recovery Start 0x00000000 0x04000000
Oct 24 15:12:52.608 ES Recovery Complete 0x00020000 0x00000000
Oct 24 15:12:52.608 ES Weight Change 0x00000000 0x00000000
    
```

EVPN ES Statistics

```

          |Adv          |Wdw
          |Cnt Last Time          Last Arg| Cnt Last Time          Last Arg
RT| 1 24/10 15:10:37.194 00000000| 0 00000000
LocalBMAC| 0 00000000| 0 00000000
ESI| 1 24/10 15:12:52.649 00000001| 0 00000000
EAD/ES| 1 24/10 15:12:55.650 000a0add| 0 00000000
EAD/EVI| 0 00000000| 0 00000000
MST-AG VPW| 0 00000000| 0 00000000
DF ElectFW| 1 24/10 15:10:36.086 00000000| 0 00000000
UpdateMAC| 0 00000000| 0 00000000
MacFlushPE| 0 00000000| 0 00000000
MacFlushCE| 0 00000000| 0 00000000
Instance| 0 00000000| 0 00000000
MP Info| 0 00000000| 0 00000000
MipConvert| 0 00000000| 0 00000000
FibReachbl| 0 00000000| 0 00000000
    
```

```

-----
Diagnostic ESI : N/A          Interface Name : N/A
Diagnostic Ifh : 0x00000000
Diagnostic Flag: 0x0000c043  DiagnosticES-RT: 0000.0000.0000
Port Key       : 0x00001a84  MAC winner       : 1
Number of EVIs : 1
Recovery Timer : 30 (global) Peering Timer   : 3 (global)
Carving Timer  : 0 (global)
AC Dbnc Timer  : 3000 (global)
RT Advertised  : 1          ESI Advertised  : 1
MSTi state mask: 0x0000    HRW MSTi Set   : 0x7
ES EAD Pulse   : 0          MP Advertised  : 0
NVE AnycastVTEP: 0         NVE Ingr-Replic: 0
Peering Done   : 1          Carving Done   : 1
Peering Chkpt  : 0          Peering Block  : 0
Rest Chkpt ES St 0         PFI Down       : No
RedundFWD      : 0          Inval NH       : 0
CarvReadvertise: 0         SubIfIMRegistrd: 1
IfIMCRegistrd  : 0         IfPFIRRegistrd : 1
CheckpointStale: 0         CheckpointRcvrd: 0
MSTIFlushFlood: 0         MIPConvertPend : 0
MSTIO Adv      : 0          MP Adv Blk'd   : 1
Checkpoint Info:
IF Type        : 1          MSTi Mask       : 0x0
Peering Done   : 1
Chkpt ES stat  : 2
NextHop Info :
192.168.0.4 [PREF:P:7fff] [1] [0000.0000.0000]
    
```

```

    192.168.0.5 [PREF:P:7fff][2][0000.0000.0000]
Forced State   : Invalid
VI member info : not present
ES EAD Update  :
Num RDs       : 1

RD: 192.168.0.4:3, Num RTs: 1      RT List:
    100:1002,

Sub-interface: PW-Ether1002.6, ifh: 0x000001b0, State is Up
MTU: 1518, BD ID: 4294967295, VRF ID: 0x60000000, VLAN ID: 0x00810600
OrigVLANID: 0x00000006, RewritVLANID: 0x00000000, Flags: 0x0
MAC: 0002.001b.0037, IPv4: 10.2.6.1, IPv6: ::
EVPN-Mapped SRG DF Role: Primary
Sub-interface: PW-Ether1002.7, ifh: 0x000001d0, State is Up
MTU: 1518, BD ID: 4294967295, VRF ID: 0x60000000, VLAN ID: 0x00810700
OrigVLANID: 0x00000007, RewritVLANID: 0x00000000, Flags: 0x0
MAC: 0002.001b.0037, IPv4: 10.2.7.1, IPv6: ::
EVPN-Mapped SRG DF Role: Backup

```

On PE4, verify the SRG role and verify the DF role. In this example, the PE3 acts as a active and the PE4 acts like a standby.

```

:
Router#show subscriber redundancy group 6
Wed Sep 20 11:44:58.794 EDT
Subscriber Redundancy Group ID: 6
  Description : <<not-configured>>

Status           : Enabled
Init-Role        : None
Negotiated-Role : Slave                Current-Role  : Slave
...

RP/0/0/CPU0:PE4#show subscriber redundancy group 207
Wed Sep 20 11:45:17.904 EDT
Subscriber Redundancy Group ID: 207
  Description : <<not-configured>>

Status           : Enabled
Init-Role        : None
Negotiated-Role : Master              Current-Role  : Master
...

Router#show evpn ethernet-segment interface PW-Ether 1002 detail
Legend:
B - No Forwarders EVPN-enabled,
C - MAC missing (Backbone S-MAC PBB-EVPN / Grouping ES-MAC vES),
RT - ES-Import Route Target missing,
E - ESI missing,
H - Interface handle missing,
I - Name (Interface or Virtual Access) missing,
M - Interface in Down state,
O - BGP End of Download missing,
P - Interface already Access Protected,
Pf - Interface forced single-homed,
R - BGP RID not received,
S - Interface in redundancy standby state,
X - ESI-extracted MAC Conflict
SHG - No local split-horizon-group label allocated
Hp - Interface blocked on peering complete during HA event
Rc - Recovery timer running during peering sequence

```

```

Ethernet Segment Id      Interface                                Nexthops

```

```

-----
0000.1002.0000.0000.1002 PE1002          192.168.0.4
                                          192.168.0.5
ES to BGP Gates      : Ready
ES to L2FIB Gates   : Ready
Main port            :
  Interface name     : PW-Ether1002
  Interface MAC      : 02a5.d7aa.c907
  IfHandle           : 0x00000130
  State              : Up
  Redundancy         : Not Defined
ESI ID               : 3
ESI type             : 0
  Value              : 0000.1002.0000.0000.1002
ES Import RT        : 0010.0200.0000 (from ESI)
Source MAC          : 0000.0000.0000 (N/A)
Topology            :
  Operational        : MH, Anycast mode
  Configured         : Anycast Single-active (default)
Service Carving     : Preferential
  Multicast          : Disabled
Convergence         :
Peering Details     : 2 Nexthops
  192.168.0.4 [PREF:P:7fff:T]
  192.168.0.5 [PREF:P:7fff:T]
Service Carving Synchronization:
  Mode               : NONE
  Peer Updates      :
    192.168.0.4 [SCT: N/A]
    192.168.0.5 [SCT: N/A]
Service Carving Results:
  Forwarders        : 1
  Elected           : 0
  Not Elected       : 0
EVPN-VPWS Service Carving Results:
  Primary           : 0
  Backup            : 1
  Non-DF            : 0
MAC Flushing mode  : STP-TCN
Peering timer      : 3 sec [not running]
Recovery timer     : 30 sec [not running]
Carving timer      : 0 sec [not running]
Revert timer       : 0 sec [not running]
HRW Reset timer    : 5 sec [not running]
Local SHG label    : 500000
Remote SHG labels  : 1
  400000 : nexthop 192.168.0.4
Access signal mode: Unsupported
    
