

Cloud Native BNG User Plane Configuration Guide for Cisco ASR 9000 Series Routers, IOS XR Release 26.1.x, 26.2.x, 26.3.x

First Published: 2026-02-28

Last Modified: 2026-09-11

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface vii

Changes to This Document vii

Communications, Services, and Additional Information vii

CHAPTER 1

New and Changed Cloud Native BNG User Plane Features 1

Cloud Native BNG User Plane Features Added or Modified in IOS XR Release 26.x.x 1

CHAPTER 2

YANG Data Models for Cloud Native BNG User Plane Features 3

Using YANG Data Models 3

CHAPTER 3

Cloud Native BNG Overview 5

Overview 5

Evolution of cnBNG 6

cnBNG Architecture 6

cnBNG Components 8

Subscriber Microservices Infrastructure 8

cnBNG Control Plane 9

cnBNG User Plane 10

License Information 10

Standard Compliance 11

Limitations and Restrictions 11

CHAPTER 4

Cloud Native BNG User Plane Overview 13

Control and User Plane Separation 13

cnBNG User Plane Overview 15

cnBNG User Plane Architecture 16

cnBNG User Plane and Control Plane Reachability 17
Software and Hardware Requirements 18
Access Types and Subscriber Types 19
Subscriber Features 20
High Availability 23
Usage Guidelines 24
Restrictions 24

CHAPTER 5 **Installing Cloud Native BNG User Plane Packages** 25
Installing and Activating the cnBNG Package on the User Plane 25

CHAPTER 6 **Configuring Cloud Native BNG User Plane and Key Features** 27
Configure cnBNG User Plane 27
Configure Basic User Plane Settings 28
Configure Access-Interface 30
Configure Loopback Interface 31
Configure DHCP 35
Configure Subscriber Gateway Address and Subnet Route 36
Configure Route Summary 39
Export Routes to Core Network 39
Configure ARP Scale Mode 40
Configure Cloud Native BNG over Pseudowire Headend 41
Multiple Control Plane associations 43
Restrictions for multiple CP associations 44
Configure and verify multiple CP associations for a UP 44
Configure multiple CP associations for a UP 44
Verify multiple CP associations to a UP 47
Adding a new CP-UP association without disrupting existing connections 66
Percentage-based QoS allocation for shared policy instances 66
Restrictions for percentage-based QoS allocation for shared policy instances 67
How SPI use percentage-based QoS allocation 67
Enable percentage-based QoS allocation for SPI 68
Verify cnBNG User Plane Configuration 70
Verify cnBNG NAL Process Information 70

Verify Control Plane Connection Status 71
Verify Subscriber Information 72
Verify cnBNG NAL Counters 81
P-BNG and cnBNG coexistence 84
Limitations of P-BNG and cnBNG coexistence 85
Enable P-BNG and cnBNG coexistence 85

CHAPTER 7 **Subscriber Management** 91
Subscriber Session Overview 91
IPoE Session 92
PPP over Ethernet (PPPoE) 93
PPPoE PPP Terminated and Aggregation Sessions (PPPoE-PTA) 94
Configure PPPoE-PTA Session 94
L2TP Access Concentrator Sessions (LAC) 95
Limitations for LAC Sessions 97
Configure LAC Sessions 98
L2TP Network Server Sessions (LNS) 99
Limitations for LNS Sessions 101
Configure LNS Sessions 101
Enable SLAAC for PPPoE Subscriber Sessions 102
Configure SLAAC for PPPoE Subscriber Sessions 104
RADIUS-Based Policing - QoS shape-rate parameterization 106
Limitations of configuring RADIUS-Based Policing 107
Configure QoS shape-rate parameterization 107
Shared Policy Instance 110
Limitations of configuring Shared Policy Instance 110
Configure Shared Policy Instance 111
Configure a QoS policy with SPI using a feature template 111
Configure a QoS policy with SPI using a RADIUS server 113
Enable percentage-based QoS allocation for SPI 115
Enhanced subscriber routing with Framed-Route Tag and Preference attributes 117
Processing multi-homed subscriber routing scenarios 118
Backward compatibility considerations 119
Restrictions for configuring Framed-Routes tag and preference attributes 119

Configure Framed-Route preference support 120
References 121
Geographical redundancy for L2TP sessions 122
Limitations of configuring geographical redundancy for L2TP sessions 122
Configure CP-GR support for L2TP tunnels and sessions 122
Synchronize L2TP tunnels between CP and UP for CP-GR 125

CHAPTER 8 **Geo Redundancy (Subscriber Redundancy Group) 127**
Overview 128
SRG Modes 129
Subscriber Session Set up Call Flow 133
Benefits of BNG Geo Redundancy 135
Supported Features in BNG Geo Redundancy 136
Unsupported Features and Restrictions for cnBNG Geo Redundancy 136
Guidelines to Configure SRG 137
Configure SRG 138
Routed subscriber sessions 144
How routed subscriber sessions work 147
Restrictions for routed subscribers 149
Configure routed subscriber sessions 150
SRG warm-standby mode 152
Configure SRG warm-standby mode 154

CHAPTER 9 **IPV6 Neighbor Discovery 157**
IPv6 Neighbor Discovery 157
Configure IPv6 Neighbor Discovery 159



Preface

The Cisco IOS XR Software Release 7.3.1 introduces the support for cloud native broadband network gateway (cnBNG) user plane for the Cisco IOS XR platform. cnBNG is an architectural evolution that is based on Control and User Plane Separation (CUPS), where the control plane (CP) and user plane (UP) run in distinct and independent environments. This book describes the cnBNG user plane functionality and related configurations on Cisco ASR 9000 Series Routers.

For details on the commands related to the cnBNG user plane, see the Cloud Native Broadband Network Gateway User Plane Command Reference for Cisco ASR 9000 Series Routers.

For details on cnBNG deployment, the control plane functionality and the related configurations, see the *Cloud Native Broadband Network Gateway Control Plane Configuration Guide*.

To know more about physical BNG on the Cisco IOS XR platform, see the *Cisco ASR 9000 Series Aggregation Services Router Broadband Network Gateway Configuration Guide*.

This preface contains these sections:

- Changes to This Document, on page vii
- Communications, Services, and Additional Information, on page vii

Changes to This Document

Date	Summary
September 2026	Republished for Release 26.3.1.
June 2026	Republished with feature updates for Release 26.2.1.
February 2026	Initial release of this document.

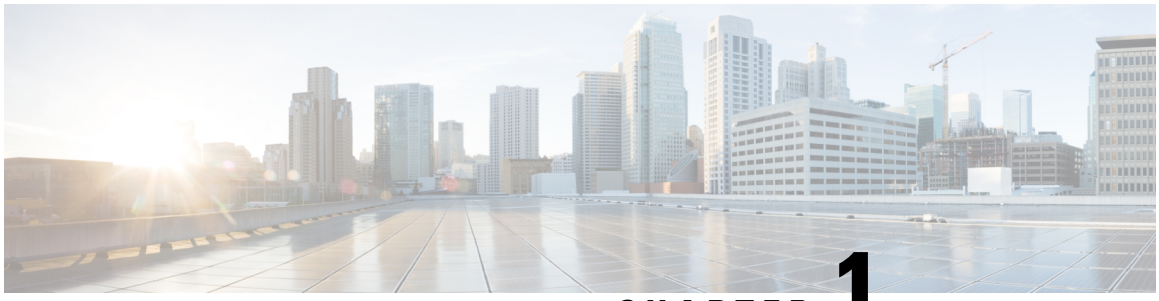
Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.

- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

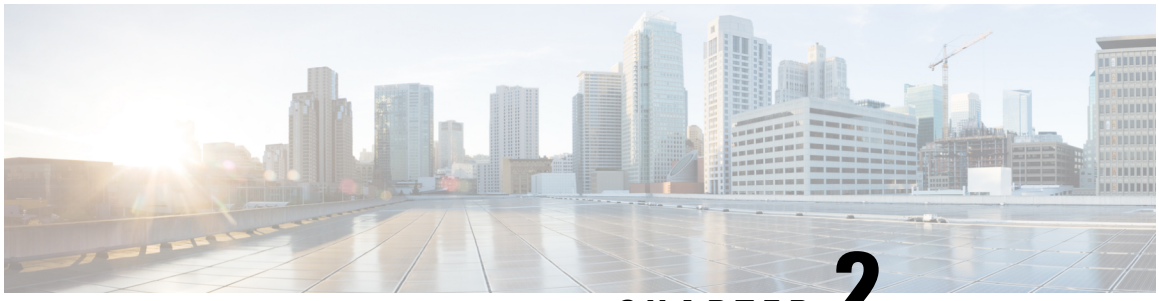
New and Changed Cloud Native BNG User Plane Features

This table summarizes the new and changed feature information for the *Cloud Native BNG User Plane Configuration Guide for Cisco ASR 9000 Series Routers*, and tells you where they are documented.

- Cloud Native BNG User Plane Features Added or Modified in IOS XR Release 26.x.x , on page 1

Cloud Native BNG User Plane Features Added or Modified in IOS XR Release 26.x.x

Feature	Description	Changed in Release	Where Documented
P-BNG and cnBNG coexistence	This feature is introduced.	Release 26.2.1	P-BNG and cnBNG coexistence, on page 84
Multiple Control Plane associations	This feature is introduced.	Release 26.2.1	Multiple Control Plane associations, on page 43
Routed subscriber sessions	This feature is modified.	Release 26.2.1	Routed subscriber sessions, on page 144



CHAPTER 2

YANG Data Models for Cloud Native BNG User Plane Features

This chapter provides information about the YANG data models for Cloud Native BNG User Plane features.

- Using YANG Data Models, on page 3

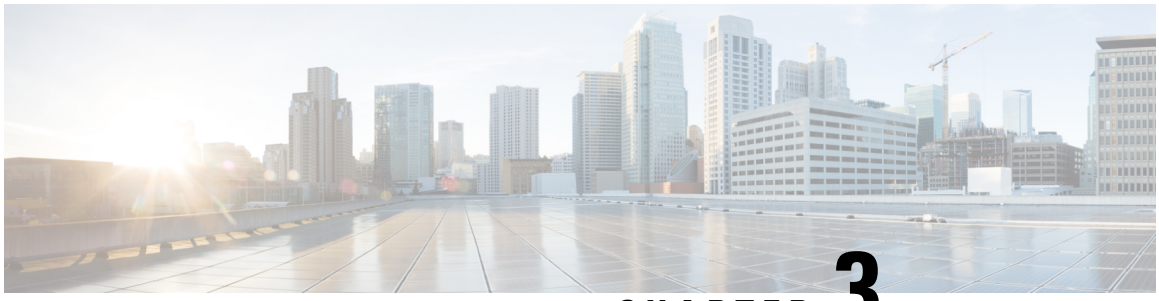
Using YANG Data Models

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the Github repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPaths. To view a comprehensive list of the data models supported in a release, navigate to the **Available-Content.md** file in the repository.

You can also view the data model definitions using the YANG Data Models Navigator tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.



CHAPTER 3

Cloud Native BNG Overview

The Cloud Native Broadband Network Gateway (cnBNG) redefines the traditional physical BNG by decoupling the subscriber management and forwarding functions of the control plane (CP) and user plane (UP) to give better flexibility and scalability for the service providers. The cnBNG architecture is based on Control and User Plane Separation (CUPS), where the CP performs the policy and charging rule function (PCRF), whereas the UP performs policy enforcement function (PEF) of the overall BNG subscriber management solution. The cnBNG solution provides optimum scale dimensioning in terms of the number of subscriber sessions and forwarding capacity and aims at rapid deployment of multi-access services for the users. It also acts as a step forward towards converging the fixed line and mobile networks at all network layers.

- Overview, on page 5
- License Information, on page 10
- Standard Compliance, on page 11
- Limitations and Restrictions, on page 11

Overview

The Broadband Network Gateway (BNG) is the access point for subscribers, through which they connect to the broadband network. When a connection is established between BNG and Customer Premise Equipment (CPE), the subscriber can access the broadband services provided by the Network Service Provider (NSP) or Internet Service Provider (ISP).

BNG establishes and manages subscriber sessions. When a session is active, BNG aggregates traffic from various subscriber sessions from an access network, and routes it to the network of the service provider.

BNG is deployed by the service provider and is present at the first aggregation point in the network, such as the edge router. An edge router, like the Cisco ASR 9000 Series Router, needs to be configured to act as the BNG. Because the subscriber directly connects to the edge router, BNG effectively manages subscriber access, and subscriber management functions such as:

- Authentication, Authorization, and Accounting (AAA) of subscriber sessions
- Address assignment
- Security
- Policy management
- Quality of Service (QoS)

Implementing the BNG provides the following benefits:

- Communicates with authentication, authorization, and accounting (AAA) server to perform session management and billing functions besides the routing function. This feature makes the BNG solution more comprehensive.
- Provides different network services to the subscriber. This enables the service provider to customize the broadband package for each customer based on their needs.

