



Geo Redundancy (Subscriber Redundancy Group)

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

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Subscriber Redundancy Group on Cloud Native BNG	Release 25.1.1	We now extend Subscriber Redundancy Group (SRG) support to PPPoE subscriber sessions. Subscriber Redundancy Group (SRG) provides flexible redundancy pairing on an access link by mirroring the subscriber session to a standby node.

Feature Name	Release Information	Feature Description
Subscriber Redundancy Group on Cloud Native BNG	Release 7.8.1	<p>You can now enable redundancy for subscriber sessions across two or more cnBNG user planes spread across different geographical locations by configuring redundancy for that subscriber group.</p> <p>Subscriber Redundancy Group (SRG) provides flexible redundancy pairing on an access link by mirroring the subscriber session to a standby node.</p> <p>When SRG is enabled, subscriber sessions are unaffected during the failure of the access link, and maintenance downtimes as the switchover happen from an active to a standby user plane automatically, or the BNG control plane assigns the active role to the user plane.</p> <p>This feature introduces these changes:</p> <p>subscriber-redundancy</p>

- [Overview, on page 2](#)
- [SRG Modes, on page 3](#)
- [Subscriber Session Set up Call Flow, on page 7](#)
- [Benefits of BNG Geo Redundancy, on page 9](#)
- [Supported Features in BNG Geo Redundancy, on page 10](#)
- [Unsupported Features and Restrictions for cnBNG Geo Redundancy, on page 10](#)
- [Guidelines to Configure SRG, on page 11](#)
- [Configure SRG, on page 11](#)
- [Routed subscriber sessions, on page 19](#)
- [SRG warm-standby mode , on page 26](#)

Overview

Using Subscriber Redundancy Group (SRG), you can now provide redundancy for the subscriber sessions across multiple BNGs located in multiple geographical locations with L3 connectivity over a shared core network through IP or MPLS routing.

SRG provides flexible redundancy pairing on access-link and performs automatic switchovers during dynamic failures or planned events such as maintenance, upgrades, and transitions.

SRG also termed Geo redundancy is a powerful technology that allows session synchronization between two nodes. An active session on one node is mirrored on a standby node, so that when the active link fails, the standby BNG can take over and continue to forward the subscriber session information without any service interruption to the user.

When the subscriber session is up on cnBNG, the control plane BNG synchronizes the state from the active to the backup User Plane (UP) cnBNG. The sessions are mirrored on the standby UP for redundancy by transferring the relevant session state from active UP to standby UP, which can then help in failover (FO) or planned switchover (SO) of sessions from one UP to another. SRG, which is a set of access-interface (or a single access-interface) is introduced in cnBNG, and all subscribers in an SRG would FO or SO as a group.

For more information about the cnBNG control plane, refer to the *Cloud Native BNG Control Plane Configuration Guide*.

CPEs are agnostic to redundancy. When you enable SRG, CPE peers with the same MAC address and node ID to fall back when there is a failover.

Control plane cnBNG initiates the SRG switchover to the standby node during:

- Access link failure
- Core network link failure
- RP failures
- Chassis failure

SRG Modes

The SRG has two modes of operation:

- Hot-standby
- Warm-standby

However, we support only Hot-Standy mode.

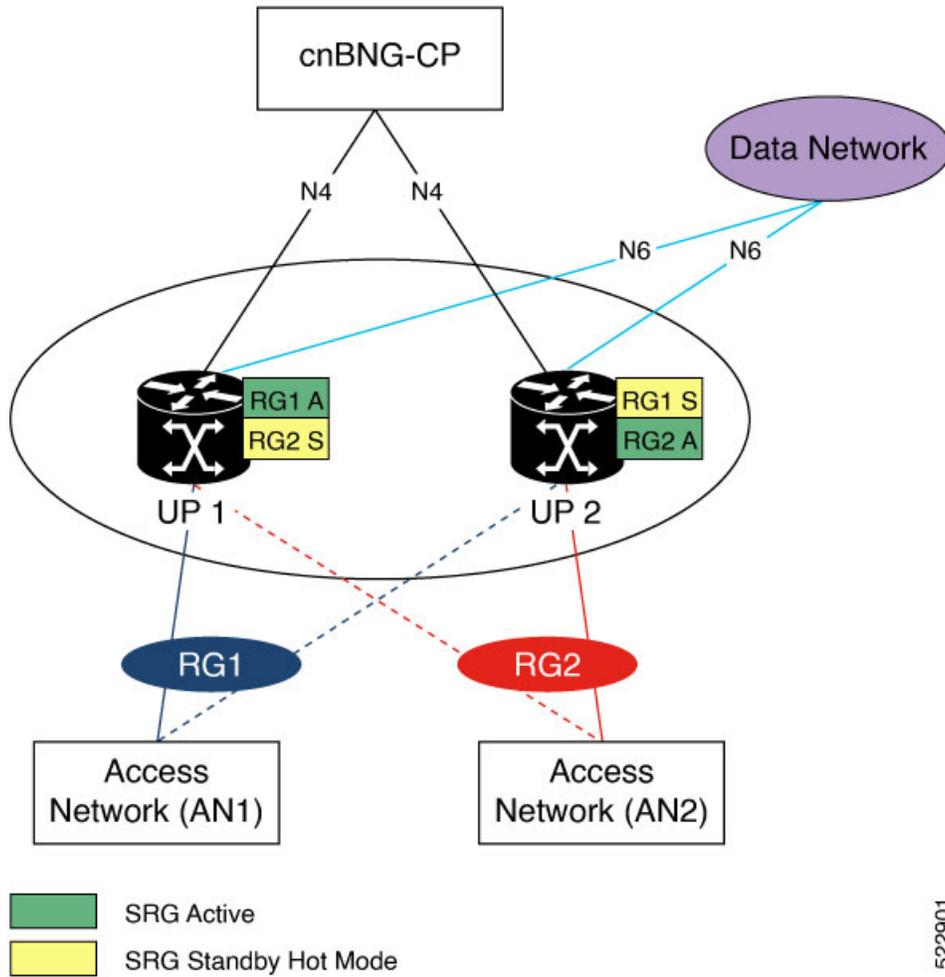
The Hot-standby mode supports 1:1 and M:N submodes.

1:1 Hot-standby

In this mode, 50 percent of the groups are in the active state and 50 percent of the groups are in the standby state.

In this topology, access network AN1 is dual homed to UP1 and UP2. All subscribers from AN1 are grouped under the RG1 group. Access network AN2 is dual homed to UP1 and UP2 and all subscribers from AN2 are grouped under the RG2 group.

cnBNG CP elects RG1 group as active in UP1 and standby in UP2. SRG is configured such that each UP is active for 50 percent of the groups and back up 50 percent of the groups in this mode.