```

Router#show l2vpn xconnect

Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
 SB = Standby, SR = Standby Ready, (PP) = Partially Programmed,
 LU = Local Up, RU = Remote Up, CO = Connected, (SI) = Seamless Inactive

XConnect			Segment 1	Segment 2
Group	Name	ST	Description	Description
evpn-headend	headend-va-1002	UP	PE1002	VA 1002,192.168.0.1 EVPN 1002,6,192.168.0.1

SB

EVPN 1002,7,192.168.0.1

UP

Geo Redundancy Features

Peer Route Disable

Peer route disable is an enhancement in BNG geo redundancy whereby the user can disable the route on geo redundancy hot-standby peer. This disabling is so that the subscriber routes are not installed in the RIB even when the subscriber sessions are brought up on the standby peer. The subscriber routes are inserted into the RIB only when the BNG Geo-Redundancy state of peer changes to active. This ensures that only the primary BNG, and not the subordinate BNG, routes the subscriber traffic in a scenario where access-interface is up on the standby peer. By disabling the routes, the hot-standby mode in BNG geo redundancy does not mandate the access-interface to be down on the standby peer any more.

To enable this feature, use the **peer route disable** command in subscriber redundancy group configuration mode.

Configuration Example

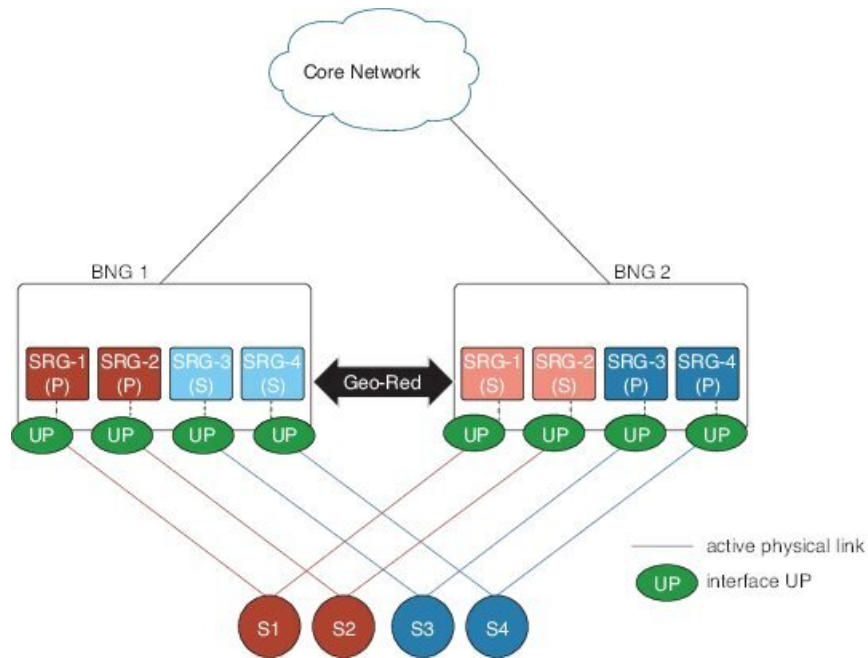
```
RP/0/RSP0/CPU0:router(config)# subscriber redundancy group 110
RP/0/RSP0/CPU0:router(config-subscr-red-group)# peer route-disable
```

Active-active Session Support for Geo Redundancy

Active-active session support for BNG geo redundancy is an enhancement where a subscriber redundancy group (SRG) can be primary on a BNG node while being subordinate on the pair BNG node, and simultaneously another SRG can be primary on the pair BNG node while being subordinate on the primary BNG node. So, a BNG node can be a primary for one SRG and at the same time subordinate for another SRG. This feature provides better load balancing for subscriber sessions across both BNG nodes.

In the case of an active-active scenario, the L2 path from the subscriber CPE towards both BNG nodes is ready to forward packets. Or in other words, the access interface protocol is UP at both BNG nodes.

Figure 11: Active-active Session for BNG Geo Redundancy



Note The pool name and the address range must be unique for each SRG group in both BNGs for the active-active configuration.

Address Pool Usage Synchronisation in BNG Geo Redundant Active-Active Nodes

The BNG geo redundancy active-active topology contains multiple SRG groups between a pair of BNG nodes. Since each BNG node will be in Primary role for some SRG groups and Subordinate role for the remaining SRG groups, it provides a load balancing of subscribers across BNG nodes.

But this mechanism brings in a lot of complexity when the IP address pool is locally configured on the BNGs, such as in the case of PPPoE, DHCP Server & ND Slaac configuration. The same address pool cannot be shared by multiple interfaces that belong to different SRG groups running in different roles. As a result, the operational overhead is high and you need to manage multiple address pools and address ranges, which might also result in an inefficient address usage.

To overcome this problem, BNG introduces address pool usage synchronisation in geo redundant active-active nodes. This feature enables the use of the same address-pool for a two BNG geo redundant topology. The pool-server running in the SERG Primary BNG assigns the IP addresses for the SRG Primary subscribers. The pool-server running in SERG Subordinate BNG acts as the proxy.

Restrictions for Address Pool Usage Synchronisation in BNG Geo Redundant Active-Active Nodes

Address pool usage synchronisation in BNG geo redundant active-active nodes is subjected to the below restrictions:

- An SERG group cannot have both the interface configuration and pool configuration. Only one of the configurations can be applied on a specific group.
- A specific address pool should be configured only under one SERG group.

Configure Address Pool Usage Synchronisation in BNG Geo Redundant Active-Active Nodes

Configuration Steps

Consider two BNG geo redundant active-active nodes. The configuration steps for configuring pools, on both the SERG primary and SERG subordinate BNG nodes, are shown below:

```
Router# configure terminal
Router(config)# pool vrf default ipv4 p1
Router(config-pool-ipv4)# network 10.10.0.1/16
Router(config-pool-ipv4)# exit
Router(config)# pool vrf default ipv6 p2
Router(config-pool-ipv4)# address-range 2001::10 2001::ffff
Router(config-pool-ipv4)# commit
```

The configuration steps for the BNG node in the SERG primary role is shown below:

```
Router(config)# session-redundancy
Router(config-sess-red)# source-interface mgmtEth 0/RP0/CPU0/0
Router(config-sess-red)# group 100
Router(config-sess-red-group)# preferred-role master
Router(config-sess-red-group)# peer 10.10.10.2
Router(config-sess-red-group)# pool-list pool-name p1
Router(config-sess-red-group)# pool-list pool-name p2
Router(config-sess-red-group)# commit
```

The configuration steps for the BNG node in the SERG subordinate role is shown below:

```
Router(config)# session-redundancy
Router(config-sess-red)# source-interface mgmtEth 0/RP0/CPU0/0
Router(config-sess-red)# group 100
Router(config-sess-red-group)# preferred-role slave
Router(config-sess-red-group)# peer 10.10.10.1
Router(config-sess-red-group)# pool-list pool-name p1
Router(config-sess-red-group)# pool-list pool-name p2
Router(config-sess-red-group)# commit
```

Running Configuration

The running configuration for the pools on both the SERG primary and SERG subordinate BNG nodes is shown below:

```
pool vrf default ipv4 p1
  network 10.10.0.1/16
!
pool vrf default ipv6 p2
  address-range 2001::10 2001::ffff
!
```

The running configuration for the BNG node in the SERG primary role is as follows:

```
session-redundancy
  source-interface MgmtEth0/RP0/CPU0/0
```

```

group 100
  preferred-role master
  peer 10.10.10.2
  pool-list pool-name p1
  pool-list pool-name p2
!
!

```

The running configuration for the BNG node in the SERG subordinate role is as follows:

```

session-redundancy
  source-interface MgmtEth0/RP0/CPU0/0
  group 100
    preferred-role slave
    peer 10.10.10.1
    pool-list pool-name p1
    pool-list pool-name p2
!
!

```

Verification

Pool specific information can be verified using the below command:

```

Router# show pool ipv4 name p1 remote-info
SERG Info:-
-----
Role           : Master
PeerDown       : False

```

Addresses allocated to Remote:-

```

-----
150.0.5.100  PPP *
150.0.5.101  PPP *
150.0.5.103  PPP *
150.0.5.104  PPP *
150.0.5.105  PPP *

```

Addresses received from Remote:-

```

-----
NONE

```

Pool information for the SERG can be verified using the below command:

```

Router# show session-redundancy group
Session Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred Slave
             H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S        : Peer Status
             I - Initialize, Y - Retry, X - Cleanup, T - Connecting
             L - Listening, R- Registered, C - Connected, E - Established
I/F-P Count: Interface or Pool Count
SS Count   : Session Count

```

```

-----
Node Name   | Group ID | Role | Flags | Peer Address | P/S | I/F-P Count
| SS Count  | Sync Pending
-----
0/RP0/CPU0 |          | 100 | Master | EMH- 10.10.10.2 | T   | 1
| 1         | 0
-----

```

```

Session Summary Count(Master/Slave/Total): 1/0/1

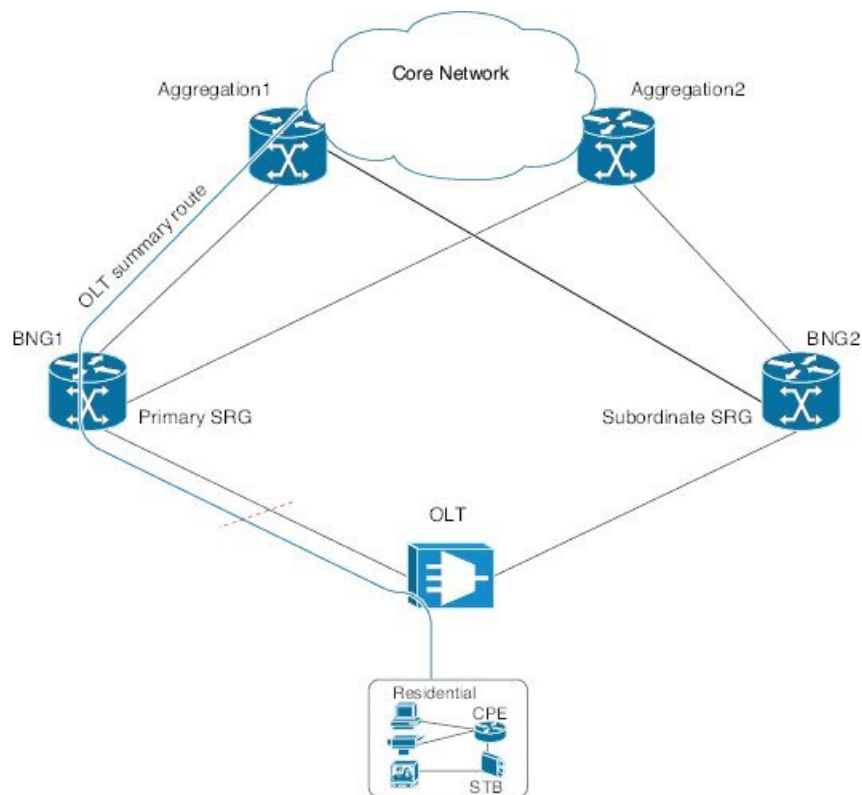
```

State Control Route for Geo Redundancy

State control route is a loss-of-signal (LOS) based solution for Broadband Network Gateway (BNG) geo redundancy, where the route advertisement to the core network is controlled based on the role of the subscriber redundancy group (SRG). Only a summary of subscriber routes from the primary SRG group get advertised to the core network with the help of routing protocols. The routes from the subordinate SRG are not advertised. When a switchover happens, the new subordinate SRG withdraws the previously advertised routes from the core. This solution provides a sub-second traffic convergence and prevents traffic back hole. This feature is mainly useful in BNG deployments where optical line terminals (OLTs) do not support access protocols (like BFDs, CFM, and so on) to detect link failures in the access or core network.

A tag value is used in order to filter out the summary route from the subscriber routes. So, only the routes having that particular tag get redistributed through the routing protocol and get advertised to the core. The remaining routes are filtered out.

Figure 12: State Control Route for BNG Geo Redundancy



Multiple State Control Routes for Each SRG

The state control route feature in BNG is extended to have multiple state control routes for each subscriber redundancy group (SRG), as opposed to allowing only single state control route for each SRG. This functionality is available for PPPoE and IPoE subscribers. Users can also add summary route for a subscriber for a specific VRF, rather than limiting it to a default VRF. This feature thereby allows service providers to terminate multiple subnets of subscribers in a particular SRG.

Configure State Control Route for Geo Redundancy

To enable state control route for BNG geo redundancy, use **state-control-route** command in subscriber redundancy group configuration mode.

To enable multiple state control routes for each subscriber redundancy group, use this command along with the **vrf** option which is available only from Cisco IOS XR Software Release 6.3.1 and later.

The **vrf** and **tag** keywords are optional parameters. If the **vrf** option is not specified, then the routes are added in the default VRF, that is, global routing table.

A maximum of 30 state control routes can be added in each subscriber redundancy group (SRG), with a limit of 10 state control routes for each route type. That is, user can have a maximum 10 IPv4 routes, 10 IANA routes and 10 IAPD routes in each SRG. In these 30 routes, user can have routes in either the default or the non-default VRF.

The route-policy with the respective tag (that is, tag 1 in this example) must be defined prior to configuring the **state-control-route**.

Configuration Example for State Control Route with Default VRF

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#subscriber redundancy group 110
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv4 192.0.2.0/9 tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv6 iana 2001:DB8::/32
tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv6 iapd 2001:DB8:1::1/32
tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#commit
```