Cisco provides two BNG solutions:

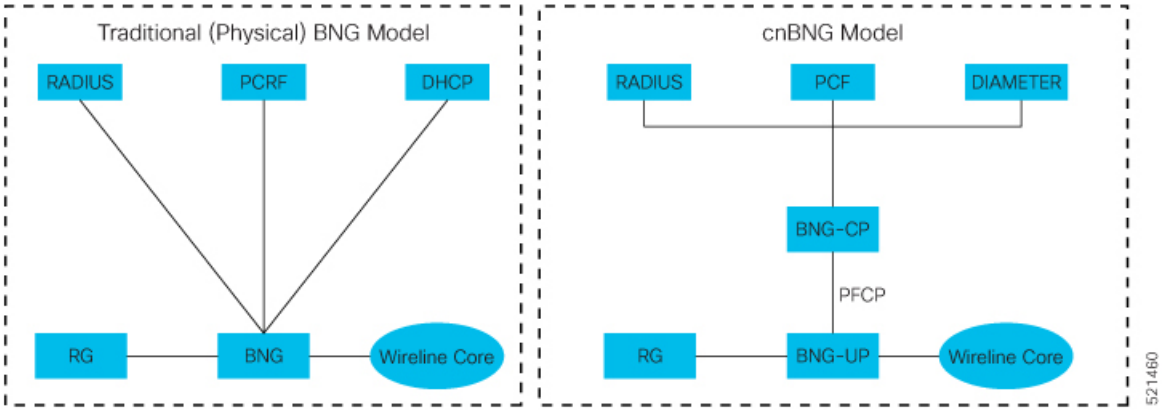
- **Physical BNG** where the BNG Control Plane (CP) and the User Plane (UP) are tightly coupled inside a Cisco IOS XR platform where the CP runs on an x86 CPU and the UP runs on a physical NPU or ASIC.

  For more information about the physical BNG, refer to the latest version of the *Broadband Network Gateway Configuration Guide* for Cisco ASR 9000 Series Routers.
- **Virtual BNG (vBNG)** where the BNG CP and UP run in separate VM-based Cisco IOS XR software on general purpose x86 UCS servers.

Evolution of cnBNG

The Cisco Cloud Native Broadband Network Gateway (cnBNG) provides a new dimension to the Control Plane and User Plane Separation (CUPS) architecture of the Broadband Network Gateway (BNG), enabling flexibility and rapid scaling for Internet Service Providers (ISPs).

Figure 1: Evolution of BNG to cnBNG



The architectural change is an evolution from an integrated traditional BNG running on a single router to a disaggregated solution, where the centralized subscriber management runs on an elastic and scalable Cloud Native Control Plane (CP) and the User Plane (UP) delivers the forwarding functionality.

cnBNG Architecture

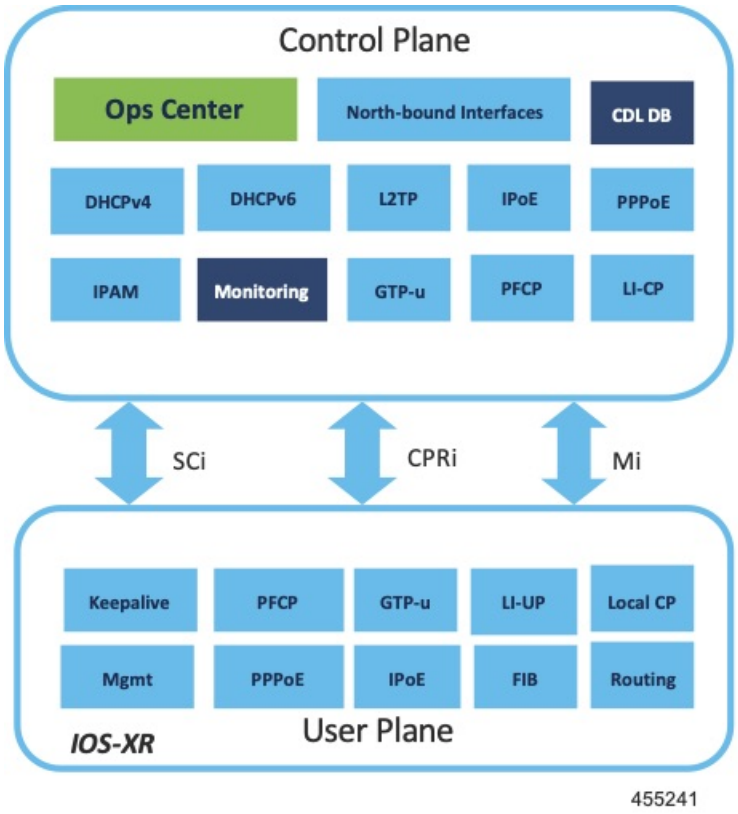
In the cnBNG architecture, the CPs and UPs are clearly and cleanly separated from each other and run in completely distinct and independent environments.

The BNG CP is moved out to a container-based microservice cloud environment.

The UP can be on any of the physical platforms that supports the BNG UP, like Cisco ASR 9000 Series Routers.

The following figure illustrates the overall cnBNG architecture.

Figure 2: cnBNG Architecture



Features and Benefits

The cnBNG supports the following features:

- **Path to convergence**: With shared Subscriber Management infrastructure, common microservices across the policy layer and shared UPs for BNG and Mobile back-haul, cnBNG paves the way for real Fixed Mobile Convergence (FMC).
- **Flexibility of scaling**: cnBNG architecture provides flexibility by decoupling the required scalability dimensions. The CP can be scaled with requirement of number of subscribers to be managed and UPs can be augmented based on the bandwidth requirements. Instead of building the CP for peak usage, the orchestrator can be triggered to deploy the relevant microservices as needed to handle the increased rate of transactions.
- **Distributed UPs**: With reduced operational complexity and minimal integration efforts with centralize CP, UPs can be distributed, closer to end-users to offload traffic to nearest peering points and CDNs. This feature reduces the core transport costs.
- **Cost effective and Leaner User planes**: With the subscriber management functions moved to cloud, you can choose cost-effective UP models for optimized deployment requirements.

The benefits of the cnBNG architecture are:

- Simplified and unified BNG CP
- Platform independent and Network Operation System (NOS) agnostic BNG CP
- Unified Policy interface across both BNG and mobility
- Common infrastructure across wireline and mobility
- Seamless migration from existing deployments
- Leverage the common infrastructure across access technologies
- Standardized model driven interface with the UP
- Data externalization for North-bound interfaces (NBI)
- Highly available and fault tolerant
- Simplified Subscriber Geo redundancy
- Horizontally scalable CP
- Independent CP and UP upgrades
- Feature agility with CI and CD
- Manageability and Operational Simplification

cnBNG Components

The cnBNG solution comprises of the following components:

Subscriber Microservices Infrastructure

The Cisco Ultra Cloud Core Subscriber Microservices Infrastructure (SMI) is a layered stack of cloud technologies that enable the rapid deployment, and seamless life-cycle operations for microservices-based applications.

The SMI stack consists of the following:

- SMI Cluster Manager—Creates the Kubernetes (K8s) cluster, creates the software repository, and provides ongoing LCM for the cluster including deployment, upgrades, and expansion.
- Kubernetes Management—Includes the K8s master and etcd functions, which provide LCM for the NF applications deployed in the cluster. This component also provides cluster health monitoring and resources scheduling.
- Common Execution Environment (CEE)—Provides common utilities and OAM functionalities for Cisco cloud native NFs and applications, including licensing and entitlement functions, configuration management, telemetry and alarm visualization, logging management, and troubleshooting utilities. Additionally, it provides consistent interaction and experience for all customer touch points and integration points in relation to these tools and deployed applications.
- Common Data Layer (CDL)—Provides a high performance, low latency, stateful data store, designed specifically for 5G and subscriber applications. This next generation data store offers HA in local or geo-redundant deployments.

- Service Mesh—Provides sophisticated message routing between application containers, enabling managed interconnectivity, additional security, and the ability to deploy new code and new configurations in low risk manner.
- NB Streaming—Provides Northbound Data Streaming service for billing and charging systems.
- NF/Application Worker nodes—The containers that comprise an NF application pod.
- NF/Application Endpoints (EPs)—The NF's/application's interfaces to other entities on the network.
- Application Programming Interfaces (APIs)—SMI provides various APIs for deployment, configuration, and management automation.

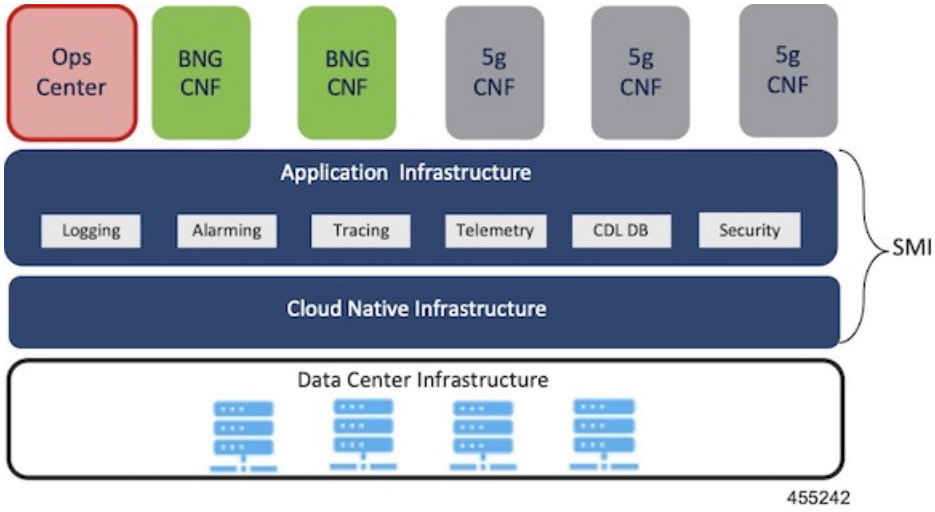
For more information on SMI components, refer to the "Overview" chapter of the *Ultra Cloud Core Subscriber Microservices Infrastructure* documentation—*Deployment Guide*.

For information on the Cisco Ultra Cloud Core, see https://www.cisco.com/c/en/us/products/collateral/wireless/packet-core/datasheet-c78-744630.html.

cnBNG Control Plane

The Cisco cnBNG CP is built on Cisco® Cloud Native Infrastructure, which is a Kubernetes-based platform that provides a common execution environment for container-based applications. This CP is built on principles of stateless microservices, to scale at-ease, introduce services much faster and more cost-effective.

Figure 3: cnBNG Control Plane Architecture



The CP runs as a Virtual Machine (VM) to adapt to existing service provider-deployed virtual infrastructure. It is built ground-up on a clean-slate architecture with a view on ‘Converged Subscriber Services’ and is aligned to 3gpp and BBF standards.

The cnBNG CP effectively manages the subscriber management functions such as:

- Authentication, authorization, and accounting of subscriber sessions
- IP Address assignment
- In-built DHCP Server

- Security
- Policy management
- Quality of Service (QoS)

Service providers can choose from wide choice of available ASR 9000 form factors, based on exact deployment requirements. The CUPS architecture allows to run these UPs in a distributed mode, to the edge of network, for early traffic offloads.

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

cnBNG User Plane

The UP delivers the forwarding functionality of the entire cnBNG solution. With the CP handling the subscriber management functionality, the cnBNG architecture enables the UP to be more distributed and interoperable with cnBNG CP with minimal integration efforts. The cnBNG Subscriber Provisioning Agent (SPA), which is the common interface between UP and CP, is bundled with the existing Cisco IOS XR image to transform an integrated physical BNG router to a cnBNG user plane.

For more information about the cnBNG UP, see the *Cloud Native BNG User Plane Overview* chapter.