1:1 Hot-standby mode supports the following submodes:

Mode	Description	Illustration
Active-active	In this mode, one cnBNG UP is active for some SRGs and its peer cnBNG UP is active for other SRGs.	<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>(A)</p> </div> <div style="text-align: center;"> <p>(B)</p> </div> </div> <div style="margin-top: 10px;"> <ul style="list-style-type: none"> 1 Partition of Sessions vlan 1 with steady-state Primary on UP₁ and standby on UP₂ 2 Partition of Sessions vlan 2 with steady-state Primary on UP₁ and standby on UP₂ 3 Partition of Sessions vlan 3 with steady-state Primary on UP₂ and standby on UP₁ 4 Partition of Sessions vlan 4 with steady-state Primary on UP₂ and standby on UP₁ </div> <p style="text-align: right; font-size: small;">522860</p>

Mode	Description	Illustration
Active-standby	In this mode, a cnBNG UP can be a dedicated standby for multiple SRGs from different cnBNG UPs that are active for those respective SRGs.	<p>1 Partition of Sessions vlan 1 with steady-state Primary on UP₁ and standby on UP₂ 2 Partition of Sessions vlan 2 with steady-state Primary on UP₁ and standby on UP₂ 3 Partition of Sessions vlan 3 with steady-state Primary on UP₁ and standby on UP₂ 4 Partition of Sessions vlan 4 with steady-state Primary on UP₁ and standby on UP₂</p>

M:N Hot-standby

In this mode, two nodes are active (M) and one node is standby (N) in the ratio M: N.

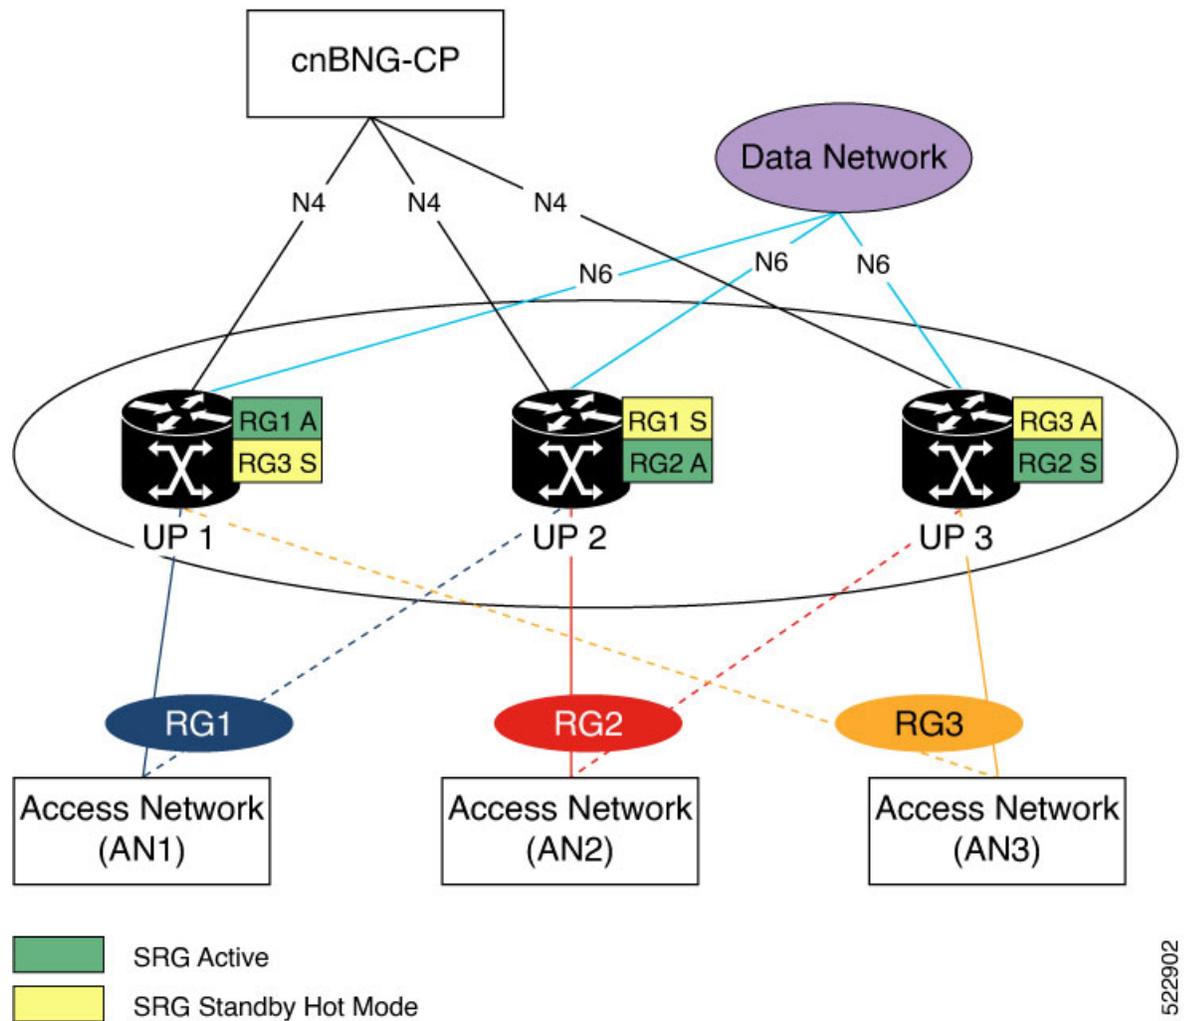
In this topology, access network AN1 is dual homed to UP1 and UP2. All subscribers from AN1 are grouped under the RG1 group. Access network AN2 is dual homed to UP1 and UP2 and all subscribers from AN2 are grouped under the RG2 group. Access network AN3 is dual homed to UP1 and UP3 and all subscribers from AN3 are grouped under RG3 group

cnBNG CP elects RG1 group as active in UP1 and standby in UP2. RG2 group is elected as active in UP2 and standby in UP3. Similarly, cnBNG elects RG3 as active in UP3 and standby in UP1.

In this example, we've three active nodes (M) and three standby nodes (N).

- M denotes an active node. As there are two active nodes, it takes the value 3.
- N denotes a standby node. As there is only one standby node, it takes the value 3.

So, M: N depicts the ratio 3: 3 (active: standby ratio).



522902

Subscriber Session Set up Call Flow

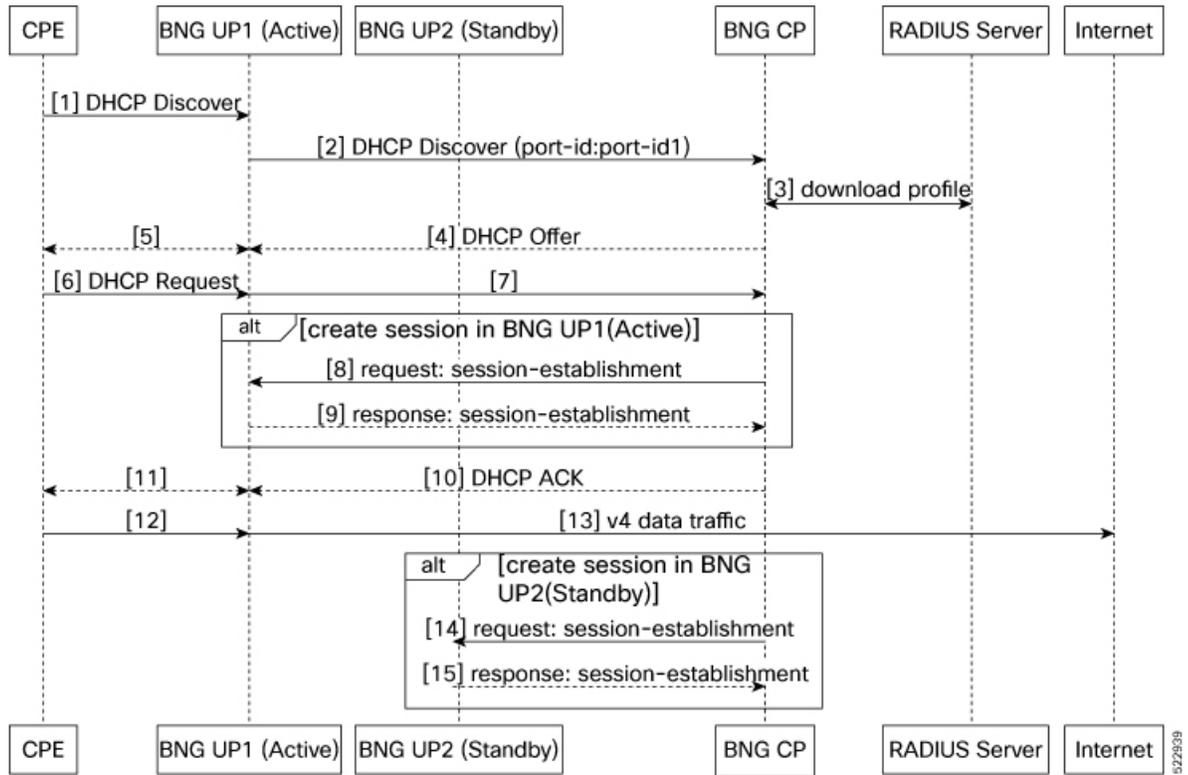
The following section graphs out the call flow and messaging between cnBNG SRG devices and the session.