Running Configuration for State Control Route with Default VRF

```
/* State control route configuration */
subscriber redundancy group 110
state-control-route ipv4 192.0.2.0/9 tag 1
state-control-route ipv6 iana 2001:DB8::/32 tag 1
state-control-route ipv6 iapd 2001:DB8:1::1/32 tag 1
!

/* Route-policy configuration */
route-policy SUB_ROUTES
if tag is 1 then
    pass
done
endif
end-policy

/* Routing protocol configuration */
router ospf core
router-id 11.11.11.11
redistribute subscriber route-policy SUB_ROUTES
address-family ipv4 unicast
area 0
interface Loopback2
!
interface GigabitEthernet0/0/0/10
!
!
```

Configuration Example for Multiple State Control Routes with Specific VRF

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#subscriber redundancy group 110
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv4 192.0.2.0/9 vrf
vrf1 tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv6 iana 2001:DB8::/32
vrf vrf1 tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#state-control-route ipv6 iapd 2001:DB8:1::1/32
vrf vrf1 tag 1
RP/0/RSP0/CPU0:router(config-subscr-red-group)#commit
```

Running Configuration for Multiple State Control Routes with Specific VRF

```
/* Multiple State control route configuration with VRF, vrf1 */
subscriber redundancy group 110
state-control-route ipv4 192.0.2.0/9 vrf vrf1 tag 1
state-control-route ipv6 iana 2001:DB8::/32 vrf vrf1 tag 1
state-control-route ipv6 iapd 2001:DB8:1::1/32 vrf vrf1 tag 1
!
```

Verification

Use this command to display the summary routes:

```
RP/0/RSP0/CPU0:router#show route subscriber

A    192.0.2.0/9 [1/0] via 0.0.0.0, 1w4d
```

Subscriber Redundancy Group Revertive Timer

The subscriber redundancy group (SRG) revertive timer feature is an enhancement in BNG geo redundancy where, based on certain conditions, the primary BNG node for which the preferred role is set as primary, automatically regains the primary role (from subordinate role) after an SRG fail-over. An auto-revertive timer starts when the preferred primary BNG becomes subordinate due to SRG fail-over and when access-tracking and core-tracking are restored. When the timer expires, the preferred primary BNG regains the primary role. This switch back to the preferred primary role is required, as the new primary SRG may not be equipped to handle the entire subscriber load in the case of a fail-over.

The SRG revertive timer starts only if all these conditions are met:

- The preferred-role of the BNG is set as primary.
- The current role of the BNG node is subordinate.
- Access-tracking is UP.
- Subscriber Redundancy Group (SRG) peering is UP.

If SRG peering is down, the role of the BNG node automatically switches back from subordinate to primary, without even starting the SRG revertive timer.

To set the SRG revertive timer, use the **revertive-timer** command in subscriber redundancy configuration mode.

Running Configuration

```
subscriber
 redundancy
  revertive-timer 5 maximum 20
!
```

Subscriber Redundancy Group-aware IPv6 Neighbor Discovery

Subscriber Redundancy Group-aware (SRG-aware) IPv6 Neighbor Discovery (ND) is an enhancement in BNG geo redundancy where, the Router Advertisement (RA) message in response to the IPv6 ND message for IPv6 deployments, is sent based on the SRG role of the parent interface. Only the primary node sends out RA message in response to the IPv6 ND message and brings up the session. The RS (Router Solicitation) or Neighbor Solicitation (NS) message is dropped on the subordinate node, but the sessions still come up in that subordinate node. That way, the routes are not advertised to the core from the standby node.

IPv6 ND sends RA on every subscriber interface. It listens to the SRG state and then stops generating protocol messages based on the SRG state. When SRG state is primary, periodic RA is performed and when SRG state is back to subordinate, the periodic RA is stopped.


Verification

```
Router# show ipv6 nd idb interface <> detail location 0/RSP0/CPU0
...
Subscriber status flag: 0x18, Supressed cache learning: FALSE
BNG nud: Disabled, Master Node: 0/1/CPU0(0x10)
Global Mac Accounting: Disabled, IDB Mac Accounting : Disabled, Marked: No
Notfn sent to iedge - Up: Yes (Apr 4 18:12:07), Down: No
                               Update: No
                               Last notif reason:No prefix available
SRG Stby Role : TRUE
```

Peer-to-peer Traffic Flow with BNG Geo Redundancy

Peer-to-peer traffic flow is an enhancement in BNG Geo Redundancy where subscribers in different subscriber redundancy groups (SRGs) in the primary and subordinate nodes can send traffic to each other through the BNG nodes. This is feasible as the primary SRGs from both the BNG nodes advertise the respective summary routes to the core.

Figure 13: Peer-to-peer Traffic Flow with BNG Geo Redundancy



Suppose, subscriber S1 is connected to BNG1 and BNG2 through OLT1. Similarly, subscriber S2 is connected to BNG1 and BNG2 through OLT2. S1 is associated with SRG1 and S2 is associated with SRG2. SRG1 is primary in BNG1 and subordinate in BNG2. Similarly, SRG2 is primary in BNG2 and subordinate in BNG1.

In this scenario, the subscriber routes are added to the main table as well as to SRG VRF table in the primary node. Whereas, in the subordinate, the routes are added only to the SRG VRF table. The primary SRG1 in BNG1 advertises the summary route of S1 to the core. Similarly, the primary SRG2 in BNG2 advertises the summary route of S2 to the core. That way, both routes are reachable by each other through the BNG peer nodes.

To enable this feature, use the **enable-fast-switchover** command in subscriber redundancy group configuration mode.

Running Configuration

```
subscriber
 redundancy
  group 110
  enable-fast-switchover
  !
  !
  !
end
```

Accounting Trigger Cause for Geo Redundancy

A new Cisco-Attribute Value Pair (AVP), **Acct-Trigger-Cause**, is introduced to send the reason of accounting start and accounting stop messages triggered during an SRG switchover. The accounting stop record, sent from the old primary BNG node, and the accounting start record, sent from the new primary BNG node, specify the Acct-Trigger-Cause to be **nas-switchover**. This, in turn, helps the backend servers to identify the reason for the new accounting trigger thereby preventing the existing accounting records of the subscriber sessions from getting deleted.

This is a sample output of the session accounting stop record on old primary BNG node:


```

RADIUS: Vendor,Cisco      [26]  41
RADIUS: Cisco AVpair      [1]   35      acct-trigger-cause=nas-switchover
RADIUS: Acct-Status-Type  [40]  6       Start[1]

```

This is a sample output of the session accounting start record on new primary BNG node:

```

RADIUS: Vendor,Cisco      [26]  41
RADIUS: Cisco AVpair      [1]   35      acct-trigger-cause=nas-switchover
RADIUS: Acct-Status-Type  [40]  6       Stop[2]

```

SRG Support for BNG SLAAC Sessions

BNG introduces the support for subscriber redundancy group (SRG) for Stateless Address Auto-Configuration (SLAAC) sessions, wherein the subordinate BNG router allocates the same Neighbor Discovery (ND) prefix as that of primary BNG router to the subscriber. This feature ensures a seamless traffic flow for SLAAC sessions in the event of a BNG switchover.

SLAAC is an IP address-assignment model in which the hosts generate their own addresses using a combination of local and router-advertised information. Routers advertise prefixes that identify the subnet(s) associated with a link. Hosts generate a unique identifier for the interface on a subnet. These two combine to form an IP address. For more details on geo redundancy on PPPoE sessions, IPv6 address assignment and SLAAC, see the *Related Topics* section below.

When the SLAAC session comes up with IPv6 prefix from ND prefix pool on subordinate SRG, IPv6 ND reserves the prefix from DAPS pool. Once the prefix reservation is successful, SLAAC session comes up on the subordinate SRG.

When SRG peer route disable feature is enabled, IPv6-ND brings up SLAAC session on the subordinate SRG without adding the subscriber IPv6 route to RIB. When the SRG role of the parent interface changes from subordinate to primary, IPv6-ND adds the subscriber IPv6 route to RIB. When the SRG role of the parent interface changes from primary to subordinate, IPv6-ND removes the subscriber IPv6 route from RIB.



Note For the seamless working of SRG for BNG SLAAC sessions, you must ensure that the configurations and IP addresses are similar on both primary and subordinate BNG routers, as it is for generic SRG functionality.

Related Topics

- [Geo Redundancy for PPPoE Sessions, on page 32](#)
- [DHCPv6 Address or Prefix Pool](#)
- [IP Address Assignment for Clients](#)
- [Scenario 1: SLAAC-Based Address Assignment](#)

SRG Support for Static Sessions

BNG introduces the support for subscriber redundancy group (SRG) for static sessions, wherein all traffic belonging to a particular VLAN sub-interface is treated as a single session. This feature ensures a seamless traffic flow for static sessions in the event of a BNG switchover.



Note Starting with Cisco IOS XR Release 7.10.2 and 7.11.1, SRG supports static IP addresses allocated by the radius for IPoE dual stack sessions.

Restrictions for SRG Support for Static Sessions

SRG support for static sessions in BNG is subjected to these restrictions:

- The bundle mac-address which is configured on the primary and subordinate BNG must be the same in order for SRG to work for static interface sessions. SRG virtual MAC address (vMAC) functionality is not supported with static interface sessions.
- IPv6-ND DAD must be disabled on the access interface for SRG to work for IPv6 static interface sessions.
- The command **clear subscriber session** is not supported on the subordinate SRG.
- SRG disable and enable operations are not supported for static interface sessions.
- SRG enable-fast-switchover and peer route-disable are not supported for static interface sessions.
- SRG support for static sessions is applicable for both RP and LC based subscribers.
- RPFO is not supported on primary BNG when static sessions are coming up on subordinate BNG.
- RPFO can lead to an SRG peering flap.
- Restart of the srg_agt process is not supported if some of the groups are disabled or if some of the groups do not have a peering group.

Configure Subscriber Redundancy Group for Static Sessions

Configuration Steps

The below section shows how to configure SRG for static sessions.

```
Router# config terminal
Router(config)# interface GigabitEthernet0/0/0/0
Router(config-if)# bundle id 1 mode on
Router(config-if)# exit
Router(config)# interface Bundle-Ether 1
Router(config-if)# mac-address a1.b1.c1
Router(config-if)# exit
Router(config)# policy-map type control subscriber POLICY1
Router(config-pmap)# event session-start match-all
Router(config-pmap-e)# class type control subscriber class-default do-all
Router(config-pmap-c)# 1 activate dynamic-template IPSUB_TEMPLATE
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# interface Bundle-Ether 1.100
Router(config-subif)# ipv4 address 20.1.1.1 255.255.255.0
```

```

Router(config-subif)# ipv6 address 1001::2/128
Router(config-subif)# ipv6 nd dad attempts 0
Router(config-subif)# encapsulation dot1q 100
Router(config-subif)# service-policy type control subscriber POLICY1
Router(config-subif)# ipsubscriber interface
Router(config-subif)# exit
Router(config)# track abc
Router(config-track)# type line-protocol state
Router(config-track-line-prot)# interface GigabitEthernet0/0/0/0
Router(config-track-line-prot)# exit
Router(config-track)# exit
Router(config)# subscriber
Router(config-subscriber)# redundancy
Router(config-subscr-red)# source-interface GigabitEthernet0/0/0/10
Router(config-subscr-red)# group 1
Router(config-subscr-red-group)# preferred-role master
Router(config-subscr-red-group)# slave-mode hot
Router(config-subscr-red-group)# peer 10.1.1.2
Router(config-subscr-red-group)# access-tracking abc
Router(config-subscr-red-group)# interface-list
Router(config-subscr-red-grp-intf)# interface Bundle-Ether1.100 id 1
Router(config-subscr-red-grp-intf)# commit

```

When bringing up scaled static sessions on subordinate SRG, below procedure must be followed:

1. Shutdown the bundle interface.
2. Configure the access-interfaces with dot1q encapsulation, IP address, subscriber control policy and ipsubscriber interface. All the vlan interfaces should be in ready state.
3. Configure the SRG with interface-list.
4. Bring up the bundle interface

Running Configuration

```

interface GigabitEthernet0/0/0/0
  bundle id 1 mode on
!
interface Bundle-Ether 1
  mac-address a1.b1.c1
!
policy-map type control subscriber POLICY1
  event session-start match-all
  class type control subscriber class-default do-all
    1 activate dynamic-template IPSUB_TEMPLATE
  !
end-policy-map
!
interface Bundle-Ether 1.100
  ipv4 address 20.1.1.1 255.255.255.0
  ipv6 address 1001::2/128
  ipv6 nd dad attempts 0
  encapsulation dot1q 100
  service-policy type control subscriber POLICY1
  ipsubscriber interface
!
track abc
  type line-protocol state
  interface GigabitEthernet0/0/0/0

```

```

!
!
subscriber
  redundancy
    source-interface GigabitEthernet0/0/0/10
    group 1
      preferred-role master
      slave-mode hot
      peer 10.1.1.2
      access-tracking abc
      interface-list
        interface Bundle-Ether1.100 id 1
!

```

Verification

Router#**show subscriber session all**

Codes: IN - Initialize, CN - Connecting, CD - Connected, AC - Activated,
ID - Idle, DN - Disconnecting, ED - End

Type	Interface	State	Subscriber IP Addr / Prefix	LNS Address (Vrf)
------	-----------	-------	-----------------------------	-------------------