License Information

cnBNG supports the following licenses:

License	Description
Application Base	Per cluster
Session (Increments)	Network-wide

These are the software license PIDs for cnBNG:

Cisco cnBNG Control Plane:

Product IDs	Description
CN-BNG-BASE-L	Base PID for cnBNG Control Plane (per cluster)
CN-BNG-100k-L	Session scale for 100,000 subscribers (network-wide) base licenses
CN-BNG-400k-L	Session scale for 400,000 subscribers (network-wide) base licenses
CN-BNG-1M-L	Session scale for 1,000,000 subscribers (network-wide) base licenses
CN-BNG-2M-L	Session scale for 2,000,000 subscribers (network-wide) base licenses

Cisco cnBNG User Planes:

Refer the ASR9000 data sheet for ordering information:
https://www.cisco.com/c/en/us/products/routers/asr-9000-series-aggregation-services-routers/datasheet-listing.html

Standard Compliance

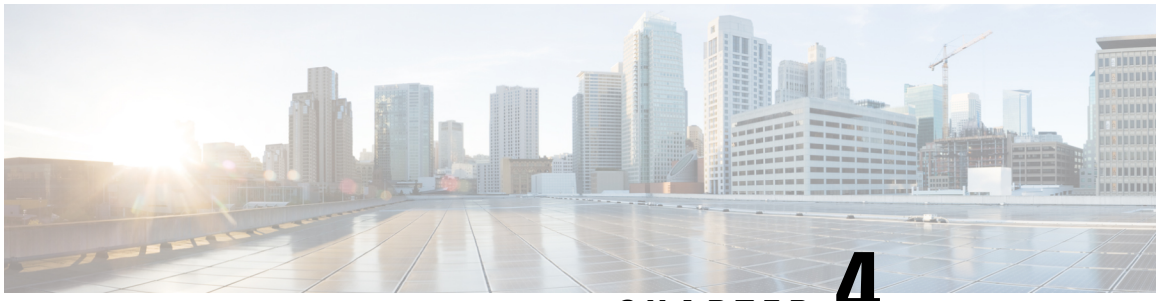
cnBNG solution is aligned with the following standard:

TR-459 Control and User Plane Separation for a disaggregated BNG

Limitations and Restrictions

The cnBNG has the following limitations and restrictions in this release:

- Only one subnet is supported per VRF.
- QoS provisioning is supported only through service.



CHAPTER 4

Cloud Native BNG User Plane Overview

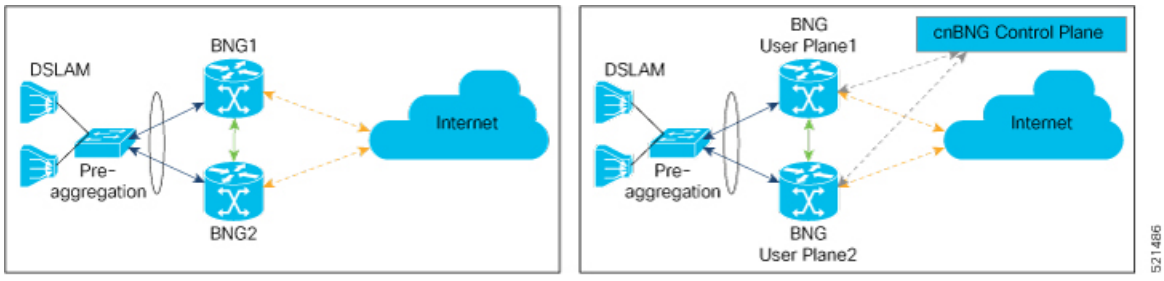
In the cnBNG architecture, which is based on Control and User Plane Separation (CUPS), the CP handles the subscriber management functionality and the UP handles the forwarding functionality of the entire BNG solution. This chapter focuses on the functionality and architecture of the cnBNG user plane.

For more details on the cnBNG control plane, see the *Cloud Native Broadband Network Gateway Control Plane Configuration Guide*.

- Control and User Plane Separation, on page 13
- cnBNG User Plane Overview, on page 15
- cnBNG User Plane Architecture, on page 16
- cnBNG User Plane and Control Plane Reachability, on page 17
- Software and Hardware Requirements, on page 18
- Access Types and Subscriber Types, on page 19
- Subscriber Features, on page 20
- High Availability, on page 23
- Usage Guidelines, on page 24
- Restrictions, on page 24

Control and User Plane Separation

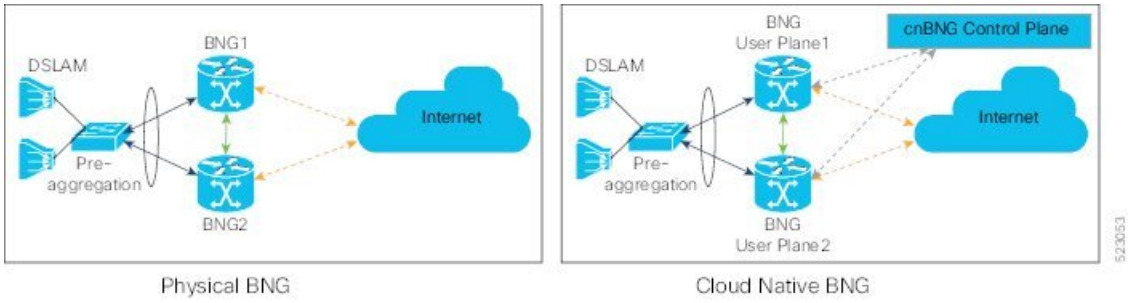
cnBNG is an architectural evolution that is based on Control and User Plane Separation (CUPS), where the CP and UP run in distinct and independent environments. cnBNG redefines the traditional physical BNG by decoupling the BNG CP and UP functions to give better flexibility and scalability for the service providers. In cnBNG, the centralized subscriber management functionality of BNG runs on CP infrastructure and the user plane delivers the forwarding functionality.



Physical BNG

Cloud Native BNG

Figure 4: Eps



In Cisco cnBNG solution, a physical Cisco IOS XR platform like Cisco ASR 9000 Series Routers provides the UP functionality. Whereas Cisco Ultra Cloud Core Subscriber Microservices Infrastructure (SMI)—a container-based microservice cloud environment, provides the CP functionality.

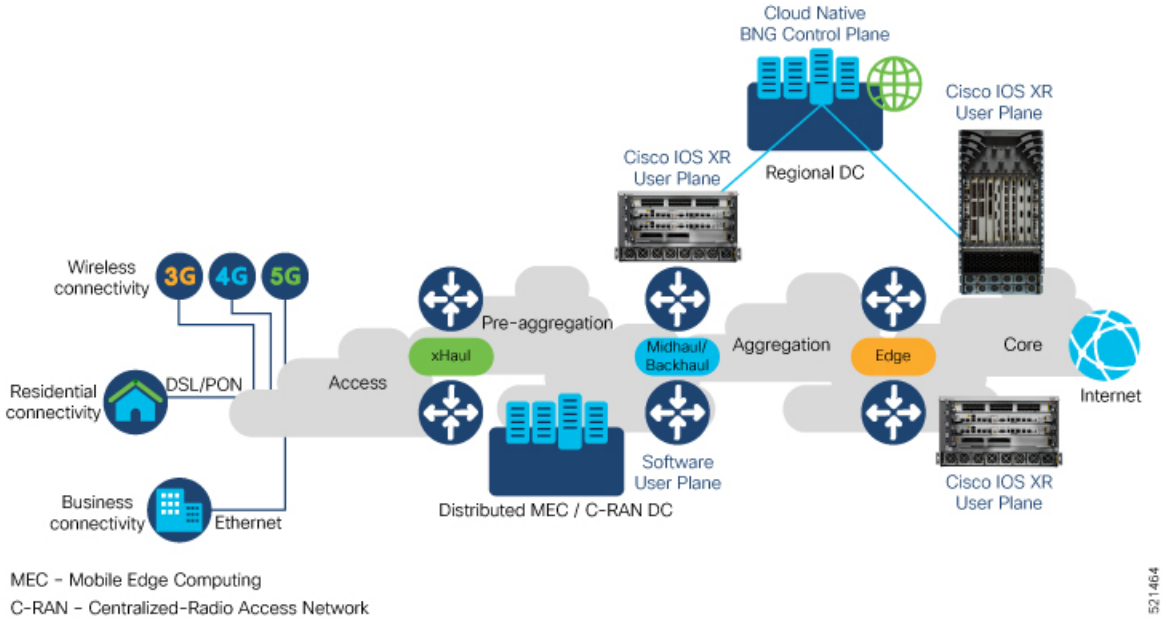
Why CUPS?

CUPS provides the capability to independently scale the CP and UP in an efficient and dynamic manner. CUPS enables network operators to optimize data center costs by hosting the CP and UP in different geographic locations. CUPS thus saves on backhaul (the access to core connection) costs by terminating data at the edge of the network. The network operators can then easily adapt to the evolving demands of mobile networks without incurring extra capital expenditures (CapEx) and operating expenditures (OpEx). The CUPS solution thus promotes a more cost-effective approach to core mobile architecture and future-proofs the network for 5G.

cnBNG User Plane Overview

Sample Network Topology for Cloud Native BNG using CUPS

Figure 5: EPS



In cloud native BNG (cnBNG), the CP provides the service policies that are sourced from the north-bound systems such as the RADIUS server or the policy and charging rules function (PCRF) node. Whereas the UP performs policy enforcement function (PEF) of the overall BNG subscriber management solution. The BNG CP protocols: RADIUS, DHCPv4, DHCPv6, PPPoE, PPP, and IP address pool management run on the CP. Whereas the non-BNG-specific protocols: IPv6 neighbor discovery (ND), ARP, routing protocols (like ISIS or BGP) that export subscriber subnet routes, and UDP or IP protocols that transport DHCPv4 or DHCPv6 payloads run on the UP.

The cnBNG UP models each subscriber as a unique flow. The system applies the subscriber features like quality of service (QoS), Hierarchial Quality of Service (HQoS),access control list (ACL), policy-based routing (PBR), lawful intercept (LI), accounting, and so on, on this flow. The DHCPv4, DHCPv6, PPPoE, and PPP protocols trigger the BNG subscriber flow. The UP presents these protocol packets to the cnBNG CP for authentication and authorization, and for evaluating policy and charging rules. Once the subscriber is accepted, the UP creates the subscriber flow and applies features on this flow. The subscriber flow can also have multiple sub-flows, and you can apply specific features to these sub-flows.

Key Features and Benefits of cnBNG User Plane

The key features and benefits of the cnBNG user plane are:

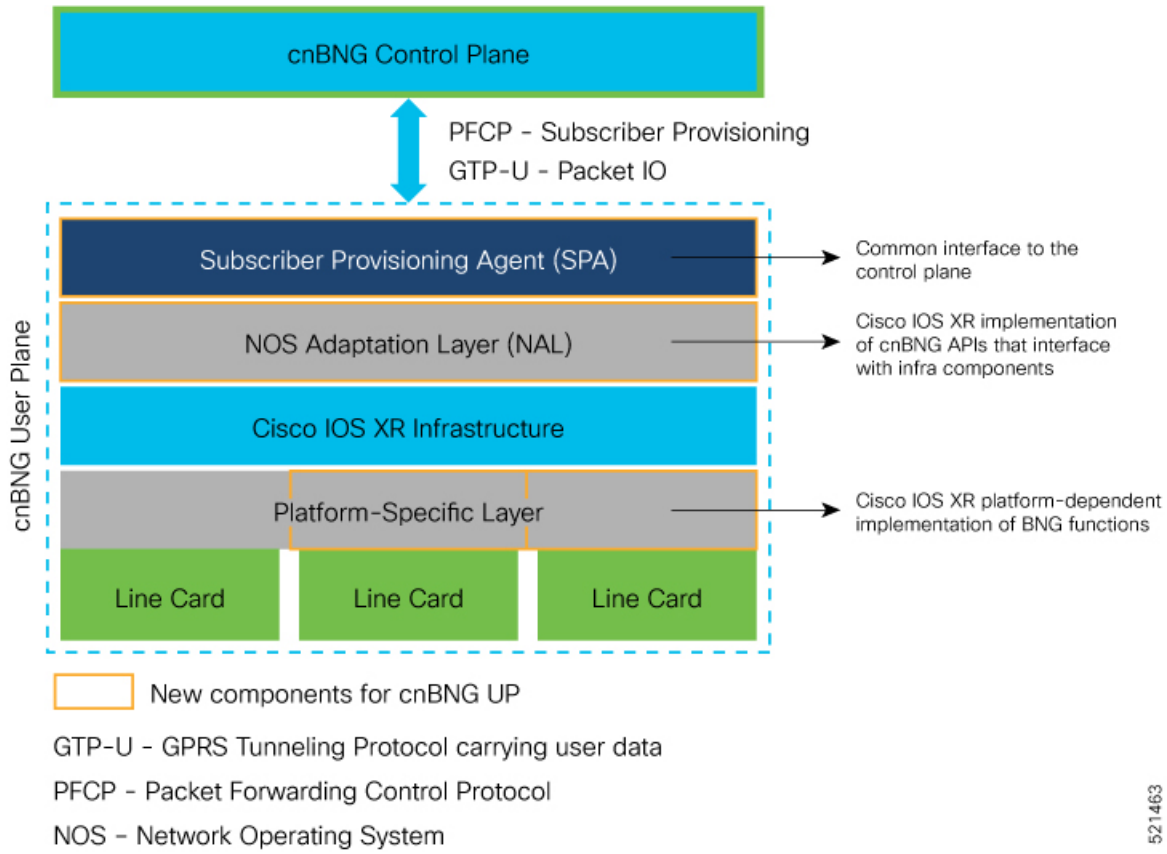
- **Distributed**: With reduced operational complexity and minimal integration efforts with centralized CP, you can distribute the UPs closer to end users. This feature helps to offload the traffic to the nearest peering points and content delivery networks (CDNs) and reduces the core transport costs.