Subscriber Session Creation Call flow for SRG

The following call flow illustrates the SRG subscriber session, where UP1 is the active node and UP2 is the standby node.

1. cnBNG CP triggers the session creation on UP1 after it receives the DHCP request from the CPE.
2. After the subscriber session is established, UP1 sends a response back to cnBNG CP with the subscriber session details.
3. cnBNG CP now sends the request to UP2 to mirror the subscriber session.

Figure 1: Subscriber Session Creation for SRG

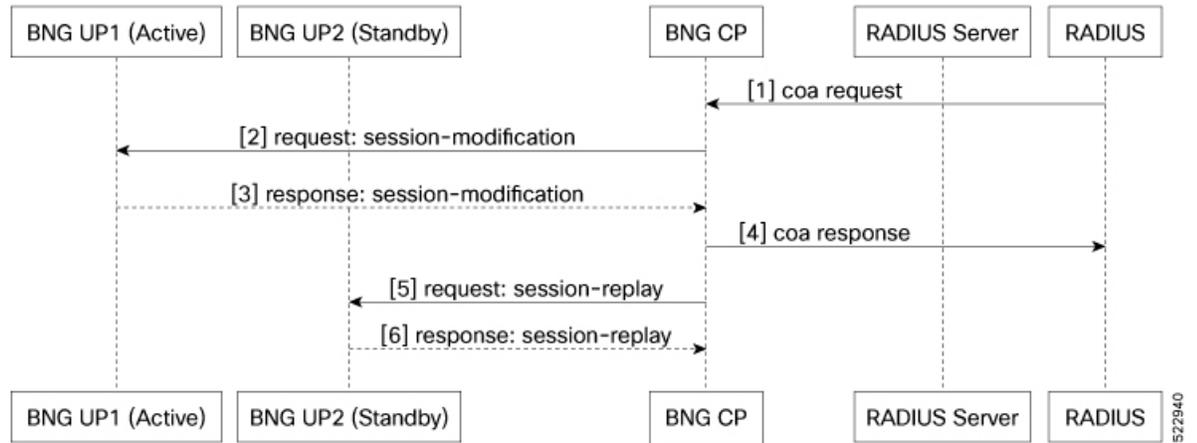


Subscriber Session attributes Modification Call flow for SRG

RADIUS Co-A (Change of Authorization) allows a RADIUS server to adjust an active client session. The following is the flow for modifying the subscriber session attributes for SRG:

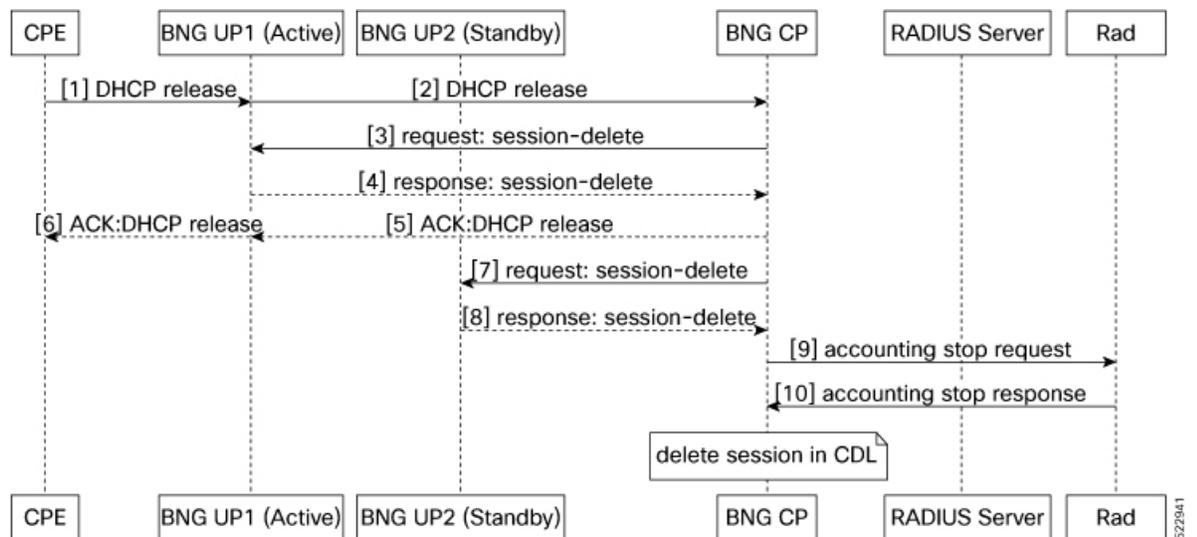
1. When there's a COA request from RADIUS, cnBNG CP triggers UP1 to modify the session attributes.
2. After the subscribers session attributes are modified, UP1 sends a response back to cnBNG CP with the subscriber session attribute details.
3. cnBNG CP now sends the COA response back to Radius and also triggers the UP2 that has a standby role.
4. UP2 modifies the session attributes and sends back the response to cnBNG CP.

Figure 2: Subscriber Session attributes Modification for SRG



Subscriber Session Deletion Call flow for SRG

Figure 3: Subscriber Session Deletion Call flow for SRG



Benefits of BNG Geo Redundancy

- Provides flexible redundancy pairing on access-link
- Supports multiple access networks such as dual-home and OLT rings
- Supports various types of subscribers such as IPv4, IPv6, and dual-stack IPoE sessions
- Supports 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
- 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

Supported Features in BNG Geo Redundancy

These access topologies are supported:

- SRG active–active mode without any access protocol
- Dual-home bundle interfaces with SRG vMAC using CFM or EFD fault detection
- Ring bundle interfaces with SRG vMAC using CFM or EFD fault detection
- MC-LAG access topology
- nV Satellite access topology

These base geo redundancy features are supported:

- RP-based subscribers with Bundle-Ether and PWHE as access interfaces
- Multiple SRG groups to different peer routers.
- Dynamic failure detection using object tracking (link up-down).
- Dampening timer supported
- Full BNG scale support (that is, half the scale number with redundancy).
- G.8032 (dual-home and ring) access technologies.
- SRG for ambiguous VLAN BNG session is supported only for IPoE subscriber sessions over bundle interface.
- PPPoE PTA is supported with SRG for RP (Bundle Ether or PWHE) based subscriber sessions.