```
-----
```

IP:STATIC	BE1.100	AC	20.1.1.1 (default)	
-----------	---------	----	--------------------	--

RP/0/RSP0/CPU0:server-1#

RP/0/RSP0/CPU0:server-1#sh subscriber session all detail

Sat Oct 28 21:30:11.303 IST

Interface: Bundle-Ether1.100

Circuit ID: Unknown

Remote ID: Unknown

Type: IP: Static

IPv4 State: Up, Sat Oct 28 21:24:06 2017

IPv4 Address: 20.1.1.1, VRF: default

Mac Address: Unknown

Account-Session Id: 00000001

Nas-Port: Unknown

User name: unknown

Formatted User name: unknown

Client User name: unknown

Outer VLAN ID: 100

Subscriber Label: 0x00000040

Created: Sat Oct 28 21:24:06 2017

State: Activated

Authentication: unauthenticated

Authorization: unauthorized

Access-interface: Bundle-Ether1.100

Policy Executed:

policy-map type control subscriber POLICY1

 event Session-Start match-all [at Sat Oct 28 21:24:06 2017]

 class type control subscriber class-default do-all [Succeeded]

 1 activate dynamic-template IPSUB_TEMPLATE [Succeeded]

Session Accounting: disabled

Last COA request received: unavailable

Router#**show ipsubscriber interface**

Interface: Bundle-Ether1.100

Type: Static

Access Interface: Bundle-Ether1.100

Subscriber Label: 0x40

IPv4 Initiator: Packet-Trigger

VLAN ID: 100

Created: Oct 28 21:24:06 (age 00:00:31)

```

IPv4 State: Up (old: Session features applied)
Last state change: Oct 28 21:24:06 (00:00:31 in current state)

Router#show subscriber redundancy group
Subscriber Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred Slave
            H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S        : Peer Status
            I - Initialize, Y - Retry, X - Cleanup, T - Connecting
            L - Listening, R- Registered, C - Connected, E - Established
I/F Count: Interface Count
SS Count  : Subscriber Session Count

```

```

-----
Node Name | Group ID | Role | Flags | Peer Address | P/S | I/F Count
| SS Count | Sync Pending
-----
0/RSP0/CPU0      1 Master EMHT 10.1.1.2      E      1
1                0

```

```

Session Summary Count(Master/Slave/Total): 1/0/1

```

```

Router#show subscriber redundancy group 1 session
Subscriber Redundancy Agent Group Session

```

```

-----
Parent Interface | MAC Address | In/Out VLAN | PPPoE/L2TP ID
-----
Bundle-Ether1.100      0000.0000.0000  0/100      0/0
-----

```

SRG Support for LC Subscribers

BNG introduces the support for subscriber redundancy group (SRG) for LC subscriber sessions. This feature ensures that if there is any access link failure or a complete LC node failure, the end user services are not impacted.

Restrictions for SRG Support for LC Subscribers

SRG support for LC subscribers in BNG is subjected to these restrictions:

- SRG is not supported for ambiguous VLAN LC subscriber sessions.
- PQoS is not supported for LC subscriber sessions with SRG.
- LC subscriber sessions cannot be cleared on the subordinate SRG.
- Clearing DHCP or DHCPv6 bindings is not supported on the subordinate SRG.
- Changing the MAC address of the interface with active sessions is not supported.
- The command **clear subscriber redundancy session group** *group-id* is not supported.
- Changing the SRG vMAC which has active sessions associated with it, is not supported.
- Configuring or unconfiguring the command **peer route-disable** with active sessions is not supported.
- Configuring or unconfiguring the command **enable-fast-switchover** with active sessions is not supported.
- If the line card is removed and re-inserted, the session accounting records for the primary line card sessions is sent with 0 statistics.

- Enabling fast switchover is not supported for subscriber framed-routes.
- Restart of the process `srg_agt` or an RSP failover can trigger SRG switchover for some groups.
- Core link failure with active sessions on both SRG nodes is not supported.
- Traffic for the primary LC sessions can drop for 30 seconds when the LC is removed and re-inserted.
- If SRG is disabled for group, then you need to wait for all the routes and sessions of that group to be deleted before SRG is enabled for that group again.
- Avoid the restart of any of the IOS XR processes simultaneously on both primary and subordinate nodes.
- The configuration for `idle-timeout` is not supported for SRG sessions.
- Clear all the sessions on LC before replacing the card with the new one.
- When both the primary and subordinate SRG nodes are being configured simultaneously, all the groups taking preferred role is not guaranteed.
- Access switch reload which results in the access links going down on both SRG nodes is not supported.
- Removing the SRG configuration on the primary node with active sessions is not supported.
- The command **admin srg switchover** is meant to be used for planned upgrades when subscriber churn is very less.
- If there is an inconsistency on the primary and subordinate nodes after SRG switchover, execute the command **subscriber redundancy synchronize group group-id** command on new primary in order to clear the inconsistency. If the subordinate node has more sessions compared to primary, execute the same command on the subordinate node. If inconsistency is still not cleared, clients will reconnect after the lease timer expires.
- Session state mismatch cannot be handled by **subscriber redundancy synchronize** command. If there is mismatch, clear the sessions administratively on the primary node. Clients will reconnect due to the lease timer expiry.

Configure Subscriber Redundancy Group for LC Subscriber Sessions

Configuration Steps

The below section shows how to configure SRG for LC subscriber sessions.

```
Router# config terminal
Router(config)# interface GigabitEthernet0/0/0/0.100
Router(config-subif)# ipv4 point-to-point
Router(config-subif)# ipv4 unnumbered Loopback0
Router(config-subif)# ipv6 enable
Router(config-subif)# service-policy type control subscriber POLICY1
Router(config-subif)# encapsulation dot1q 100
Router(config-subif)# ipsubscriber ipv4 l2-connected
Router(config-if-ipsub-ipv4-l2conn)# initiator dhcp
Router(config-if-ipsub-ipv4-l2conn)# exit
Router(config-subif)# ipsubscriber ipv6 l2-connected
Router(config-if-ipsub-ipv6-l2conn)# initiator dhcp
Router(config-if-ipsub-ipv6-l2conn)# exit
Router(config-subif)# exit
Router(config)# track access-link1
Router(config-track)# type line-protocol state
```

```

Router(config-track-line-prot)# interface GigabitEthernet0/0/0/0
Router(config-track-line-prot)# exit
Router(config-track)# exit
Router(config)# subscriber
Router(config-subscriber)# redundancy
Router(config-subscr-red)# source-interface GigabitEthernet0/1/0/10
Router(config-subscr-red)# group 1
Router(config-subscr-red-group)# preferred-role master
Router(config-subscr-red-group)# virtual-mac 00a1.00b1.00c1
Router(config-subscr-red-group)# peer 55.1.0.1
Router(config-subscr-red-group)# access-tracking access-link1
Router(config-subscr-red-group)# state-control-route ipv4 192.0.2.0/8 tag 1
Router(config-subscr-red-group)# state-control-route ipv6 iana 2001:DB8::/32 tag 1
Router(config-subscr-red-group)# state-control-route ipv6 iapd 2001:DB8:1::1/32 tag 1
Router(config-subscr-red-group)# interface-list
Router(config-subscr-red-grp-intf)# interface GigabitEthernet0/0/0/0.100 id 1
Router(config-subscr-red-grp-intf)# commit

```

Running Configuration