- **Cost-effective and leaner**: With the subscriber management functions moved to cloud, you can choose cost-effective UP models for optimized deployment requirements.

cnBNG User Plane Architecture

The Cisco IOS XR platforms have a distributed hardware architecture that uses a switch fabric to interconnect a series of chassis slots. Each slot can hold one of several types of line cards (LCs). Each line card in these routers has integrated input-output and forwarding engines. The system can identify and handle the subscriber flow either on the route processor (RP) or on the LC. This architecture thereby provides multiple levels of redundancy and scalability for the subscriber management functionality in cnBNG.

Figure 6: Cloud Native BNG User Plane Architecture



The cnBNG UP architecture is designed to interoperate with cnBNG CP with minimal integration efforts. The main components of cnBNG user plane on the Cisco IOS XR platform are:

- **Subscriber Provisioning Agent (SPA)**—is the common interface to the control plane that is bundled with the existing Cisco IOS XR image. This interface helps to have a minimal configuration requirement to transform from an integrated physical BNG router to a cnBNG user plane. SPA consists of a transport layer at the top that interfaces with the CP, and an API layer at the bottom that isolates the network operating system (NOS) and the CP. This isolation from the NOS helps to make the control plane hardware-agnostic and portable across multicloud environments.

  The functionality of SPA includes:

- The support for standard PFCP port for a single UP connection.
- The support for nonstandard ports for both PFCP and GTPv1-U for multiple connections.
- UP to CP keep alive (KA) to detect any communication channel faults between CP-UP.

- **NOS Adaptation Layer (NAL)**—translates the CP instructions or messages coming to the UP to Cisco IOS XR-defined format. It is the Cisco IOS XR implementation of cnBNG APIs that interfaces with Cisco IOS XR infra components for various functions. These functions include input-output of packets, interface creation and deletion, subscriber feature provisioning, route operations, subscriber interface statistics and notifications. NAL also manages the subscriber flow on the Cisco IOS XR platform and handles the high availability (HA) requirements of Cisco IOS XR infrastructure.

  The **cnbng-nal** is the internal process that provisions all the above NAL functionalities. For details on commands that are related to NAL, see the *Configuring Cloud Native BNG User Plane* chapter and the *Verifying Cloud Native BNG User Plane Configurations* chapter.

- **Platform-specific Layer**—is the API adaptation layer that helps to plug-in different types of hardware architectures to the common Cisco IOS XR infrastructure. This layer in turn helps to extend the user plane functionality to other Cisco IOS XR platforms without altering the basic infrastructure.

  Platform-specific layer defines the forward API calls that each underlying platform of the user plane has to implement. The system uses these APIs to provision the following BNG functions:

  - IPoE subscriber traffic classification—for L2-connected subscribers that are based on port, MAC, and VLAN.
  - PPPoE subscriber traffic classification—for L2-connected subscribers that are based on port and PPPoE session-ID.

cnBNG User Plane and Control Plane Reachability

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
cnBNG User Plane and Control Plane Reachability over IPv6 Networks	Release 24.3.1	You can now establish communication between the user and control plane cnBNG over IPv6 networks, allowing users to seamlessly migrate to an IPv6-only core.

Starting from Release 24.3.1, we support IPv6 transport for the interface between cnBNG and user plane.

For network functionality and interoperability, it is essential to ensure reachability between user plane and control plane using both IPv4 and IPv6. You can set up the user plane to reach control plane either using IPv4 or IPv6 transport, and not both. However, control plane supports multiple user planes configured with IPv4 or IPv6 transport simultaneously.

The protocols used for control plane and user plane communication for exchanging subscriber provisioning messages and protocol packets are:

- PFCP
- GTP-U

Software and Hardware Requirements

The support for cnBNG user plane functionality on Cisco ASR 9000 Series Routers is compatible with the following line card (LC), route switch processors (RSPs), and modular port adapters (MPAs).

Note The Cisco IOS XR Software Release 7.3.1 supports cnBNG UP only on Cisco ASR 9000 High Density 100GE Ethernet line cards. See the table for the list of supported PIDs.

Table 2: Software and Hardware Requirements for cnBNG User Plane

Cisco IOS XR Software Release	LC	RSP	Chassis	MPA
Cisco IOS XR Software Release 25.2.1	A99-4HG-FLEX-X-SE A9-4HG-FLEX-X-SE	Not Applicable	Not Applicable	Not Applicable
Cisco IOS XR Software Release 7.11.1	Not Applicable	Not Applicable	ASR-9902-FC (ASR 9902, 2RU, Flexible Consumption Chassis) ASR-9903-FC (ASR 9903 1.6T Flexible Consumption Chassis)	Not Applicable
Cisco IOS XR Software Release 7.3.1	• A9K-24X10GE-1G-SE • A9K-48X10GE-1G-SE • A9K-4X100GE-SE • A9K-MOD200-SE • A9K-MOD200-CM • A9K-MOD400-SE • A9K-MOD400-CM	• A9K-RSP880-SE • A9K-RSP880-LT-SE • A99-RP-SE and A99-RP2-SE (on the Cisco ASR 9912 and the Cisco ASR 9922 chassis) • A99-RSP-SE (on the Cisco ASR 9910 and the Cisco ASR 9906 chassis) • A9K-RSP5-SE	Not Applicable	• A9K-MPA-1X100GE • A9K-MPA-2X100GE • A9K-MPA-4X10GE • A9K-MPA-8X10GE • A9K-MPA-20X10GE • A9K-MPA-20X1GE • A9K-MPA-1X40GE • A9K-MPA-2X40GE

Access Types and Subscriber Types

Access Types

The cnBNG user plane on Cisco IOS XR platform supports sub-interface Bundle-Ethernet access type with these encapsulations:

- **Single Dot1q**—which is the IEEE 802.1Q networking standard to support VLANs on an Ethernet network.
- **Double-tagged VLANs**—where two VLAN ID tags (inner tag and outer tag) are inserted into a single data frame. This encapsulation enables users to use their own VLANs inside the VLAN provided by the service provider.
- **Ambiguous VLANs**—that use a range or group of VLAN IDs that enables you to create multiple sessions on a single access-interface.

Subscriber Types

The IP subscriber sessions that connect through a Layer-2 aggregation network are called **L2-connected** sessions. Subscriber sessions where an IPv4 address and an IPv6 address co-exist for the same subscriber are called **dual-stack** subscriber sessions.

The cnBNG UP on Cisco IOS XR platform supports two types of **L2-connected dual-stack** subscriber sessions:

- **IPoE-DHCP dual-stack sessions**: In IP over Ethernet (IPoE) sessions, subscribers run IPv4 or IPv6 on the CPE device and connect to the BNG through an L2 aggregation network. These sessions rely on the DHCP protocol for assigning IP address for the subscriber.
- **PPPoE-DHCPv6 dual-stack PTA sessions**: The PPP over Ethernet (PPPoE) subscriber session is established using the point-to-point protocol (PPP) that runs between the CPE and BNG. These sessions rely on the standard PPP negotiations for subscriber authentication and IP address assignment.

  In a PPP Termination and Aggregation (PTA) session, the PPP encapsulation terminates on BNG. After that the BNG routes the traffic to the service provider using IP routing.

Subscriber Features

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Increased Granularity for Cloud Native BNG Traffic Management with Hierarchical QoS (H-QOS)	Release 7.4.2	This feature allows you to specify QoS behavior at multiple policy levels for Internet Protocol over Ethernet (IPoE), Point-to-Point Protocol over Ethernet (PPPoE), PPP Termination and Aggregation (PTA), and LNS (L2TP Network Server) sessions and provides a high degree of granularity in traffic management. Use the first level of the traffic policy, the parent traffic policy, to control the traffic at the main interface or sub-interface level. Use the second level, the child traffic policy, for additional control over a specific traffic stream or class.
Multiple Framed IPv4 and IPv6 Routes for Customer PremisesEquipment (CPE)	Release 7.4.2	You can configure multiple framed routes for IPv4 and IPv6 traffic across CPE. This functionality allows you to route multiple customer networks through a single customer broadband connection, thus enabling the LAN network subscriber to use a different subnet from WAN.

Feature Name	Release Information	Feature Description
Cloud Native BNG feature extension to 5th Generation Line Card	Release 7.4.2	The A99-32X100GE-X-SE, A9K-20HG-FLEX-SE , and A9K-8HG-FLEX-SE line cards now support the following Cloud Native BNG functionalities: • IPoE subscriber sessions that run both IPv4, IPv6 on the CPE device. • PPP over Ethernet PPP Termination and Aggregation (PPPoE PTA) sessions • DHCPv6 support for PPPoE sessions • Cloud Native BNG over Bundle Ether interface • Quality of Service (QoS) • Policy-based Routing (PBR) • Access Control List (ACL) • ICMP unreachable • Lawful Intercept (LI) This enhancement enables BNG features to leverage the higher throughput of the 5th generation of line cards.

This section lists the set of subscriber features that cnBNG user plane on the Cisco IOS XR platform supports.

- **IPv4 or IPv6**:
  - **Maximum Transmission Unit (MTU)**—that defines the maximum size of each packet that you can transmit during the subscriber session.
  - **Unicast Reverse Path Forwarding (URPF)**—that ensures that the system does not accept any traffic on the subscriber interface from malformed or forged IP source addresses.
  - **Internet Control Message Protocol (ICMP)**—a supporting protocol that networking devices use to send error messages and operational information to the originator of transmission.
- **Access Control List (ACL)**—that performs packet filtering to control the traffic flow into and out of network interfaces. It helps to define the access rights such as, filtering the content, blocking access to various resources and so on, for a subscriber .

  Supported ACL types are:
  - Input ACL (IPv4 or IPv6)

- Output ACL (IPv4 or IPv6)

- **QoS**:
  - **Policing (input and output)**—that allows you to control the maximum rate of traffic sent or received on an interface. It also allows to partition a network into multiple priority levels or class of service (CoS).
  - **Shaping (output)**—that allows you to control the traffic flow that exits an interface to match its transmission to the speed of the remote target interface. It also ensures that the traffic conforms to policies contracted for it.
  - **Policy Merging**—that merges multiple QoS policies on a single subscriber. The UP supports a maximum of 6 policy-maps and 10 class-maps, including the default ones.

- **Hierarchical QoS (H-QoS)**—that allows you to specify QoS behavior at multiple policy levels, which provides a high degree of granularity in traffic management. cnBNG supports the following two-level hierarchical policy for deploying QoS:
  - Parent policy:
  - Child policy

  H-QoS is applied on the router interface using nested traffic policies. The first level of traffic policy, the parent traffic policy, is used for controlling the traffic at the main interface or sub-interface level. The second level of traffic policy, the child traffic policy, is used for additional control over a specific traffic stream or class. The child traffic policy, is a previously defined traffic policy, that is referenced within the parent traffic policy. Two-level H-QoS is supported on both ingress and egress directions on all line cards and on physical or bundle main interfaces and sub-interfaces.

  To know more about H-QOS, refer the *Configuring Hierarchical Modular QoS* chapter in the *Modular QoS Configuration Guide for Cisco ASR 9000 Series Routers*

- **HTTP Redirect using PBR** (for input policy)—that redirects subscriber traffic to a destination other than to its original destination. Policy-based Routing (PBR) makes packet forwarding decisions based on the policy configurations, instead of routing protocols.

- **Accounting:** cnBNG UP supports periodic accounting for these accounting types:
  - **Session Accounting**—which is the statistics of a subscriber session.
  - **Service Accounting:**—which is the statistics for each service (collection of features) that is enabled for a subscriber.

Note
- You cannot enable service accounting without session accounting.
- You cannot have different periodicity for session and service accounting.

- **Lawful Intercept** (for mediation device in default or non-default VRF)—that allows Law Enforcement Agencies (LEA) to conduct electronics surveillance as authorized by judicial or administrative order.

- **Multiple Framed Route**—allows a large number of customer networks to reach through framed routes through a single broadband connection. Framed Route is supported on both IPoE and PPPoE sessions.

There’s no limit enforced to the number of framed routes per session. You don't have to configure or enable Framed Route through command line interface as it is downloaded from RADIUS.

Read more about these features in the *Cisco ASR 9000 Series Aggregation Services Router Broadband Network Gateway Configuration Guide*.

Supported Parameter Limit for Subscriber Features

The cnBNG user plane on Cisco IOS XR platform supports a maximum of:

- 32 IP subnet pools
- 32 secondary IP addresses
- Eight QoS services
- Eight class-maps
- Six actions for multi-action change-of-authorization (MA-CoA)

  (MA-CoA is a feature which enables the service providers to activate and deactivate multiple subscriber services using a single CoA request).