Unsupported Features and Restrictions for cnBNG Geo Redundancy

cnBNG Geo Redundancy does not support the following:

- IPoE packet-triggered sessions
- Multicast
- LC subscriber sessions with or without SRG is not supported with PPPoE-PTA, LAC, IPoE Dual stack sessions

- ARP scale mode

Guidelines to Configure SRG

- At least one VLAN group must be configured to create SRG.
- 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.
- You must use the 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 one or more subscriber summary routes corresponding to that SRG's subscriber address or subnet pools.
- 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.
- Synchronization of the framed IPv6 prefix addresses in the SRG scenario is not supported on satellite bundle access interfaces in dual-homed satellite topology.
- 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 or subnet pool 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 reroute the traffic from the subordinate to the primary (for example, a tunnel or interchassis link) for transient or prolonged intervals.
- Routing convergence and its correlation with access failures or convergence is a key to the overall end-to-end service impact for subscribers. Multiple options exist to achieve subsecond intervals.

Configure SRG

Perform the following task to configure SRG:

```
/* Configure SRG and associate it with the access interface */
Router#configure
Router(config)#cnbng-nal location 0/0/CPU0
```

```

Router(config-cnbng-nal-local)#subscriber-redundancy
Router(config-cnbng-nal-sub-red)#group group1
Router(config-cnbng-nal-srg-grp)#virtual-mac 0aaa.0bbb.0c01
Router(config-cnbng-nal-srg-grp)# core-tracking core1
Router(config-cnbng-nal-srg-grp)#access-tracking track1
Router(config-cnbng-nal-srg-grp)#access-interface-list
Router(config-cfg-srg-grp-intf)#interface Bundle-Ether1.1
Router(config-cfg-srg-grp-intf)# exit
Router(config-cfg-srg-grp)# fast-switchover-disable
Router(config-cfg-srg-grp)# exit

Router(config-cnbng-nal-sub-red)#group group2
Router(config-cnbng-nal-srg-grp)#virtual-mac 0aaa.0bbb.0a02
Router(config-cnbng-nal-srg-grp)#core-tracking core1
Router(config-cnbng-nal-srg-grp)#access-tracking track1
Router(config-cnbng-nal-srg-grp)#access-interface-list
Router(config-cfg-srg-grp-intf)#interface Bundle-Ether1.2
Router(config-cfg-srg-grp-intf)# exit
Router(config-cfg-srg-grp)# fast-switchover-disable
Router(config-cfg-srg-grp)# exit

Router(config-cnbng-nal-srg-grp-red)#group group3
Router(config-cnbng-nal-srg-grp)#virtual-mac 0aaa.0bba.0a03
Router(config-cnbng-nal-srg-grp)#core-tracking core1
Router(config-cnbng-nal-srg-grp)#access-tracking track1
Router(config-cnbng-nal-srg-grp)#access-interface-list
Router(config-cfg-srg-grp-intf)#interface Bundle-Ether1.3
Router(config-cfg-srg-grp-intf)#exit
Router(config-cfg-srg-grp)#fast-switchover-disable
Router(config-cfg-srg-grp)#exit

Router#show running-config cnbng-nal location 0/0/CPU0
cnbng-nal location 0/0/CPU0
hostidentifier RTR1
up-server ipv4 10.11.11.1 gtp-port 15002 pfc-p-port 15003 vrf default
cp-server primary ipv4 10.11.11.2
auto-loopback vrf test
    interface Loopback1
    !
!
auto-loopback vrf default
    interface Loopback0
    !
!
disconnect-history file-logging-enable
spa-req-resp-history file-logging-enable
disable-secondary-address-notification
cp-association retry-count 5
ipoe fsol-flow-control 60
pppoe fsol-flow-control 60
subscriber-redundancy
    group group1
        virtual-mac 0aaa.0bbb.0c01
        core-tracking core1
        access-tracking track1
        access-interface-list
            interface Bundle-Ether1.1
            !
        fast-switchover-disable
    !
    group group2
        virtual-mac 0aaa.0bbb.0a02
        core-tracking core1

```

```

access-tracking track1
access-interface-list
  interface Bundle-Ether1.2
!
fast-switchover-disable
!

```

Verification

The following show output shows the list of SRG groups that you created and its role:

```
Router#show cnbng-nal srg-group
```

```

=====
Location: 0/0/CPU0
=====

```

Group-name	SRG role	Access OT	Core OT	Subs Count	V4 routes	V6 routes
group1	Active	Up	Up	1	2	2
group2	Active	Up	Up	1	2	2
group3	Active	Up	Up	1	2	2

```
Total Entries : 3
```

```
Summary
-----
```

Category	Total	Active	Standby	None
Groups	3	3	0	0
Subscribers	4	4	0	0
V4 subnet routes	16	16	0	0
V6 subnet routes	16	16	0	0

The following show output displays the detailed information about SRG that includes group name, role, ID, subscriber count, and so on.

```
Router#show cnbng-nal srg-group detail
```

```

=====
Location: 0/0/CPU0
=====

```

```

SRG group name           : group1
SRG group admin state    : UP_CP_Configured
SRG group state          : Up
SRG role                  : Active
SRG ID                   : 0x00000001
SRG VRF name             : -NA- (fast-switchover disabled)
Last SRG role update time : Oct 18 14:38:56.290388
Virtual mac              : 0AAA.0BBB.0C01
V4 Table Id              : 0x00000000
V6 Table Id              : 0x00000000
V4 Proto Id              : 0x0000ffff
V6 Proto Id              : 0x0000ffff
Subscriber count          : 1
IPV4 route count         : 2
IPV6 route count         : 2
Damping timer interval   : 120 Sec
Subnet route tag         : 0
Route export on Standby enable : False

```

```

Fast switchover enable           : False
Ready for role change           : Yes [Success]
FSM State                       : UNKNOWN
Update Request State            : IDLE
Sub disconnect resp pend        : NA

Access tracking object
-----
Object name                     : track1
Tracking state                   : Up
Last tracking state update time : Oct 18 14:38:39.822489

Core tracking object
-----
Object name                     : core1
Tracking state                   : Up
Last tracking state update time : Oct 18 14:38:39.821638

Access Interfaces
-----
Bundle-Ether1.1

IM counters
-----
Total entries                   : 1
Pending                         : 0
On-hold                         : 0
Total errors                    : 0

RIB counters
-----
Total entries                   : 0
Pending                         : 0
Total errors                    : 0

STATS counters
-----
Total entries                   : 0
Pending                         : 0
Total errors                    : 0
Stats state                     : IDLE

Flags
----
Value: [0x00000000]
None

Checkpoint Flags
-----
Value: [0x00000000]
None

CP Recon data
-----
Duration                       : 0 secs
Replay reqs in progress        : 0
Replay subs in progress        : 0
CP Recon Flags                  : 0x0

Subscriber transaction Info
-----
Subscribers in transaction     : 0
Subscribers in AF down queue   : 0
Subscribers in disc queue      : 0

```

Group role switchover stats Info

```

-----
Last stats interaction time(A->S): : 0.0 secs
Last stats interaction time(S->A): : 0.0 secs
Max stats interaction time(A->S): : 0.0 secs (NA)
Max stats interaction time(S->A): : 0.0 secs (NA)

