



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.

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 11
- [Geo Redundancy for PPPoE Sessions](#), on page 28
- [BNG Geo Redundancy with Satellite](#), on page 31
- [Geo Redundancy Features](#), on page 34
- [Deployment Models for BNG Geo Redundancy](#), on page 41

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:

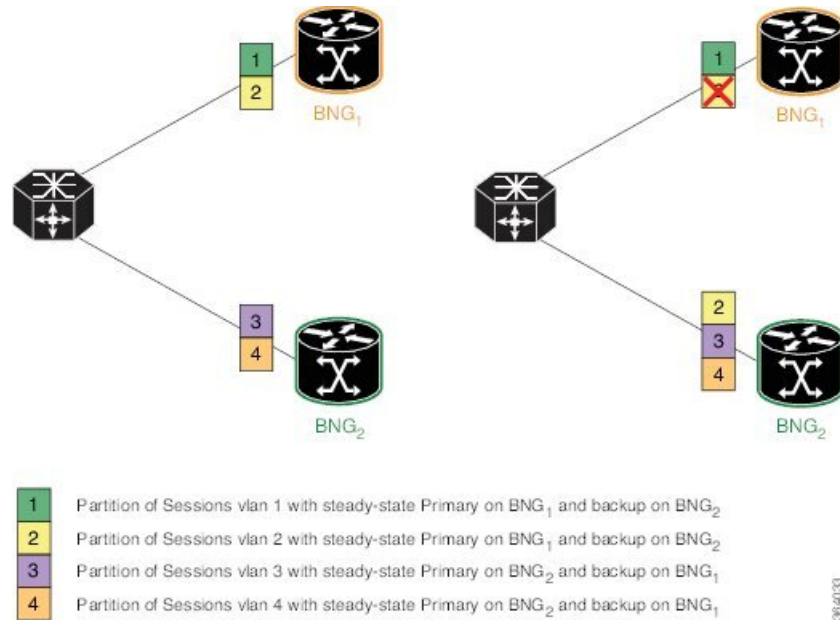
- An interface failure gets detected (using object-tracking of the access-interface) through MC-LAG.
- MC-LAG and SRG for each VLAN on BNG2 gets the primary role.
- All 32K sessions are switched to BNG2.
- BNG2 sees a session termination count of 32K.

- 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.
- 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).

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.
- Session initiation through DHCPv4 or DHCPv6 protocol.

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
- SRG for ambiguous VLAN BNG sessions.
- SRG between Cisco IOS XR 64-bit BNG node and 32-bit BNG node.
- Both RP and LC subscribers do not support enabling fast switchover for subscriber framed-routes.
- Excessive Punt Flow Trap (EPFT)
- IPoE and PPPoE sessions do not support `idle-timeout` configuration for both RP and LC-based subscribers.

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

- Warm-standby subordinate mode.
- Line card (LC) based subscribers (that is, using physical port sub-interfaces).
- DHCP server mode.
- Pseudowire Headend (PWHE)

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 access-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 are 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.

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
```

- Setup access object tracking for SRG and summary subscriber route:

```
track mclag-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.


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).



Note When a SERG peering is down, the communication between the SERG agents is lost. So the role of the subordinate routers automatically switches to the primary role. The routers on both sides act as a primary router, but only the SERG group configured with the preferred role as primary can provide addresses to subscribers establishing sessions.

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*.

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.

2. 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
  !
!
```

- 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
Pending Update  : 0
Client:
  IPv6ND        : 0
  DHCPv6        : 0
Slave-Upd-Fail : 0
Pending Delete  : 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	00030001ca011bba000000000000000000000000	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	0
Invalid DeRegistration	0	0
Connection Up Count	2	0
Connection Down Count	0	0
Message CallBack Count	2	0
Message Received	4	0
Command Message Received	0	0
Session Message Received	4	0
Peer Done	2	0

Peer Statistics:

	Ok	NotOk
Timer Handler	0	0
Invalid Registration	0	0
Invalid DeRegistration	0	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
  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          3
  Delete           :          0          2
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
  Update           :          0          0
  Delete           :          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 0
Connection Down : 0 0
Cleanup : 0 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
  !
!
```

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

Peer Statistics:

	Ok	NotOk
Timer Handler	0	0
Invalid Registration	0	0
Invalid DeRegistration	0	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 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:		

```
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
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         :           7
Clear             - All           :           0
Clear             - Selected      :           0
Replay           - All           :           1
Replay           - Selected      :           0
Session
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 20.

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.

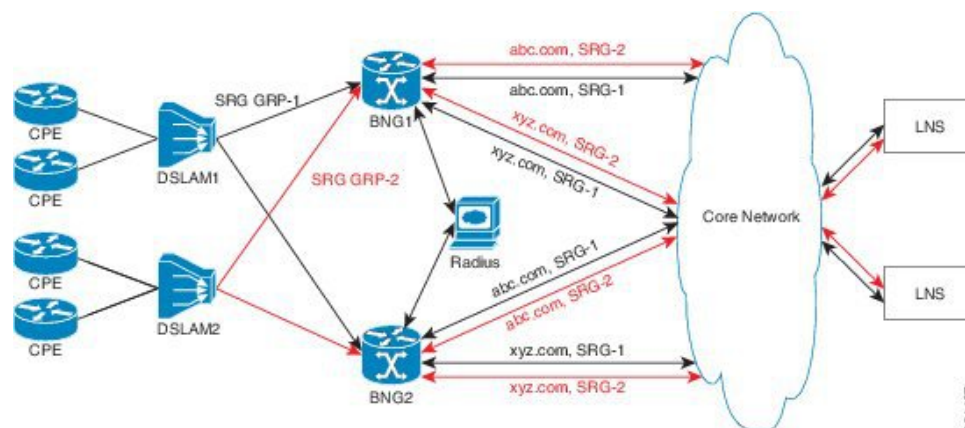
PPPoE-PTA Geo Redundancy

Geo redundancy behavior for the PPPoE-PTA sessions remains the same as for basic geo redundancy set up, except that the keeplives are disabled on the subordinate BNG node. The keeplives 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 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

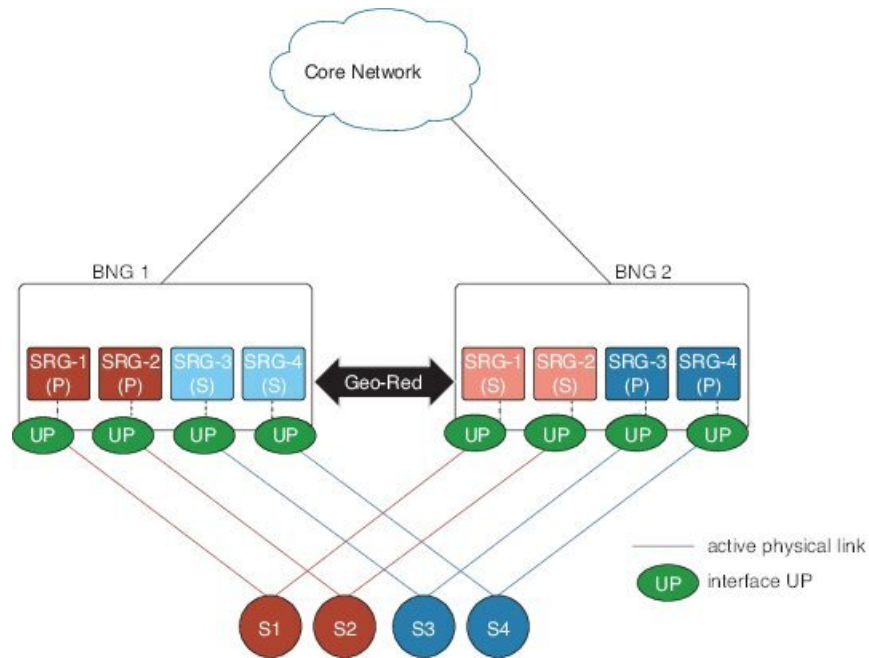
```
RP/0/RSP0/CPU0:router(config)# subscriber redundancy group 110
RP/0/RSP0/CPU0:router(config-subscr-red-group)# peer route-disable
```

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 10: 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.

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 11: 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 12: 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 generates 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 28](#)
- [DHCPv6 Address or Prefix Pool](#)
- [IP Address Assignment for Clients](#)
- [Scenario 1: SLAAC-Based Address Assignment](#)

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.