High Availability

High Availability (HA) enables network-wide protection by providing fast recovery from faults that may occur in any part of the network. The cnBNG user plane does not delete the subscriber state, summary subnet route state, subscriber route state, and so on, in a stable system except in a few scenarios. These scenarios can be either explicit execution of CLI commands to clear the session, process restart of peer process, mark and sweep procedure (an internal clean-up process which detects and reclaims the memory that is used by unused objects) of *cnbng-nal* process, route processor fail over (RPFO), or deletion of parent interface.

This section describes the expected behavior if a high availability event such as a router reload or RPFO occurs on the cnBNG user plane:

- NAL restores the last stable (check-pointed) session state with best effort after the HA event.
- To ensure data and session integrity between NAL and peer processes, the system triggers a mark and sweep procedure during *cnbng-nal* process restart. During this process, the NAL might not be able to restore the sessions due to unforeseen issues from the feature or from the IOS XR infra components. In that case, the system deletes those sessions and sends a notification to the CP.
- The *cnbng-nal* process restart does not initiate automatic reconciliation procedure between the CP and the UP. The CP triggers this explicitly using a CLI configuration.
- The Cisco IOS XR platform has active and standby hardware level support (active RSP and standby RSP) on logical interface subscriber. The data sync between these nodes is not real time. The *cnbng-nal* process periodically syncs for various internal data on a best-effort basis. There can be a few cases where session data is out of sync between route processors which leads to session recreate failure after RFPO. These cases maybe for recently created sessions or inflight sessions. The system deletes those sessions at UP and sends a notification to the CP.
- The CP acts on these notification events to make sure that the subscriber state is in sync. If not, it leads to out of sync sessions between the UP and the CP in such scenarios.

- During process restart or RPFO, mark and sweep procedure might lead to subscriber session deletion on UP.
- The UP might not push the final statistics if a process restart or RPFO happens while a subscriber or service deletion is in progress. In that case, the CP considers the last collected statistics PCRF or back-end statistics as the final statistics.

Usage Guidelines

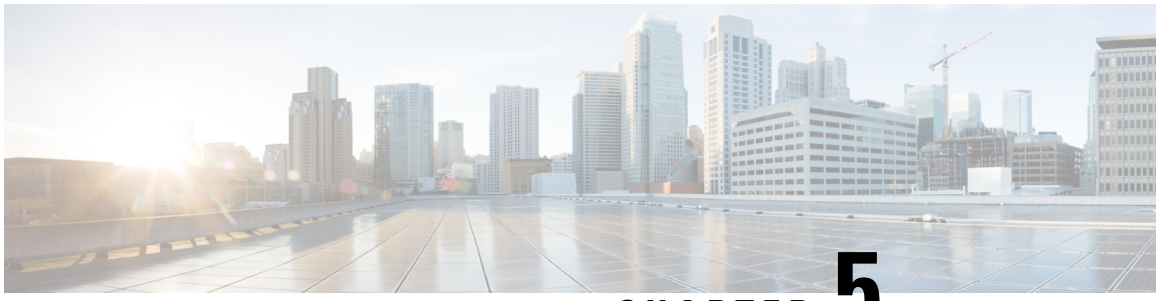
These guidelines apply to using the cnBNG user plane functionality on the Cisco IOS XR platform:

- You must not perform these actions on the fly while active subscriber sessions are present on the router:
  - Removal of configurations
  - Enabling or disabling service accounting
  - Deletion or modification of parent interface properties (such as IPv4 or IPv6 address, MTU, DHCPv4 initiator, PPPoE, DHCPv6 initiator, enable L2TP and LNS, and so on)
  - Deactivation of cnBNG package
  - Deletion or modification of VRF and loopback
  - Modification of service profile, and IPv4 or IPv6 address
- PPP keep alive (KA)—the user plane generates the PPP KA messages to the CPE to make packet transport more efficient between the CP and the UP. You must ensure that the duration of PPP keep alive is large enough (in tens of minutes) to have better CPU performance in scenarios with large subscriber scale.
- If an update request having service deactivation fails, the UP reactivates the service as part of rollback and starts the statistics afresh from zero.
- The CP-UP communication loss might cause the CP and UP to be in out of sync. There is no automated recovery mechanism for such scenarios.

Restrictions

The cnBNG user plane functionality on the Cisco IOS XR platform does not support these functionalities:

- Standalone PPP use case with cnBNG enabled
- Per pool or VRF tag support for IP pool subnet routes installed with a specific tag for the entire UP
- Back-to-back RPFO or switchover without graceful shutdown
- LC-based subscribers
- Subscriber redundancy group (SRG) for Bundle-Ethernet
- Enabling service accounting without session accounting
- Different periodicity for session and service accounting



CHAPTER 5

Installing Cloud Native BNG User Plane Packages

This chapter describes the procedure for installing cloud native BNG user plane packages on Cisco IOS XR platform.

- Installing and Activating the cnBNG Package on the User Plane, on page 25

Installing and Activating the cnBNG Package on the User Plane

Before you begin:

You must follow these guidelines before installing cnBNG package on the user plane:

- The cnBNG user plane functionality requires two packages to be installed on the router—the BNG support package (**asr9k-bng-supp-x64-*.rpm**) and the cnBNG package **(asr9k-cnbng-x64-*.rpm**).
- You can install cnBNG as an optional package on the router. The standard Cisco Golden ISO (GISO) image does not contain the cnBNG package.
- The physical BNG package (**asr9k-bng-x64*.rpm**) and the cnBNG package (**asr9k-cnbng-x64*.rpm**) are mutually exclusive. You cannot install both the packages on the router. The install operation fails if tried.
- You must uninstall and remove the physical BNG package and reboot the router prior to installing the cnBNG package on a router which is already being used as a physical BNG.
- You can either activate the BNG support package and the cnBNG package together as a single step or activate the BNG support package first and then activate the cnBNG package.
- The system does not support standalone PPP use case with cnBNG enabled. You must remove any PPP configuration before activating cnBNG on the router.

Installing and Activating the cnBNG Package on the User Plane

- **Step 1**: Install both the BNG support package and the cnBNG package from the RPM location to the router

  Use the **install add source** command.

```
Router#install add source tftp://209.165.200.225/test-path/
asr9k-bng-supp-x64-1.0.0.0-r73105I.x86_64.rpm asr9k-cnbng-x64-1.0.0.0-r73105I.x86_64.rpm
```

This step adds the BNG support package (**asr9k-bng-supp-x64-1.0.0.0-r73105I.x86_64.rpm**) and the cnBNG package (**asr9k-cnbng-x64-1.0.0.0-r73105I.x86_64.rpm**) from the source location of the RPMs (**tftp://209.165.200.225/test-path/**) to the router.

- **Step 2:** Activate the packages

  Use the **install activate** *activate-id* command.

  Where, *activate-id* is the ID that you see on the router console once the **install add** operation in the previous step is completed.

```
Router#install activate 1
```

  This step activates both the BNG support package and the cnBNG package which were installed as part of step 1.

- **Step 3:** Verify the activated packages.

  Use the **show install active** command.

```
Router#show install active
Sun Apr 19 09:49:34.041 UTC
Node 0/RSP0/CPU0 [RP]
  Boot Partition: xr_lv0
  Active Packages: 5
        asr9k-xr-7.3.1.05I version=7.3.1.05I [Boot image]
        asr9k-bng-supp-x64-1.0.0.0-r73105I
        asr9k-cnbng-x64-1.0.0.0-r73105I

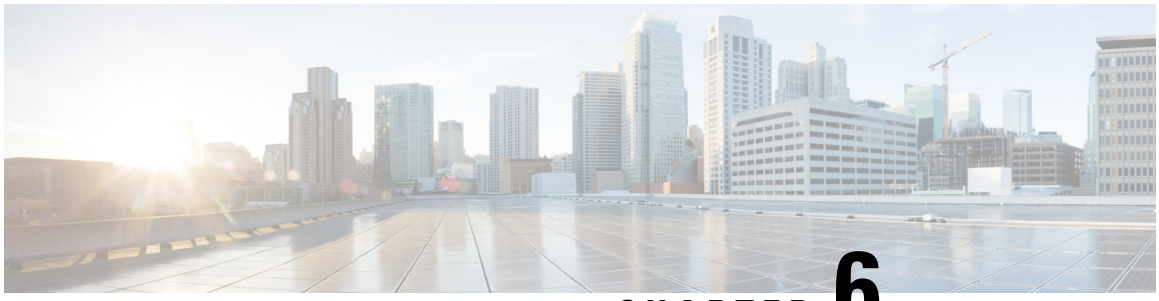
Node 0/0/CPU0 [LC]
  Boot Partition: xr_lv0
  Active Packages: 5
        asr9k-xr-7.3.1.05I version=7.3.1.05I [Boot image]
        asr9k-bng-supp-x64-1.0.0.0-r73105I
        asr9k-cnbng-x64-1.0.0.0-r73105I

Node 0/1/CPU0 [LC]
  Boot Partition: xr_lv0
  Active Packages: 5
        asr9k-xr-7.3.1.05I version=7.3.1.05I [Boot image]
        asr9k-bng-supp-x64-1.0.0.0-r73105I
        asr9k-cnbng-x64-1.0.0.0-r73105I

Node 0/3/CPU0 [LC]
  Boot Partition: xr_lv0
  Active Packages: 5
        asr9k-xr-7.3.1.05I version=7.3.1.05I [Boot image]
        asr9k-bng-supp-x64-1.0.0.0-r73105I
        asr9k-cnbng-x64-1.0.0.0-r73105I
```

  The *Active Packages* parameter in the show command output lists the BNG support package and the cnBNG package. This shows successful activation of the packages.

This step completes the installation and activation of cnBNG package on the user plane.



CHAPTER 6

Configuring Cloud Native BNG User Plane and Key Features

This chapter describes the configuration procedures to achieve the cnBNG user plane functionality on Cisco ASR 9000 Series Routers.

For details on cnBNG user plane commands, see the *Cloud Native BNG Command Reference for Cisco ASR 9000 Series Routers*.

- Configure cnBNG User Plane, on page 27
- Multiple Control Plane associations, on page 43
- Percentage-based QoS allocation for shared policy instances, on page 66
- Verify cnBNG User Plane Configuration, on page 70
- P-BNG and cnBNG coexistence, on page 84

Configure cnBNG User Plane

Before you begin:

You must follow these guidelines for configuring cnBNG user plane:

- You must perform a complete reimage followed by a reboot of the router if you are switching between physical BNG to cnBNG, or the other way around.
- Ensure that the cnBNG package is installed and activated on the user plane. See the *Installing Cloud Native BNG User Plane Packages* chapter for detailed procedure.
- The system does not support the removal of configurations while active sessions are present. You must delete all active sessions and dissociate the CP-UP connection prior to any configuration change or commit replace procedure.

Configuration Procedure

You must perform the following tasks for the UP to spawn the NAL process, to establish connection with the CP, and to provision the subscriber requests.

Configure Basic User Plane Settings

The basic user plane configuration for cnBNG involves these high-level tasks:

- Configuring the server endpoints of CP to which UP can send PFCP or GTP-U messages to enable cnBNG on the router.
- Configuring a loopback interface for each VRF.
- Configuring a route tag for subscriber summary routes.
- Configuring the access-interface to enable IPoE and PPPoE subscribers.

The cnBNG endpoint configurations on the UP are delivered to the cnBNG SPA component for initiating connection with the CP.

Configuration Procedure

This section describes the steps for the basic user plane configuration, which include certain mandatory and optional configurations.

Mandatory Configurations:

- Specifying a unique name for the UP-server instance.
- Specifying the details of the UP server (such as IP address, GTP port, and PFCP port) to which the CP can send PFCP or GTP-U messages.
- Specifying the details of CP server to which the UP can send PFCP or GTP-U messages.
- Specifying the retry count for CP-UP association.
- Enabling secondary address programming.
- Specifying a name for the auto-loopback VRF.
- Configuring a loopback interface to associate with the above VRF.
- Specifying a primary address for the loopback interface.

Optional Configuration:

- Configuring a route summary tag for the routes to add in the routing table

Configuration Example

Configuration example with IPv4 transport.