```

Event history

```

-----
| Event Name                               | Time Stamp                               | S, M
|-----|-----|-----|
| Group create                             | Oct 18 14:38:39.820086                  | 0, 0
| V4 backup vrf create                     | Oct 18 14:38:39.820245                  | 0, 0
| V6 backup vrf create                     | Oct 18 14:38:39.820271                  | 0, 0
| Role active                              | Oct 18 14:38:56.290385                  | 0, 0
| Role active start                        | Oct 18 14:38:56.290388                  | 0, 0
| Role active end                          | Oct 18 14:38:56.290446                  | 0, 0
| CP action add                            | Oct 18 14:38:56.290447                  | 0, 0
| Notify: State Up                         | Oct 18 14:38:56.341312                  | 0, 0
| State change ack'ed                      | Oct 18 14:38:56.341434                  | 0, 0

```

```

=====
SRG group name           : group2
SRG group admin state   : UP_CP_Configured
SRG group state         : Up
SRG role                 : Active
SRG ID                  : 0x00000002
SRG VRF name            : -NA- (fast-switchover disabled)
Last SRG role update time : Oct 18 14:38:57.804402
Virtual mac              : 0AAA.0BBB.0A02
V4 Table Id             : 0x00000000
V6 Table Id             : 0x00000000
V4 Proto Id             : 0x0000ffff
V6 Proto Id             : 0x0000ffff
Subscriber count         : 1
IPV4 route count        : 2
IPV6 route count        : 2
Damping timer interval  : 120 Sec
Subnet route tag        : 0
Route export on Standby enable : False
Fast switchover enable  : False
Ready for role change    : Yes [Success]
FSM State                : UNKNOWN
Update Request State     : IDLE
Sub disconnect resp pend : NA

```

Access tracking object

```

-----
Object name              : track1
Tracking state           : Up
Last tracking state update time : Oct 18 14:38:39.823154

```

Core tracking object

```

-----
Object name              : core1
Tracking state           : Up
Last tracking state update time : Oct 18 14:38:39.823144

```

Access Interfaces

```

-----
Bundle-Ether1.2

```

IM counters

```

-----
      Total entries           : 1
      Pending                 : 0
      On-hold                 : 0
      Total errors            : 0

RIB counters
-----
      Total entries           : 0
      Pending                 : 0
      Total errors            : 0

STATS counters
-----
      Total entries           : 0
      Pending                 : 0
      Total errors            : 0
      Stats state             : IDLE

Flags
-----
Value: [0x00000000]
None

Checkpoint Flags
-----
Value: [0x00000000]
None

CP Recon data
-----
      Duration                 : 0 secs
      Replay reqs in progress  : 0
      Replay subs in progress  : 0
      CP Recon Flags           : 0x0

Subscriber transaction Info
-----
      Subscribers in transaction : 0
      Subscribers in AF down queue : 0
      Subscribers in disc queue   : 0

Group role switchover stats Info
-----
      Last stats interaction time(A->S): : 0.0 secs
      Last stats interaction time(S->A): : 0.0 secs
      Max stats interaction time(A->S): : 0.0 secs (NA)
      Max stats interaction time(S->A): : 0.0 secs (NA)

Event history
-----
| Event Name                               | Time Stamp                | S, M
|-----|-----|-----|
| Group create                             | Oct 18 14:38:39.822756    | 0, 0
| V4 backup vrf create                     | Oct 18 14:38:39.822846    | 0, 0
| V6 backup vrf create                     | Oct 18 14:38:39.822937    | 0, 0
| Role active                             | Oct 18 14:38:57.804399    | 0, 0
| Role active start                         | Oct 18 14:38:57.804402    | 0, 0
| Role active end                           | Oct 18 14:38:57.804448    | 0, 0
| CP action add                             | Oct 18 14:38:57.804448    | 0, 0
| Notify: State Up                          | Oct 18 14:38:57.855062    | 0, 0
| State change ack'ed                      | Oct 18 14:38:57.855170    | 0, 0
=====

```

The following example shows a sample configuration for MC-LAG on both BNG devices.

```

/* This the configuration on BNG1:*\:
redundancy
  iccp
    group 100
      mlacp node 1
      mlacp system mac 1001.2000.3015
      mlacp system priority 1
      mlacp connect timeout 0
      member
        neighbor 188.1.1.1
      !
      backbone
        interface Bundle-Ether131
        !
        isolation recovery-delay 300
      !
    group 500
      member
        neighbor 188.1.1.1
      !
      backbone
        interface Bundle-Ether131
        !
        isolation recovery-delay 100
        nv satellite
        system-mac 1001.4000.3015
      !
    !
  interface Bundle-Ether10
    mtu 9000
    lacp switchover suppress-flaps 100
    mlacp iccp-group 100
    mlacp switchover recovery-delay 60
    mlacp port-priority 10
    mac-address 10f3.1153.5a01
    bundle wait-while 0
    bundle load-balancing hash dst-ip
    bundle minimum-active links 1
    load-interval 30
  !
  interface Bundle-Ether10.1
    service-policy output SVLAN subscriber-parent resource-id 0
    vrf pppoe_vrf
    ipv4 point-to-point
    ipv4 unnumbered Loopback517
    arp learning disable
    ipv4 unreachableables disable
    ipv6 nd cache-limit 5000
    ipv6 nd other-config-flag
    ipv6 nd managed-config-flag
    ipv6 address 2006::c353/128
    ipv6 address 2007::c353/128
    ipv6 enable
    ipv6 unreachableables disable
    encapsulation dot1q 10 second-dot1q 1
    cnbng-nal ipv6 nd
      ra-initial 0 4
      slaac
    !
  ipsubscriber
    ipv4 l2-connected

```

```

    initiator dhcp
    !
    ipv6 l2-connected
    initiator dhcp
    !
    !
    pppoe enable
    !

interface TenGigE0/2/0/0
nv
  satellite-fabric-link satellite 300
  redundancy
    iccp-group 500
    !
  remote-ports TenGigE 0/0/0-30
  !
  !
  !

/* This the configuration on BNG2:*
redundancy
iccp
  group 100
    mlacp node 2
    mlacp system mac 1001.2000.3015
    mlacp system priority 1
    mlacp connect timeout 0
  member
    neighbor 177.1.1.1
    !
  backbone
    interface Bundle-Ether431
    !
    isolation recovery-delay 300
    !
  group 500
  member
    neighbor 177.1.1.1
    !
  backbone
    interface Bundle-Ether431
    !
    isolation recovery-delay 100
  nv satellite
    system-mac 1001.4000.3015
    !
    !

interface Bundle-Ether10
  mtu 9000
  lacp switchover suppress-flaps 100
  mlacp iccp-group 100
  mlacp switchover recovery-delay 60
  mlacp port-priority 5
  mac-address 10f3.1153.5a01
  bundle wait-while 0
  bundle load-balancing hash dst-ip
  bundle minimum-active links 1
  load-interval 30
  !

interface Bundle-Ether10.1
  service-policy output SVLAN subscriber-parent resource-id 0
  vrf pppoe_vrf
  ipv4 point-to-point

```

```
ipv4 unnumbered Loopback517
arp learning disable
ipv4 unreachable disable
ipv6 nd cache-limit 5000
ipv6 nd other-config-flag
ipv6 nd managed-config-flag
ipv6 address 2007::c353/128
ipv6 enable
ipv6 unreachable disable
encapsulation dot1q 10 second-dot1q 1
cnbng-nal ipv6 nd
  ra-initial 0 4
  slaac
!
ipsubscriber
  ipv4 l2-connected
  initiator dhcp
!
  ipv6 l2-connected
  initiator dhcp
!
!
pppoe enable
!
interface TenGigE0/1/0/7
  nv
  satellite-fabric-link satellite 300
  redundancy
  iccp-group 500
  !
  remote-ports TenGigE 0/0/0-30
  !
!
!
```

Routed subscriber sessions

The routed subscriber session is a type of subscriber session that

- forwards subscriber traffic through a Layer 3 network
- connects IP subscribers to the cnBNG through a routed access network
- utilizes DHCP-initiated connections, and
- identifies the subscribers by IP addresses instead of MAC addresses, applying policies and services.