```

interface GigabitEthernet0/0/0/0.100
  ipv4 point-to-point
  ipv4 unnumbered Loopback0
  ipv6 enable
  service-policy type control subscriber POLICY1
  encapsulation dot1q 100
  ipsubscriber ipv4 l2-connected
    initiator dhcp
  !
  ipsubscriber ipv6 l2-connected
    initiator dhcp
  !
!
track access-link1
  type line-protocol state
  interface GigabitEthernet0/0/0/0
  !
!
subscriber
  redundancy
    source-interface GigabitEthernet0/1/0/10
    group 1
      preferred-role master
      virtual-mac 00a1.00b1.00c1
      peer 55.1.0.1
      access-tracking access-link1
      state-control-route ipv4 192.0.2.0/8 tag 1
      state-control-route ipv6 iana 2001:DB8::/32 tag 1
      state-control-route ipv6 iapd 2001:DB8:1::1/32 tag 1
      interface-list
        interface GigabitEthernet0/0/0/0.100 id 1
    !
  !
!
!
!

```

Verification

```

Router# show ipsubscriber access-interface brief
Codes: UP - Up, DOWN - Down, DELETED - Deleted State, UNKNOWN - Unknown State,
       PKT - Packet Trigger Initiation, DHCP - DHCP Initiation,
       PKTv6 - Packet Trigger Initiation for IPv6, DHCPv6 - DHCPv6 Initiation

```

```

Interface          Proto          DHCP          Pkt Trigger DHCPv6          PktTrigIPv6
State
-----
Gi0/0/0/0.100     DHCP,DHCPv6          1          0          1
0          UP

```

```

Router# show subscriber redundancy group
Subscriber Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred Slave
           : H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S       : Peer Status
           : I - Initialize, Y - Retry, X - Cleanup, T - Connecting
           : L - Listening, R- Registered, C - Connected, E - Established
I/F Count: Interface Count
SS Count  : Subscriber Session Count

```

```

Node Name | Group ID | Role | Flags | Peer Address | P/S | I/F Count
| SS Count | Sync Pending
-----
0/0/CPU0 |          | 1 Master | EMHT | 55.1.0.1 | E | 1
          | 1          | 0

```

```


Session Summary Count(Master/Slave/Total): 1/0/1

```

Subscriber Redundancy Group for Pseudowire Headend Subscribers

Cisco IOS XR Software Release 6.6.3 introduces the support for subscriber redundancy group (SRG) for pseudowire headend (PWHE) subscribers on BNG routers. The geographical redundancy feature helps in failover or planned switchover of subscriber sessions from one BNG router to another. This feature in turn allows service providers to support geographical redundancy for PWHE subscriber sessions on BNG. Currently, the SRG feature for PWHE subscribers is supported for BNG IPoE sessions.

Figure 14: Sample Network Topology of SRG for PWHE Subscriber Sessions on BNG




The geo redundancy solution exists at the BNG nodes, BNG1 and BNG2, in this topology. These BNG nodes terminate IPoE subscriber sessions in primary or subordinate mode. The access aggregation nodes (AA 1 and AA 2, in this topology) are connected to BNG nodes through PWHE. The PWHE provides Layer 2 loop avoidance technologies in case the link between the aggregation nodes is to be utilized for failure scenarios. The PWHE terminates access PW into a Layer 2 domain from aggregation node to the BNG. Each PW-Ethernet sub-interface on the BNG represents a ring of access aggregation. The Virtual Private Wire Service (VPWS)-Ethernet transport solution (virtual connection type 5 cross-connect) exists between the BNG and the aggregation node. The SRG is configured to track PW-Ether interface on primary and subordinate modes. When PW-Ether interface goes down on primary, the SRG triggers the switchover to the subordinate.

Link Failure Scenarios in BNG PWHE with SRG

This section illustrates some of the access link failure scenarios in BNG PWHE topology with SRG. Scenario 1 depicts a link failure between aggregation node 1 and BNG1. Scenario 2 depicts a link failure between the access aggregator 1 and aggregation node 1. In these scenarios, the SRG 1 is in primary mode on BNG1 and at the same time in subordinate mode on BNG2. In both these cases, the PWHE interface goes down on BNG1. The SRG 1 then does a switchover to BNG2.

The handling of core link failure scenarios with PWHE sessions remain the same as in regular SRG scenarios.

Figure 15: Link Failure Scenarios in BNG PWHE topology with SRG



368743

Restrictions for SRG for PWHE subscribers on BNG

The SRG support for PWHE subscribers on BNG is subjected to these restrictions:

- Supported only on the SE variants of Cisco ASR 9000 Enhanced Ethernet Line Cards and Cisco ASR 9000 High Density 100GE Ethernet Line Cards, with RSP880, RSP880-LT, and RSP5.
- Supported only for DHCP-initiated IPoE (IPv4, IPv6 or dual-stack) L2-connected sessions with BNG as DHCPv4 or DHCPv6 proxy.
- Once SRG switchover is triggered, the next SRG switchover must not be triggered within the recommended value of revertive timer (approximately 10 to 12 minutes). If back-to-back switchover happens, the primary and subordinate sessions might go out of sync.
- Not supported for L2VPN hot-standby pseudowire (HSPW) deployment model solution.
- Scale limit is subjected to the number of interfaces in the generic interface list of PWHE.

Configure SRG for Pseudowire Headend Subscribers on BNG

Configuration Example

You must perform these tasks in order to configure SRG for pseudowire headend subscribers on BNG:

- Configure L2VPN

- Configure MPLS
- Configure OSPF
- Configure PWHE
- Configure subscriber redundancy

The below example shows the configuration steps done on the primary BNG router. Repeat these steps on the subordinate BNG router as well with the respective values.

```

/* Configure L2VPN */
Router#configure
Router(config)#l2vpn
Router(config-l2vpn)#pw-class atom
Router(config-l2vpn-pwc)#encapsulation mpls
Router(config-l2vpn-pwc-mpls)#exit
Router(config-l2vpn-pwc)#exit
Router(config-l2vpn)#xconnect group pwhe
Router(config-l2vpn-xc)#p2p bng-pwhe
Router(config-l2vpn-xc-p2p)#interface PW-Ether 1
Router(config-l2vpn-xc-p2p)#neighbor ipv4 192.0.2.1 pw-id 101
Router(config-l2vpn-xc-p2p-pw)#pw-class atom
Router(config-l2vpn-xc-p2p-pw)#commit

/* Configure MPLS */
Router(config)#mpls ldp
Router(config-ldp)#log
Router(config-ldp-log)#neighbor
Router(config-ldp-log)#nsr
Router(config-ldp-log)#graceful-restart
Router(config-ldp-log)#exit
Router(config-ldp)#router-id 203.0.113.1
Router(config-ldp)#interface TenGigE 0/1/0/10
Router#commit