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)#hostidentifier asr9k-1
Router(config-cnbng-nal-local)#cp-server primary ipv4 198.51.100.1
Router(config-cnbng-nal-local)#up-server ipv4 192.0.2.1 gtp-port 15002 pfcp-port 15003 vrf
 default
Router(config-cnbng-nal-local)#secondary-address-update-enable
Router(config-cnbng-nal-local)#cp-association retry-count 10
Router(config—cnbng-nal-local)#auto-loopback vrf test
Router(config—cnbng-nal-local-auto-loopback-vrf)#interface Loopback2
```

```
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#primary-address 127.0.0.1
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#exit
Router(config—cnbng-nal-local-auto-loopback-vrf)#exit
/* Auto-loopback configuration for default VRF */
Router(config—cnbng-nal-local)#auto-loopback vrf default
Router(config—cnbng-nal-local-auto-loopback-vrf)#interface Loopback1
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#primary-address 10.0.0.1
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#exit
Router(config—cnbng-nal-local-auto-loopback-vrf)#exit
Router(config-cnbng-nal-local)#route-summary tag 4
Router(config-cnbng-nal-local)#commit
```

Configuration example with IPv6 transport.

Starting from Release 24.3.1, we support IPv6 transport for the interface between cnBNG and user plane.

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)#hostidentifier asr9k-1
Router(config-cnbng-nal-local)#cp-server primary ipv6 1::2
Router(config-cnbng-nal-local)#up-server ipv6 1:1
Router(config-cnbng-nal-local)#secondary-address-update-enable
Router(config-cnbng-nal-local)#cp-association retry-count 10
Router(config-cnbng-nal-local)#auto-loopback vrf test
Router(config—cnbng-nal-local-auto-loopback-vrf)#interface Loopback2
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#primary-address 1::3
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#exit
Router(config—cnbng-nal-local-auto-loopback-vrf)#exit

/* Auto-loopback configuration for default VRF */
Router(config—cnbng-nal-local)#auto-loopback vrf default
Router(config—cnbng-nal-local-auto-loopback-vrf)#interface Loopback1
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#primary-address 1::4
Router(config—cnbng-nal-local-auto-loopback-vrf-int)#exit
Router(config—cnbng-nal-local-auto-loopback-vrf)#exit
Router(config-cnbng-nal-local)#route-summary tag 4
Router(config-cnbng-nal-local)#commit
```

Running Configuration

```
Router#show running-config cnbng-nal location 0/RSP0/CPU0
cnbng-nal location 0/RSP0/CPU0
 hostidentifier asr9k-1
 up-server ipv4 192.0.2.1 vrf default
  gtp-port 15002
  pfcp-port 15003
 cp-server primary ipv4 198.51.100.1
 secondary-address-update-enable
cp-association retry-count 10
auto-loopback vrf test
 interface Loopback2
  primary-address 127.0.0.1
 !
!
auto-loopback vrf default
 interface Loopback1
  primary-address 10.0.0.1
 !
!
```

```
route-summary tag 4
!
```

Running configuration for IPv6 Transprot

```
Router#show running-config cnbng-nal location 0/RSP0/CPU0
cnbng-nal location 0/RSP0/CPU0
 hostidentifier asr9k-1
 up-server ipv6 1::1 vrf default
  gtp-port 15002
  pfcp-port 15003
 cp-server primary ipv4 1::2
 secondary-address-update-enable
cp-association retry-count 10
auto-loopback vrf test
 interface Loopback2
  primary-address 1::3
 !
!
auto-loopback vrf default
 interface Loopback1
  primary-address 1::4
 !
!
route-summary tag 4
!
```

Configure Access-Interface

This section describes how to configure the access-interface and to enable PPPoE on the cnBNG user plane.

Configuration Example

```
Router#configure
Router(config)#interface Bundle-Ether1.1
Router(config-subif)#ipv4 point-to-point
Router(config-subif)#ipv4 unnumbered Loopback1
Router(config-subif)#ipv6 enable
Router(config-subif)#encapsulation dot1q 1
Router(config-subif)#ipsubscriber
Router(config-cnbng-nal-ipsub)#ipv4 l2-connected
Router(config-cnbng-nal-ipsub-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-l2conn)#exit
Router(config-cnbng-nal-ipsub)#ipv6 l2-connected
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#exit
Router(config-cnbng-nal-ipsub)#exit

/* Enable PPPoE */
Router(config-subif)#pppoe enable
Router(config-subif)#commit
```

Running Configuration

```
Router#show running-config interface be1.1
interface Bundle-Ether1.1
 ipv4 point-to-point
 ipv4 unnumbered Loopback1
```

```
 ipv6 enable
 encapsulation dot1q 1
 ipsubscriber
  ipv4 l2-connected
   initiator dhcp
  !
  ipv6 l2-connected
   initiator dhcp
  !
 !
 pppoe enable
!
```

Configure Loopback Interface

This section describes how to configure the loopback interface for cnBNG user plane.



Note You must not configure any IP address under loopback interface.

Configuration Example

```
Router#configure
Router(config)#interface loopback 2
Router(config-if)#ipv6 enable
Router(config-if)#commit
```

Running Configuration

```
Router#show running-config interface loopback 2
interface Loopback2
 ipv6 enable
!
```

Enable Multiple Loopback Interfaces for Cloud Native BNG

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Multiple Loopback Interfaces for Cloud Native BNG	Release 7.6.2	By configuring multiple loopbacks under the same VRF and splitting a large network into smaller chunks, you can increase the number of IP addresses using IP subnet pools, thus improving address allocation and managing network bandwidth efficiently. In earlier releases, Cloud Native BNG supported only a single loopback for a given VRF.

You can configure multiple loopback interfaces under *cnbng-nal* configuration mode for default and non-default VRF.



Note A maximum of 47 loopback interfaces are supported under the *cnbng-nal* configuration mode.

Configuration for default VRF

```
Router#configure
Router(config)#cnbng-nal location 0/rsp0/CPU0
Router(config-cnbng-nal-local)#auto-loopback vrf default
Router(config-cnbng-nal-auto-loopback-vrf)#interface loopback 0
Router(config-cnbng-nal-auto-lb-vrf-int)#primary-address 10.0.0.1
Router(config-cnbng-nal-auto-lb-vrf-int)#exit
Router(config-cnbng-nal-auto-loopback-vrf)#interface loopback 1
Router(config-cnbng-nal-auto-lb-vrf-int)#primary-address 10.0.0.2
Router(config-cnbng-nal-auto-lb-vrf-int)#exit
Router(config-cnbng-nal-auto-loopback-vrf)#interface loopback 2
Router(config-cnbng-nal-auto-lb-vrf-int)#primary-address 10.0.0.3
Router(config-cnbng-nal-auto-lb-vrf-int)#
```

Running Configuration

```
/* Configuration for default VRF */
cnbng-nal location 0/rsp0/CPU0
 auto-loopback vrf default
  interface loopback 0
    primary-address 10.0.0.1
    !
  interface loopback 1
    primary-address 10.0.0.2
    !
  interface loopback 2
    primary-address 10.0.0.3
    !
```

Configuration for non-default VRF

```
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# hostidentifier cnBNG-SRG1
Router(config-cnbng-nal-local)# up-server ipv4 10.1.1.1 vrf default
Router(config-cnbng-nal-local)# cp-server primary ipv4 201.201.201.65
Router(config-cnbng-nal-local)# auto-loopback vrf vrftwo
Router(config-cnbng-nal-auto-loopback-vrf)# interface Loopback16
Router(config-cnbng-nal-auto-lb-vrf-int)# primary-address 10.0.0.1
Router(config-cnbng-nal-auto-lb-vrf-int)# interface Loopback17
Router(config-cnbng-nal-auto-lb-vrf-int)# primary-address 10.0.0.2
Router(config-cnbng-nal-auto-lb-vrf-int)# interface Loopback18
Router(config-cnbng-nal-auto-lb-vrf-int)# primary-address 10.0.0.3
Router(config-cnbng-nal-auto-lb-vrf-int)#
```

Running Configuration

```
cnbng-nal location 0/RSP0/CPU0
 hostidentifier cnBNG-SRG1
 up-server ipv4 19.1.1.1 vrf default
 cp-server primary ipv4 201.201.201.65
 auto-loopback vrf RJIL-VRF-OLT-MGMT
  interface Loopback16
   primary-address 1.1.0.1
  !
  interface Loopback17
   primary-address 2.1.0.1
  !
  interface Loopback18
   primary-address 3.1.0.1
```

Verification

```
/* Verification for default VRF */
Router#show cnbng-nal dynamic-routes summary

Location: 0/RSP0/CPU0
-----------------------

Counter Name            Value
--------------------------------
V4 OC Entries           64
V6 OC Entries           0
V4 Primary Entries      0
V4 Secondary Entries    64      <<< This the total secondary addresses pushed from CP for
each subnet
V4 RIB Entries          64      <<< This is the total number of v4 subnets pushed from CP
V6 RIB Entries          128     <<< This is the total number of v4 subnets pushed from CP
OC replay entry count   96
```

The following example shows the details of each route entry:

```
/* Verification for non-default VRF */
Router#show cnbng-nal dynamic-routes afi ipv4

Location: 0/RSP0/CPU0
-----------------------

Index                    : 1
Interface                : Loopback24
VRF                      : vrfone
AFI                      : IPv4
Prefix                   : 10.0.0.0/20
Secondary address        : 10.0.0.1
Route tag                : 41
```

```
State                    : RIB_REQ_COMPLETE
SRG group name           :
Route metric             : 0


[Event History]
| Event Name                              | Time Stamp

| Route OC request sent                   | Jun 15 15:29:10.144
| Added secon V4 addrs on lb              | Jun 15 15:29:10.144
| Route update succeed                    | Jun 15 15:29:10.144
| V4 route add success                    | Jun 15 15:29:10.144
==========================================

Router#show cnbng-nal dynamic-routes afi ipv6
Thu Jun 16 08:06:25.312 GMT

Location: 0/RSP0/CPU0
----------------------

Index                    : 1
Interface                : Loopback20
VRF                    : vrfone
AFI                      : IPv6
Prefix                   : 4001::/52
Secondary address        : NA
Route tag                : 31
State                    : RIB_REQ_COMPLETE
SRG group name           :
Route metric             : 0


[Event History]
| Event Name                              | Time Stamp

| Added secon V6 addrs on lb              | Jun 15 15:29:05.152
| Skip V6 rt install (standby)            | Jun 15 15:29:05.152
| Route update succeed                    | Jun 15 15:29:05.152
==========================================

/* Verification for non-default VRF */
Router#nshow cnbng-nal dynamic-routes afi ipv4
Mon Aug  8 05:13:59.576 GMT

Location: 0/RSP0/CPU0
----------------------

Index                    : 1
Interface                : Loopback16
VRF                      : vrftwo
AFI                      : IPv4
Prefix                   : 10.0.0.0/20
Secondary address        : 10.0.0.1
Route tag                : 32
State                    : RIB_REQ_COMPLETE
SRG group name           : group32
Route metric             : 0


[Event History]
| Event Name                              | Time Stamp

| Route OC request sent                   | Aug  4 15:39:24.288
| Added secon V4 addrs on lb              | Aug  4 15:39:24.288
```

```
| Route update succeed                    | Aug  4 15:39:24.416
| V4 route add success                    | Aug  4 15:39:24.416
==========================================

Index                   : 2
Interface               : Loopback17
VRF                     : vrftwo
AFI                     : IPv4
Prefix                  : 10.0.0.0/20
Secondary address       : 10.0.0.1
Route tag               : 32
State                   : RIB_REQ_COMPLETE
SRG group name          : group32
Route metric            : 0


[Event History]
| Event Name                              | Time Stamp

| Route OC request sent                   | Aug  4 15:39:26.976
| Added secon V4 addrs on lb              | Aug  4 15:39:26.976
| Route update succeed                    | Aug  4 15:39:26.976
| V4 route add success                    | Aug  4 15:39:26.976
==========================================
```

Configure DHCP

This section describes the steps to configure DHCP for cnBNG BNG user plane.

The basic DHCP configurations include these steps:

- Creating a cnBNG profile
- Assigning the cnBNG profile to access-interfaces

Configuration Example

```
Router(config)#dhcp ipv4
/* Create a cnBNG profile */
Router(config-dhcpv4)#profile cnbng_1 cnbng
Router(config-dhcpv4-cnbng-profile)#exit
/* Assign the cnBNG profile to access-interfaces */
Router(config-dhcpv4)#interface bundle-Ether 1.1 cnbng profile cnbng_1
Router(config-dhcpv4)#interface bundle-Ether 2.1 cnbng profile cnbng_1
Router(config-dhcpv4)#commit
```

Similarly, you can configure the DHCP IPv6 profiles.