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
Routed subscriber sessions	Release 25.1.1	<p>You can now enhance scalability by managing subscriber sessions over a routed network, allowing IP subscribers to connect to the cnBNG through a routed access network.</p> <p>Routed subscriber sessions enable dynamic IP management and provisioning across the IP cloud. By identifying subscribers with IP addresses instead of MAC addresses, it offers greater flexibility and efficiency in managing network resources, eliminates the need for direct Layer 2 connections, and simplifies access design with Layer 3 access.</p> <p>This feature introduces these changes:</p> <ul style="list-style-type: none"> • The prefix-len and src-ip-dual-lookup keywords are introduced in the initiator dhcp command. • The Group for routed subscribers field is added to the show cnbng-nal srg-group command output. • The routed type is added to the show cnbng-nal subscriber command. • The routed subscriber session counters are added to the show cnbng-nal counters command output. • The next-hop IP field is added to the show cnbng-nal dynamic-routes command output. • The ipoe-routed keyword is added to the show cnbng-nal access-interface if-type command. • The ipoe-routed keyword is added to the clear cnbng-nal subscriber sub-type command. <p>Yang Data Models:</p> <ul style="list-style-type: none"> • Cisco-IOS-XR-subscriber-nal-cfg.yang • Cisco-IOS-XR-cnbng-nal-oper.yang <p>(see GitHub, YANG Data Models Navigator)</p>

Why routed subscriber sessions are necessary

The current IP-MPLS network is designed to handle both wireless and wireline traffic. However, subscriber growth demands a scalable Layer 3 solution to manage effectively beyond the existing Layer 2 setup

Routed subscriber sessions provide a framework for efficiently managing and provisioning subscriber sessions on a large scale. This approach supports deploying Layer 3 solutions at the aggregation layer, crucial for meeting growing demand and ensuring seamless service delivery across networks

With routed subscriber session support, cnBNG provisions subscribers anywhere in the IP cloud, overcoming the restriction of Layer 2 connectivity.

Benefits of routed subscriber sessions

Routed subscriber sessions offer several advantages that enhance network scalability and management:

- Enables scalable solutions by allowing subscriber sessions to be initiated and managed over a routed network, bypassing the constraints of Layer 2 connectivity
- Supports subscriber provisioning across the IP cloud, eliminating the need for direct Layer 2 connections to the cnBNG.
- Utilizes DHCP for IP address allocation, enabling dynamic management of IP addresses and subscriber sessions.

How routed subscriber sessions work

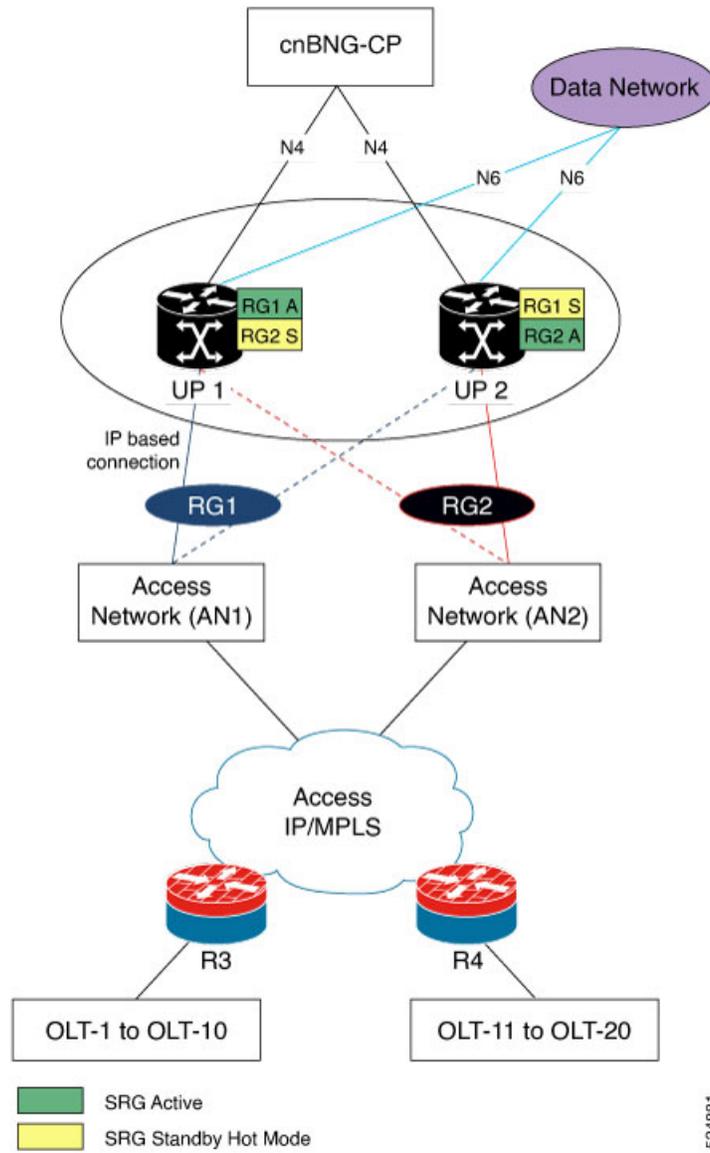
Routed subscriber sessions allow devices to connect to the cnBNG through routed devices. Instead of using MAC addresses, this session identification relies on IP addresses. The Customer Premises Equipment (CPE) uses Network Address Translation (NAT) to hide connected devices. The BNG sees a subscriber session linked to the CPE's WAN interface.

Summary

The key components involved in this process are:

- User Plane (UP): Receives discover or solicit packets and forwards them to the Control Plane (CP).
- Control Plane (CP): Manages subscriber data and sends DHCP offers to establish session.
- RADIUS Server: Provides authorization and DHCP class information.
- Connectivity: Establishes a layer 3 IP network connection between the access network (AN1/AN2) and UP for routed subscriber sessions.