/* Configure OSPF */
Router(config)#router ospf srg_pw_test
Router(config-ospf)#router-id 203.0.113.1
Router(config-ospf)#address-family ipv4 unicast
Router(config-ospf)#area 0
Router(config-ospf-ar)#interface Loopback 101
Router(config-ospf-ar-if)#exit
Router(config-ospf-ar)#interface TenGigE 0/1/0/10
Router(config-ospf-ar-if)#commit

/* Configure PWHE */
Router(config)#generic-interface-list pwhe1
Router(config-gen-if-list)#interface TenGigE 0/1/0/10
Router(config-gen-if-list)#commit

Router(config)#interface PW-Ether1
Router(config-if)#attach generic-interface-list pwhe1
Router(config-if)#exit
Router(config)#interface PW-Ether1.1
Router(config-subif)#ipv4 unnumbered Loopback1
Router(config-subif)#ipv6 address 2001:DB8::1/64
Router(config-subif)#ipv6 enable
Router(config-subif)#service-policy type control subscriber TEST_POLICY
Router(config-subif)#encapsulation dot1q 1
Router(config-subif)#ipsubscriber ipv4 l2-connected

```

```

Router(config-if-ipsub-ipv4-l2conn)#initiator dhcp
Router(config-if-ipsub-ipv4-l2conn)#exit
Router(config-subif)#ipsubscriber ipv6 l2-connected
Router(config-if-ipsub-ipv6-l2conn)#initiator dhcp
Router(config-if-ipsub-ipv6-l2conn)#commit

Router(config)#track access1
Router(config-track)#type line-protocol state
Router(config-track-line-prot)#interface PW-Ether1.1
Router(config-track-line-prot)#commit

/* Configure Subscriber Redundancy */
Router(config)#subscriber
Router(config-subscriber)#redundancy
Router(config-subscr-red)#source-interface Loopback500
Router(config-subscr-red)#hold-timer 5
Router(config-subscr-red)#group 1
Router(config-subscr-red-group)#virtual-mac 0000.0000.0001
Router(config-subscr-red-group)#peer 203.0.113.2
Router(config-subscr-red-group)#core-tracking core
Router(config-subscr-red-group)#access-tracking access1
Router(config-subscr-red-group)#enable-fast-switchover
Router(config-subscr-red-group)#state-control-route ipv4 198.51.100.0/21 vrf default tag
10
Router(config-subscr-red-group)#state-control-route ipv6 iana 2001:DB8:A:B::1/64 vrf default
tag 10
Router(config-subscr-red-group)#state-control-route ipv6 iapd 2001:DB8:A:B:ABCD::1/64 vrf
default tag 10
Router(config-subscr-red-group)#interface-list
Router(config-subscr-red-grp-intf)#interface PW-Ether1.1 id 1
Router(config-subscr-red-grp-intf)#commit

Router(config)#ipsubscriber ipv4 l2-connected
Router(config-gen-if-list)#interface TenGigE 0/1/0/10
Router(config-gen-if-list)#commit

Router(config)#interface PW-Ether1
Router(config-l2vpn-xc-p2p-pw)#attach generic-interface-list pwhe1

```



Note For efficient tracking, it is recommended to track the member links under the generic interface list (GIL) rather than the parent PW interface itself, as it gives faster convergence.

Running Configuration

```

l2vpn
pw-class atom
  encapsulation mpls
  protocol ldp
  transport-mode ethernet
  !
!
xconnect group pwhe
  p2p bng-pwhe
  interface PW-Ether1
  neighbor ipv4 192.0.2.1 pw-id 101
  pw-class atom
  !
!

```

```

!
mpls ldp
  log
  neighbor
  nsr
  graceful-restart
!
router-id 203.0.113.1
interface Bundle-Ether1
!
interface TenGigE0/1/0/10
!
!
router ospf srg_pw_test
  router-id 203.0.113.1
  address-family ipv4 unicast
  area 0
    interface Bundle-Ether1
    !
    interface Loopback101
    !
    interface TenGigE0/1/0/10
    !
  !
!
generic-interface-list pwhe1
  interface Bundle-Ether1
  interface TenGigE0/1/0/10
!
interface PW-Ether1
  mtu 1514
  attach generic-interface-list pwhe1
!
interface PW-Ether1.1
  service-policy output spd subscriber-parent resource-id 0
  ipv4 unnumbered Loopback1
  ipv6 address 2001:DB8::1/64
  ipv6 enable
  service-policy type control subscriber TEST_POLICY
  encapsulation dot1q 1
  ipsubscriber ipv4 l2-connected
    initiator dhcp
  !
  ipsubscriber ipv6 l2-connected
    initiator dhcp
  !
!
track access1
  type line-protocol state
  interface PW-Ether1.1
!
!
subscriber
  redundancy
    source-interface Loopback500
    hold-timer 5
    group 1
      preferred-role master
      virtual-mac 0000.0000.0001
      peer 203.0.113.2
      core-tracking core
      access-tracking access1
      enable-fast-switchover
      state-control-route ipv4 198.51.100.0/21 vrf default tag 10

```

```

state-control-route ipv6 iana 2001:DB8:A:B::1/64 vrf default tag 10
state-control-route ipv6 iapd 2001:DB8:A:B:ABCD::1/64 vrf default tag 10
interface-list
  interface PW-Ether1.1 id 1
!
!
!

```

Verification

```

Router#show subscriber redundancy group
Thu Jan 24 11:57:03.766 UTC
Subscriber Redundancy Agent Group Summary
Flags      : E - Enabled, D - Disabled, M - Preferred Master, S - Preferred slave
            H - Hot Mode, W - Warm Mode, T - Object Tracking Enabled
P/S       : Peer Status
            I - Initialize, Y - Retry, X - Cleanup, T - Connecting
            L - Listening, R- Registered, C - Connected, E - Established
I/F Count: Interface Count
SS Count  : Subscriber Session Count

```

```

-----
Node Name | Group ID | Role | Flags | Peer Address | P/S | I/F Count | SS Count |
Sync Pending
-----
Router    | 1        | Master | EMHT | 203.0.113.2 | E    | 1         | 0
0

```

Related Topics

[Subscriber Redundancy Group for Pseudowire Headend Subscribers, on page 64](#)

Associated Commands

- [access-tracking](#)
- [state-control-route](#)
- [subscriber redundancy](#)
- [virtual-mac](#)

Deployment Models for BNG Geo Redundancy

Multiple access networks are considered for BNG geo redundancy deployment scenarios. Some of the sample use cases are:

- Multi-chassis Link Aggregation (MC-LAG) - Two BNG boxes that are point-of-attachment (POA) devices, connected through MC-LAG either to a single Dual Homed Device (DHD) or to a DHD-pair using MC-LAG.
- Multiple Spanning Tree - Access Gateway (MST-AG):
 - Dual Homed Device using Bundle Interfaces - A single DHD with one bundle interface each to the two BNGs in active-active mode.

- Ethernet Access Network-Ring - A physical ring (open or closed) that connects multiple OLTs (or L2 devices in general) to the two BNGs in active-active mode.