Running Configuration

```
Router#show run dhcp ipv4
Wed Oct 14 16:48:56.814 UTC
dhcp ipv4
 profile cnbng_1 cnbng
 !
 interface Bundle-Ether1.1 cnbng profile cnbng_1
 interface Bundle-Ether2.1 cnbng profile cnbng_1
```

```
!


Router#show run dhcp ipv6
Wed Oct 14 16:49:19.095 UTC
dhcp ipv6
 profile cnbng_1 cnbng
 !
 interface Bundle-Ether1.1 cnbng profile cnbng_1
 interface Bundle-Ether2.1 cnbng profile cnbng_1
!
```

Configure Subscriber Gateway Address and Subnet Route

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
Disable Notifications for Dynamic Programming of Subscriber Gateway Address	Release 7.4.2	This feature allows you to disable the notifications exchanged internally between software components when the user plane (UP) of a cloud-native BNG (cnBNG) network programs the gateway address for its subscriber. It prevents excessive notifications when many active subscribers are on the UP, thus reducing the overhead on UP resources. The feature introduces the following command: • **disable-secondary-address-notification**

In cnBNG, the IP address management is more dynamic. Hence, the loopback interface for IPoE or PPPoE subscribers isn't provisioned in the user profile of the subscriber with static configuration. cnBNG user plane selects the loopback based on the subnet allocated to a loopback dynamically at cnBNG user plane.

Note For every VRF, one loopback must be present on the UP.

Consider this example,

```
On RSP0:
Tue Jul 28 05:55:13.015 UTC
cnbng-nal location 0/RSP0/CPU0
hostidentifier asr9k-1
up-server ipv4 192.0.2.1 vrf default
cp-server primary ipv4 198.51.100.1
auto-loopback vrf default
  interface Loopback1
```

```
   primary-address 10.0.0.1
  !
!
On RSP1:
Tue Jul 28 05:56:13.015 UTC
cnbng-nal location 0/RSP1/CPU0
hostidentifier asr9k-1
up-server ipv4 192.0.2.1 vrf default
cp-server primary ipv4 198.51.100.1
auto-loopback vrf default
  interface Loopback1
   primary-address 10.0.0.1
  !
!
```

In this example, the CP assigns 10.11.12.0/24 as subnet, and 10.11.12.1/32 as gateway address to subscribers under the default VRF. This gateway address serves as the DHCPv4 server address for DHCPv4 OFFER or ACK messages. The *cnbng-nal* process uses Operations Center (OC) to configure this gateway address as secondary IP address on the loopback and route provision APIs to program the entry in the L3 routing table.

Note The system supports a maximum of 32 secondary IP addresses under an interface.

```
Router#show ipv4 interface loopback 1
Tue Jul 28 05:29:58.741 UTC
Loopback1 is Up, ipv4 protocol is Up
  Vrf is default (vrfid 0x60000000)
  Internet address is 10.0.0.1/32
  Secondary address 10.11.12.1/32


Router#show route vrf all ipv4 subscriber
A    10.11.12.0/24 [1/0] via 0.0.0.0, 00:10:29
```

Note The dynamic programming of the subnet (secondary gateway) under the loopback causes a major churn on the UP if large scale of active subscribers is present on the node. Hence, the secondary address programming is disabled, by default.

Enable Secondary Address Programming

It's mandatory to enable the secondary address programming on cnBNG user plane. To enable that, use the **secondary-address-update enable** command under the cnbng-nal configuration mode.

Configuration Example

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal)#secondary-address-update enable
Router(config-cnbng-nal)#commit
```

Running Configuration

```
Router#show running-config cnbng-nal location 0/RSP0/CPU0
cnbng-nal location 0/RSP0/CPU0
 secondary-address-update enable
!
```

Note From Release 7.4.2 onwards, you can disable internal notifications on the UP while it programs the secondary address on the loopback interface by configuring the command **disable-secondary-address-notification**.

Disable Notifications for Dynamic Programming of Subscriber Gateway Address

In a cnBNG network, the CP assigns the gateway address for each subscriber. The UP dynamically programs gateway address assigned to each subscriber as a secondary IP address on its loopback interface. During this configuration, UP internally exchanges notification messages between various software components. The more the number of active subscribers on the UP, the more the notifications. To preserve valuable time and resources of the UP, you can disable notifications using the command **disable-secondary-address-notification** in the **cnbng-nal-local** config mode.

Configuration Example

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)#disable-secondary-address-notification
Router(config-cnbng-nal-local)#commit
```

Running Configuration

The following running configuration on cnBNG UP includes basic UP configuration as well:

```
Router#show running-config cnbng-nal location 0/RSP0/CPU0
cnbng-nal location 0/1/CPU0
 hostidentifier RTR1
 auto-loopback vrf test
  interface Loopback1
   primary-address 10.1.1.1
  !
 !
 auto-loopback vrf default
  interface Loopback0
   primary-address 10.30.30.1
  !

 !

 up-server ipv4 10.11.11.1 gtp-port 15002 pfcp-port 15003 vrf default
 cp-server primary ipv4 10.11.11.2
 enable-test-server
 disconnect-history file-logging-enable
 secondary-address-update enable
 disable-secondary-address-notification
 route-summary tag 111
```

```
 cp-association retry-count 5
!
```

Configure Route Summary

This section describes the steps to configure route summary for the cnBNG user plane.

The NAL handles the following routes:

- Individual subscriber routes
- Summary routes for subscriber pool subnet

The subscriber routes are part of the subscriber provisioning message, which includes:

- WAN IP address (/32 or /128 subnet)
- LAN IP (prefix delegation)

The summary routes are for the subscriber pool subnet which are exported to the core network to download traffic towards the subscriber. On physical BNG, the subscriber pool subnets were configured as static routes and redistributed through BGP or IGP. With cnBNG and auto-loopback selection, these subnets for the subscribers are added dynamically to the loopback. Every time a new subscriber pool subnet is added to the loopback, the same is added to the RIB with the tag that is provided by the CP. If tag is '0', the NAL uses the tag configured under the cnbng-nal. Routes with this tag can be exported to the core using the Routing Protocol for Low-Power and Lossy Networks (RPLs).

To configure route summary, use the **route-summary** command under the cnbng-nal configuration mode.

Configuration Example

```
Router#configure
Router(config)#cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal)#route-summary tag 10
Router(config-cnbng-nal)#commit
```

Running Configuration

```
Router#show running-config cnbng-nal location 0/RSP0/CPU0
cnbng-nal location 0/RSP0/CPU0
 route-summary tag 10
!
```

After the first subnet is installed on NAL, the following routes are added to the system:

```
A 10.11.12.0/24 [1/0] via 0.0.0.0, 0d01h
```

Export Routes to Core Network

This section describes how to export routes to core network as part of enabling cnBNG user plane functionality.

Configuration Example

```
Router#configure
Router(config)#route-policy test-policy-cnbng
Router(config-rpl)#if tag eq 10 then
Router(config-rpl-if)#set community (123:100)
Router(config-rpl-if)#done
Router(config-rpl-if)#endif
Router(config-rpl)#end-policy
Router(config)#commit

Router(config)#router ospf 10
Router(config-ospf)#vrf test-vrf-cnbng
Router(config-ospf-vrf)#redistribute subscriber route-policy test-policy-cnbng
Router(config-ospf-vrf)#commit
```

Running Configuration

```
Router#show running-config route-policy test-policy-cnbng
route-policy test-policy-cnbng

     if tag eq 10  then

       set community (123:100)

       done

     endif

   end-policy
   !

Router#show running-config router ospf
router ospf 10
 vrf test-vrf-cnbng
  redistribute subscriber route-policy test-policy-cnbng
  !
```

Configure ARP Scale Mode

This section describes the steps to configure ARP scale mode for the cloud-native BNG user plane.

To disable interface entry creation by ARP for each subscriber interface on the data plane (line cards), you must enable ARP scale mode for the subscriber using the **arp scale-mode-enable** command in subscriber configuration mode.

Configuration Example

```
Router#configure
Router(config)#subscriber
Router(config-subscriber)#arp scale-mode-enable
Router(config-subscriber)#commit
```

Running Configuration

```
Router#show running-config subscriber
Sat Aug 22 06:36:21.422 UTC
subscriber
arp scale-mode-enable
!
```

Configure Cloud Native BNG over Pseudowire Headend

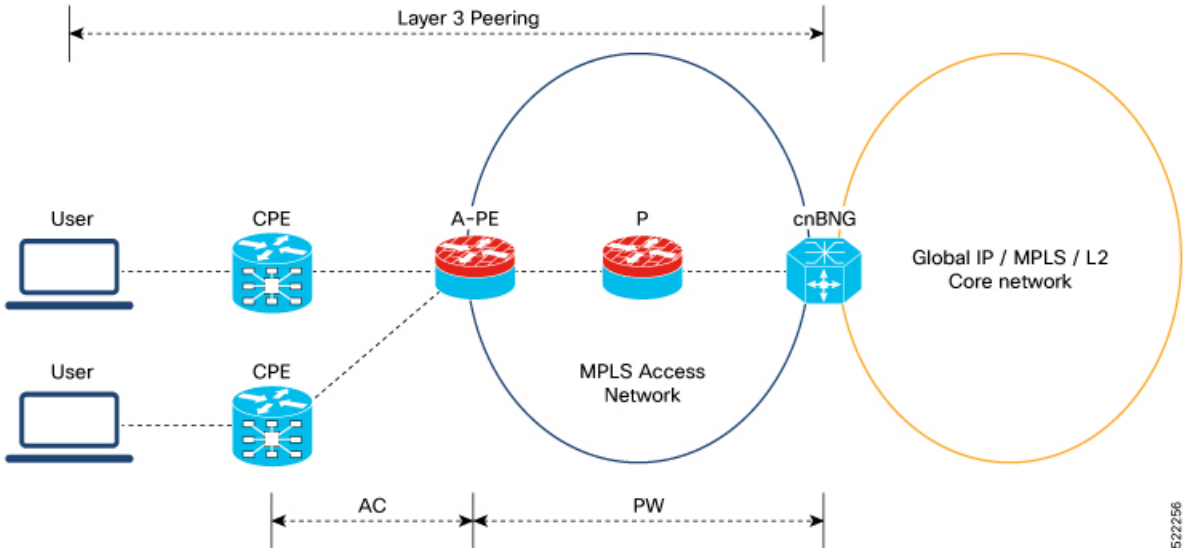
Table 6: Feature History Table

Feature Name	Release Information	Feature Description
Establishing Cloud Native BNG Sessions over Pseudowire Headend (PWHE)	Release 7.4.2	This feature establishes Cloud Native BNG subscriber sessions on PWHE interfaces. PWHE enables an easy and scalable mechanism for tunneling cnBNG traffic into a common IP, MPLS, or L2 network.

Cloud Native BNG provides subscriber support over Pseudowire Headend (PWHE). PWHE provides L3 connectivity to customer edge nodes through a pseudowire connection. PWHE terminates the L2VPN circuits that exist between the access-provide edge (A-PE) nodes, to a virtual interface, and performs routing on the native IP packet. Each virtual interface can use one or more physical interfaces towards the access cloud to reach customer Router through the A-PE nodes.

This figure shows a sample topology for Cloud Native BNG over Pseudowire Headend:

Figure 7: Sample Topology for Cloud Native BNG over Pseudowire Headend:



Restrictions

You can not configure eight ECMP links on the same PE device.

Configuration Example

This section provides the sample configurations for BNG over Pseudowire Headend:

The following is the sample configuration to allow IPOE or PPPOE subsciber to bring up from the PWHE access interface on the cnBNG:

```
Router#configure
Router(config)#interface PW-Ether100.102
Router(config-subif)#ipv4 unnumbered Loopback100
Router(config-subif)#ipv6 enable
Router(config-subif)#load-interval 30
Router(config-subif)#ipsubscriber
Router(config-cnbng-nal-ipsub)#ipv4 l2-connected
Router(config-cnbng-nal-ipsub-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-l2conn)#exit
Router(config-cnbng-nal-ipsub)#ipv6 l2-connected
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#exit
Router(config-cnbng-nal-ipsub)#exit
Router(config-subif)#pppoe enable
Router(config-subif)#encapsulation ambiguous dot1q any second-dot1q 102
Router(config-subif)#commit
```

This example shows the configuration of DHCPv4 on PWHE interfaces:

```
Router#configure
Router(config)#dhcp ipv4
Router(config-dhcpv4)#profile cn4 cnbng
Router(config-dhcpv4-cnbng-profile)#exit
Router(config-dhcpv4)#interface PW-Ether100.102 cnbng profile cn4
Router(config-dhcpv4)#interface PW-Ether100.103 cnbng profile cn4
Router(config-dhcpv4)#interface PW-Ether100.104 cnbng profile cn4
Router(config-dhcpv4)#commit
```

This example shows the configuration of DHCPv6 on PWHE interface:

```
Router#configure
Router(config)#dhcp ipv6
Router(config-dhcpv6)#profile cn4 cnbng
Router(config-dhcpv6-cnbng-profile)#exit
Router(config-dhcpv6)#interface PW-Ether100.102 cnbng profile cn6
Router(config-dhcpv6)#interface PW-Ether100.103 cnbng profile cn6
Router(config-dhcpv6)#commit
```

Running Configuration

The following example displays the running configuration of pw-ether interface.

```
Router#show running-config interface PW-Ether 100.102
Thu Feb  3 11:33:58.450 IST
interface PW-Ether100.102
 ipv4 unnumbered Loopback100
 ipv6 enable
 load-interval 30
 ipsubscriber
  ipv4 l2-connected
   initiator dhcp
  !
  ipv6 l2-connected
```

```
   initiator dhcp
  !
 !
 pppoe enable
 encapsulation ambiguous dot1q any second-dot1q 102
!
```

Configure DHCPv4 on PWHE interface:

```
Router#show run dhcp ipv4
Thu Feb  3 11:55:01.903 IST
dhcp ipv4
 profile cn4 cnbng
 !
 interface PW-Ether100.102 cnbng profile cn4
 interface PW-Ether100.103 cnbng profile cn4
 interface PW-Ether100.104 cnbng profile cn4
!
```

Configure DHCPv6 on PWHE interface:

```
Router#show run dhcp ipv6
Thu Feb  3 11:55:07.906 IST
dhcp ipv6
 profile cn6 cnbng
 !
 interface PW-Ether100.102 cnbng profile cn6
 interface PW-Ether100.103 cnbng profile cn6
!
```

Multiple Control Plane associations

Multiple Control Plane association is a deployment capability that

- allows a single User Plane (UP) to simultaneously connect to multiple Control Plane (CP) instances
- enables Subscriber Redundancy Groups (SRG) to be mapped flexibly across different CPs, and
- improves scalability, redundancy, and migration flexibility in Broadband Network Gateway (BNG) environments.