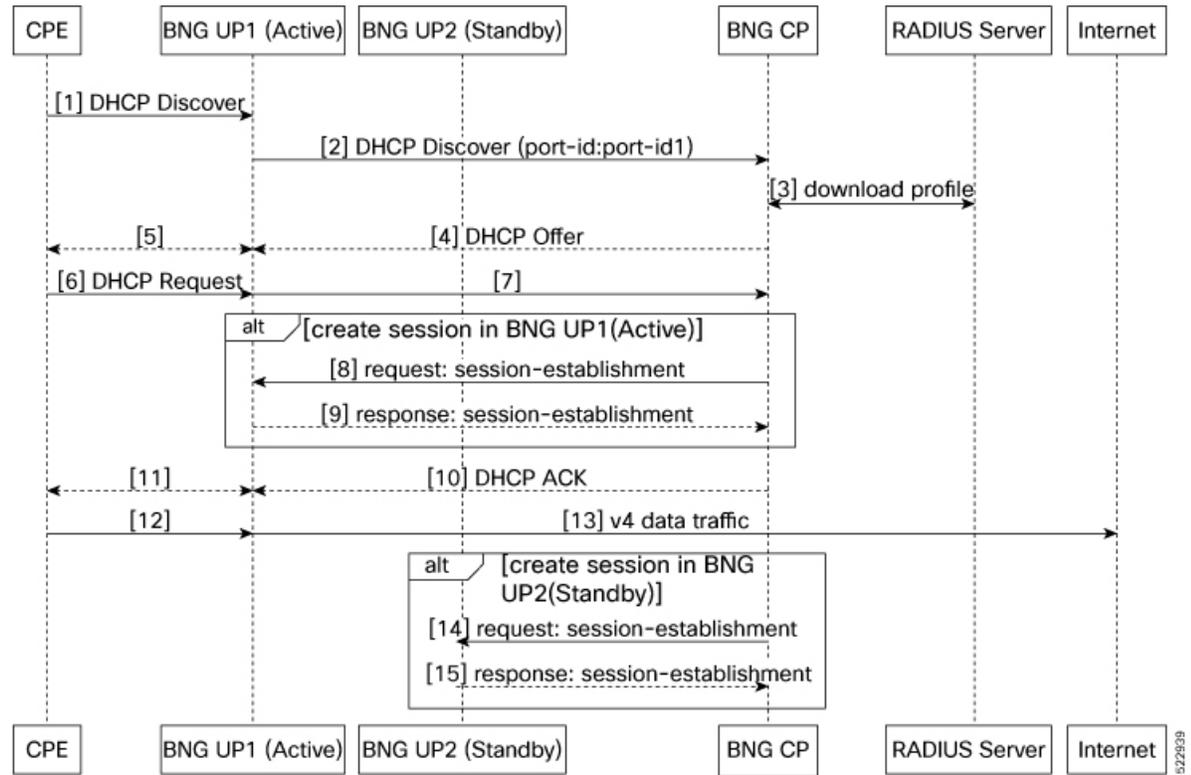
Figure 4: Topology for routed subscriber



Workflow

These stages describe the subscriber session creation for routed subscriber sessions:

Figure 5: Subscriber Session Creation for Routed Subscribers



1. Initial Packet Handling:

- a. A discover packet originates from the CPE and is sent to the BNG UP.
- b. CP authenticates the subscriber, downloads the subscriber's profile using RADIUS, and sends a DHCP offer to CPE upon successful authentication.
- c. CPE sends a DHCP request to the CP.

2. Session Establishment:

- a. After the DHCP request that comes from the CPE is received, CP sends a session creation request to UP1.
- b. UP1 creates a subscriber interface and acknowledges session creation. CP sends DHCP acknowledgment to the CPE via UP.

3. Replication and Standby: After the session is created on the active UP, it is replicated on the standby UP (UP2).

4. Outgoing Traffic: DHCP packets from the control plane use the access interface to send the packets towards CPE.

Restrictions for routed subscribers

These restrictions are applicable for routed subscribers:

- Supported only on the bundle main interface and not on other access interfaces and VLAN sub-interfaces.
- Supported only for V6 address family (IANA and IAPD). V4 AFI is not supported.
- For bandwidth capacity expansion, multiple interfaces would be needed in bundle interface and not as ECMP.
- The underlay network can support both IPv4 and IPv6, but routed subscriber sessions will only support IPv6 traffic.
- The next-hop IP address for the subscriber must remain constant and not be learned through a recursive route. To prevent disruptions, ensure the connectivity between access interface and next-hop IP interface doesn't change.

Guidelines for configuring routed subscribers

Follow these guidelines while configuring routed subscribers.

- If subscribers are deleted from BNG CP, DHCP bindings on relay devices should be cleared manually.
- Configuration changes for dynamic route next-hop are not allowed. Changes in configurations must be done after bringing down all the subscribers and clearing dynamic routes belonging to that SRG group.
- Route updates can lead to a change in the CPE reachability next-hop IP address. Changes to the next-hop IP address for existing subscribers are not supported.

Configure routed subscriber sessions

Enable routed subscriber sessions on the cnBNG UP and allow subscriber connections through a routed access network.

Follow these steps to configure routed subscriber sessions:

Procedure

Step 1 Enable routed subscriber sessions on an access interface.

Enable the configuration of prefix-length for Identity Association for Prefix Delegation (IAPD) IP addresses. Ensure that the IAPD route prefix length (**prefix-len**) matches the prefix length configured under the access interface, if not cnBNG CP request is rejected.

Note

Routed subscriber access interface configuration should only be applied to the bundle main interface.

Example:

```
Router#configure
Router(config)#interface bundle-ether 1
Router(config-if)#ipsubscriber
Router(config-cnbnng-nal-ipsub)#ipv6 routed
Router(config-cnbnng-nal-ipsub-ipv6-routed)#initiator dhcp prefix-len 20 src-ip-dual-lookup
Router(config-cnbnng-nal-ipsub-ipv6-routed)#
```

- Step 2** Configure state control next hop IP under SRG group to specify next-hop IP while installing state control routes and subscriber routes for routed subscribers.

The next-hop IP configured must be reachable by subscriber VRF. Next-hop support is limited to IPv6 address-family.

Example:

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)#subscriber-redundancy
Router(config-cnbng-nal-sub-red)#group group1
Router(config-cnbng-nal-srg-grp)#state-control-next-hop-ip ipv6 2002:4888:11:0:4:4:1:1
Router(config-cfg-srg-grp-intf)#exit
```

- Step 3** Execute the `show cnbng-nal srg-group group1 detail` command to verify the SRG group contains routed subscribers.

Example:

```
Router#show cnbng-nal srg-group group1 detail
Fri Jan 24 14:03:20.351 IST

=====
Location: 0/0/CPU0
=====

SRG group name                : group1
SRG group admin state         : UP_CP_Configured
SRG group state               : Up
SRG role                      : Active
SRG ID                        : 0x00000001
SRG VRF name                  : **srg_1
Last SRG role update time     : Jan 23 15:16:42.915881
Virtual mac                   : 0000.0000.0000
V4 Table Id                   : 0xe0000012
V6 Table Id                   : 0xe0800012
V4 Proto Id                   : 0x00000001
V6 Proto Id                   : 0x00000002
Subscriber count               : 1
IPV4 route count              : 0
IPV6 route count              : 5
Damping timer interval        : 120 Sec
Subnet route tag              : 0
Route export on Standby enable : False
Fast switchover enable        : True
Ready for role change         : Yes [Success]
FSM State                     : COMPLETE
Update Request State          : IDLE
Stats state                   : IDLE
Sub disconnect resp pend      : NA
Sub count for keep alive start : 0
Group for routed subscribers  : TRUE
```

- Step 4** Execute the `show cnbng-nal subscriber all detail` command to view the subscriber type and sub-type for routed subscribers.

Example:

```
Router#show cnbng-nal subscriber all detail
Thu Jan 23 15:24:17.857 IST
=====
Location: 0/0/CPU0
=====
Interface:                Bundle-Ether1.ip2147483664
UPID:                     0x80000010
```

```

CPID: 0x00000002
Type: IPoE
  Sub-type: Routed
    Routed IPv4 Prefix: 0.0.0.0
    Routed IPv6 Prefix: 1::2
IPv4 Address: 0.0.0.0
IPv4 Framed Route: NA
IPv6 IANA Address: cafe::bad1
IPv6 IAPD Prefix: cafe:bed1::/64
IPv6 Slaac Prefix: ::/0
CPE link local Address: ::
IPv6 Framed Route: NA
IPv6 State: UP, Thu Jan 23 15:16:49 2025

```

Step 5 Execute the `show cnbng-nal subscriber all summary` command to view the classification of IPoE subscribers as L2 connected and routed subscribers.

Example:

```
Router#show cnbng-nal subscriber all summary
```

```

=====
Location: 0/0/CPU0
=====

```

Type	PPPoE	IPoE		LAC	LNS
====	=====	L2-Conn	Routed	===	===
		=====	=====		
Session Counts by State:					
initializing	0	0	0	0	0
connecting	0	0	0	0	0
connected	0	0	0	0	0
activated	0	0	1	0	0
idle	0	0	0	0	0

Step 6 Optionally, use the `clear cnbng-nal subscriber sub-type ipoe-routed` command if you want to clear the routed subscriber session records.

Example:

```
Router# clear cnbng-nal subscriber sub-type ipoe-routed
```

SRG warm-standby mode

Subscriber Redundancy Group (SRG) warm-standby is a subscriber redundancy mode that

- pre-creates and maintains subscriber session context in process memory on standby User Plane (UP) nodes, without programming sessions into the hardware
- enables rapid activation of subscriber services after a switchover event, and

- improves cost effectiveness and scalability by allowing a single standby UP node to provide redundancy for multiple UP instances.

Table 3: Feature history

Feature Name	Release Information	Description
SRG warm-standby mode	Release 25.4.1	You can now reduce failover recovery time and lower costs by using a single standby UP node for multiple UP instances. The standby node keeps subscriber session context in memory (without programming hardware), then quickly replays it to hardware during switchover, enabling fast service restoration, better hardware utilization, and improved scalability.

Session management between active and standby nodes

In SRG warm-standby mode, when a subscriber session is established, the Control Plane (CP) ensures that both the active and standby UP nodes are aware of the subscriber. However, only the active node programs subscriber sessions into hardware, while the standby node holds the session context in memory, without programming it into hardware.

Switchover process and service restoration

During a switchover event, the standby node takes over by rapidly programming the cached subscriber session contexts, previously stored only in memory, into its hardware. Since these sessions are already pre-authenticated and have assigned IP addresses, the process is significantly faster than a full session creation. However, there is a brief delay as the sessions must be programmed onto hardware during the switchover. This approach minimizes service disruption, restores services within a few seconds, and is especially effective for PPPoE subscriber sessions, in line with standards such as BBF TR-459.

Supported session types

SRG warm-standby mode supports only subscriber sessions that use Point-to-Point Protocol over Ethernet (PPPoE).

Comparison: Warm standby vs. hot standby vs. non-redundant deployment

Here's a comprehensive comparison of warm-standby, hot-standby, and non-redundant deployment scenarios:

Aspect	Warm standby	Hot standby	Non-redundant deployment
Session storage	Session context stored in memory only (not pre-programmed in hardware on standby node)	Sessions pre-programmed in hardware on standby node	Sessions exist only on the active node
Hardware utilization	High efficiency: active node uses full hardware capacity for active sessions; standby node's hardware is not reserved for session programming until switchover	Lower efficiency: both active and standby nodes reserve hardware capacity for sessions (typically 50% each)	Highest efficiency for a single node, but no redundancy or failover capability
Switchover process	On failure, standby node quickly programs session context from memory into hardware, then resumes service	On failure, standby node takes over instantly since sessions are hardware-ready	No switchover possible; service is interrupted until manual recovery

Aspect	Warm standby	Hot standby	Non-redundant deployment
Restoration time	A few seconds (some delay due to hardware programming at switchover)	Nearly instantaneous (minimal/no delay)	Prolonged outage; service only resumes after manual intervention
Service continuity	High, with minimal disruption (brief delay during switchover)	Highest, with almost seamless continuity	None; subscribers lose service during outages
Scalability	High—standby node can hold large number of session contexts in memory, supporting multiple active nodes	Limited—standby node's capacity matches active; hardware reserved for standby sessions	Limited to hardware capacity of a single node
Cost efficiency	More cost-effective; standby resources shared, maximizing hardware utilization	Less cost-effective; redundant hardware must be reserved for standby	Most cost-effective for initial deployment, but high operational risk
Best use case	Environments requiring high scalability and efficient hardware use with acceptable brief failover delay	Scenarios where the fastest possible failover and minimal disruption are critical	Cost-sensitive environments where redundancy is not required and brief outages are acceptable

Configure SRG warm-standby mode

Enable SRG warm-standby mode to provide rapid subscriber failover and service restoration in cnBNG networks.

Procedure

Step 1 Use the `standby-mode warm` command in user-plane configuration mode on the CP to enable SRG warm-standby mode:

Example:

```
user-plane
instance 1
  user-plane asr9k-1
  peer-address ipv4 1.1.1.1
  subscriber-profile subs-upf
  subscriber-redundancy
  group Group1
  standby-mode warm
  exit
exit
exit
exit
exit
```

Step 2 Use the `show subscriber redundancy srg-peer-id` command on the CP to verify the SRG warm-standby configuration.

Example:

```
bng# show subscriber redundancy srg-peer-id Peer1 detail
Tue Mar 18 12:19:14.448 UTC+00:00
subscriber-details
```

```

{
  "subResponses": [
    {
      "PeerID": "Peer1",
      "GroupID": "Group1",
      "Domain": "d1",
      "UP List": {
        "asr9k-1": {
          "N4 State": "Connected",
          "Srg State": "Up",
          "Srg Role": "Active",
          "Interface map": {
            "BundleEther1.100": 1,
            "BundleEther1.200": 2
          }
        },
        "asr9k-2": {
          "N4 State": "Connected",
          "Srg State": "Up",
          "Srg Role": "Standby",
          "Interface map": {
            "BundleEther1.100": 1,
            "BundleEther1.200": 2
          },
          "Standby Mode": "Warm"
        }
      },
      "Preferred Active": "asr9k-1"
    }
  ]
}

```

Step 3 Use the **show cnbng-nal srg-group** and **show cnbng-nal subscriber warm** commands on the UP to monitor and troubleshoot SRG warm-standby functionality.

Example:

```

Router# show cnbng-nal srg-group srg-mode warm
Thu Nov 27 11:28:09.151 IST

```

```

=====
Location: 0/0/CPU0
=====

```

Group-name	SRG role	Access OT	Core OT	Subs Count	V4 routes	V6 routes
group1	Active	Up	Up	1	2	2
group2	Standby	Up	Up	0	2	2
group3	Active	Up	Up	0	2	2

```
Total Entries : 3
```

```
Summary
-----
```

Category	Total	Active	Standby	None
Groups	3	2	1	0
Subscribers	1	1	0	0
V4 subnet routes	6	4	2	0
V6 subnet routes	6	4	2	0

```

Router# show cnbng-nal subscriber warm cpid 2 location 0/0/CPU0
Thu Nov 27 11:29:40.870 IST

```

■ Configure SRG warm-standby mode

```
=====  
Location: 0/0/CPU0  
=====  
Codes: IA - Inactive, AG - Activating, AC - Active,  
  
CPID(hex)  State  Mac Address      Subscriber IP Addr / Prefix (Vrf)  
-----  
2          AC    1234.1234.4555  10.0.0.2 default  
                                     10.0.0.0/24 (Framed IPv4)  
                                     10.0.4.0/24 (Framed IPv4)  
                                     befe::bad1 (IANA)  
                                     befe:bed1::/64 (IAPD)  
                                     beca::/64 (Slaac PD)  
                                     2001:db8:2:/24 (Framed IPv6)  
                                     2001:db8:4:/24 (Framed IPv6)
```