Table 7: Feature History

Feature Name	Release Information	Description
Multiple Control Plane associations	Release 26.2.1	You can now improve cnBNG redundancy and flexibility by enabling a single User Plane (UP) to connect to multiple Control Plane (CP) instances at the same time. This allows Subscriber Redundancy Groups (SRG) to be mapped more flexibly across different CPs.

Traditionally, a UP could connect to only one CP instance, which limited the ability to support complex SRG peering or to optimize resource utilization in large, distributed networks. With this enhancement, a UP can

maintain up to 10 logical associations with different CP instances or SRG instances. Each logical association acts as a separate UP context from the CP perspective.

Benefits of multiple CP associations

These are the key benefits:

- Enhanced network resilience through improved redundancy.
- Reduced operational complexity by simplifying migration processes.
- Support for scalable deployment across different environments.

Restrictions for multiple CP associations

These restrictions apply when configuring multiple CP associations with a single UP:

- Only IP over Ethernet (IPoE) hot standby SRG sessions are supported across multiple UP instances. PPPoE, LNS, and LAC session types are not supported in this mode.
- Warm standby is not supported when using logical UP instances.
- Support for non-SRG subscribers is available only on the default instance (ID 0).
- Simultaneous reconciliation of subscriber, route, and SRG information is not supported. For example, if subscriber reconciliation is in progress, route resynchronization and SRG group synchronization cannot start until subscriber reconciliation is complete.
- Changing a UP name on the CP while the association is up is not supported.
- Two UP instance IDs cannot share the same **cp-server** IP address. Configurations with duplicate **cp-server** IP addresses will be rejected.

Configure and verify multiple CP associations for a UP

To configure the Multiple CP associations feature, you must complete these tasks:

Configure multiple CP associations for a UP

Set up one or more logical associations between a UP and multiple CP instances to enable advanced redundancy and scalability.

Use this task when you want a UP node to support multiple CP clusters or SRG instances simultaneously, such as during migration or when optimizing SRG group mapping.

Before you begin

Ensure you have these prerequisites:

- Confirm the number and addresses of CP instances you wish to connect.

Follow these steps to configure multiple CP associations for a UP.

Procedure

Step 1 On the UP, configure the default (global) association parameters (instance-id 0).

Example:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# hostidentifier asr9k-1-cp-1-inst-0
Router(config-cnbng-nal-local)# up-server ipv4 10.1.1.51 vrf default gtp-port 15002 pfcp-port 15003
Router(config-cnbng-nal-local)# cp-server primary ipv4 192.0.2.1
Router(config-cnbng-nal-local)# cp-association retry-count 6
```

By default, the device uses instance ID 0 if no **up-instance-id** is specified.

Note
You can also configure **up-server** at the global level. If **up-server** is not configured locally, the individual UP instance inherits the global **up-server** setting.

Step 2 Configure additional logical UP instances to support multiple CP-UP association.

Example:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# up-instance-id 1
Router(config-cnbng-nal-local)# hostidentifier asr9k-1-cp-1-inst-1
Router(config-cnbng-nal-local)# up-server ipv4 10.1.1.52 vrf default gtp-port 15002 pfcp-port 15003
Router(config-cnbng-nal-local)# cp-server primary ipv4 192.0.2.2
Router(config-cnbng-nal-local)# cp-association retry-count 6
```

Repeat this step for each additional logical UP instance. Assign a unique **up-instance-id** (1–15) and **hostidentifier** per instance.

Example:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# up-instance-id 2
Router(config-cnbng-nal-local)# hostidentifier asr9k-1-cp-1-inst-2
Router(config-cnbng-nal-local)# up-server ipv4 10.1.1.53 vrf default gtp-port 15002 pfcp-port 15003
Router(config-cnbng-nal-local)# cp-server primary ipv4 192.0.2.3
Router(config-cnbng-nal-local)# cp-association retry-count 6
```

Step 3 On the UP, configure SRG groups to be mapped to UP instances.

Example:

Associating a group to default UP instance:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# subscriber-redundancy
Router(config-cnbng-nal-sub-red)# group group2001
Router(config-cnbng-nal-srg-grp)# virtual-mac 0000.1111.03e9
Router(config-cnbng-nal-srg-grp)# access-tracking track_group2001
Router(config-cnbng-nal-srg-grp)# access-interface-list
Router(config-cnbng-nal-srg-grp-intf)# interface Bundle-Ether40.2001
```

Example:

Associating a group to UP instance 1:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# up-instance-id 1
Router(config-cnbng-nal-local)# subscriber-redundancy
Router(config-cnbng-nal-sub-red)# group group2002
Router(config-cnbng-nal-srg-grp)# virtual-mac 0000.1111.03ea
Router(config-cnbng-nal-srg-grp)# access-tracking track_group2002
Router(config-cnbng-nal-srg-grp)# access-interface-list
Router(config-cnbng-nal-srg-grp-intf)# interface Bundle-Ether40.2002
```

Step 4 On the CP, configure each UP instance as a separate UP.

Important
The **user-plane** hostname configured on the CP must exactly match the **hostidentifier** configured for the corresponding UP instance on the UP.

Example:

```
user-plane
  instance 0
    user-plane asr9k-1-cp-1-inst-0
      peer-address ipv4 10.1.1.51
      subscriber-profile SUBS_PROF_1
      exit
      subscriber-redundancy
        group group2001
          preferred-role-active
          peer-identifier peer2001-x
          port-id-map port-name Bundle-Ether40.2001 1
        exit
      exit

  instance 1
    user-plane asr9k-1-cp-1-inst-1
      peer-address ipv4 10.1.1.52
      subscriber-profile SUBS_PROF_1
      exit
      subscriber-redundancy
        group group2002
          peer-identifier peer2002-x
          port-id-map port-name Bundle-Ether40.2002
        exit
      exit
```

Step 5 (Optional) On the UP, configure per-instance flow control thresholds.

Example:

```
Router(config-cnbng-nal-instance-id)# up-cp-notification flow-control 100
Router(config-cnbng-nal-instance-id)# up-cp-stats flow-control 150
```

Adjust per-instance flow and statistics thresholds as needed for scale and performance requirements.

Step 6 (Optional) On the UP, configure global flow control thresholds if required.

Use only if you require system-wide aggregate limits.

Example:

```
Router(config-cnbng-nal-local)# up-cp-notification global-flow-control 500
Router(config-cnbng-nal-local)# up-cp-stats global-flow-control 250
```

Step 7 On the UP, execute the **cnbng-nal cp-association start [up-instance-id 1] location 0/RSP0/CPU0** command to bring up the association.

Verify multiple CP associations to a UP

Ensure each logical UP instance is correctly associated and operating with its SRG CP peer. Diagnose and resolve issues related to SRG mapping, session and route allocation, and association failures.

Procedure

Step 1 Verify CP-UP association status for a specific or all UP instances.

Example:

```
Router# show cnbng-nal cp connection status up-instance-id 1

Tue Jun  2 08:25:19.093 IST
Location: 0/0/CPU0

UP Instance ID: 1

User-Plane configurations:
--------------------------
 IP   : 21.21.21.1
 GTP Port  : 15002
 PFCP Port  : 15003
 VRF   : default

Control-Plane configurations:
-----------------------------
 PRIMARY IP  : 21.21.21.2
 GTP Port  : 2152
 PFCP Port  : 8805

 Association retry count: 5

 Connection Status: Up
 Connection Status time stamp:  Mon Jun  1 19:23:18 2026

 Connection Prev Status: Init
 Connection Prev Status time stamp:  Mon Jun  1 19:20:48 2026

 Association status: Active
 Association status time stamp: Mon Jun  1 19:23:18 2026

 GTPu socket bind rc: Success
 PFCP socket bind rc: Success

[NAL-CP-INST-EVENTS]

| Event Name                    | Time Stamp              | User Data
| Association Active            | Jun  1 19:23:18.090052  | -- NA --
| Transport Up                  | Jun  1 19:23:18.091136  | -- NA --

Router# show cnbng-nal cp connection status up-instance-id all

Tue Jun  2 08:25:26.696 IST
Location: 0/0/CPU0
```

```
==========================================================================
UP Instance ID: 1

User-Plane configurations:
--------------------------
 IP   : 21.21.21.1
 GTP Port  : 15002
 PFCP Port  : 15003
 VRF   : default

Control-Plane configurations:
-----------------------------
 PRIMARY IP  : 21.21.21.2
 GTP Port  : 2152
 PFCP Port  : 8805

 Association retry count: 5

 Connection Status: Up
 Connection Status time stamp:  Mon Jun  1 19:23:18 2026

 Connection Prev Status: Init
 Connection Prev Status time stamp:  Mon Jun  1 19:20:48 2026

 Association status: Active
 Association status time stamp: Mon Jun  1 19:23:18 2026

 GTPu socket bind rc: Success
 PFCP socket bind rc: Success

[NAL-CP-INST-EVENTS]

| Event Name                    | Time Stamp              | User Data
| Association Active            | Jun  1 19:23:18.090052  | -- NA --
| Transport Up                  | Jun  1 19:23:18.091136  | -- NA --

UP Instance ID: 2

User-Plane configurations:
--------------------------
 IP   : 31.31.31.1
 GTP Port  : 15004
 PFCP Port  : 15005
 VRF   : default

Control-Plane configurations:
-----------------------------
 PRIMARY IP  : 31.31.31.2
 GTP Port  : 2152
 PFCP Port  : 8805

 Association retry count: 5

 Connection Status: Up
 Connection Status time stamp:  Mon Jun  1 19:23:18 2026

 Connection Prev Status: Init
 Connection Prev Status time stamp:  Mon Jun  1 19:20:48 2026

 Association status: Active
 Association status time stamp: Mon Jun  1 19:23:18 2026

 GTPu socket bind rc: Success
 PFCP socket bind rc: Success
```

```
[NAL-CP-INST-EVENTS]

| Event Name                    | Time Stamp              | User Data
| Association Active            | Jun  1 19:23:18.092553  | -- NA --
| Transport Up                  | Jun  1 19:23:18.093636  | -- NA --
```

Expected result: Each UP instance reports an **ESTABLISHED** state with its assigned CP group. If state is **FAILED** or **ATTEMPTING**, proceed to troubleshooting.

Step 2 Verify SRG group assignments and operational roles.

Example:

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

Tue Jun  2 08:25:37.200 IST
====================
Location: 0/0/CPU0
====================

Group-name   SRG role  Access OT  Core OT  Subs Count V4 routes V6 routes UP Inst ID
-------------------------------------------------------------------------------

Router# show cnbng-nal srg-group up-instance-id 1

Tue Jun  2 08:25:41.905 IST
====================
Location: 0/0/CPU0
====================

Group-name  SRG role  Access OT  Core OT  Subs Count V4 routes V6 routes UP Inst ID
-------------------------------------------------------------------------------
group1       Active    Up         Up       1          2         2         1

Total Entries : 1

Summary
-------
Category                  Total      Active     Standby    None
-------------------------------------------------------------------
Groups                    1          1          0          0
Subscribers               1          1          0          0
V4 subnet routes          2          2          0          0
V6 subnet routes          2          2          0          0

Router# show cnbng-nal srg-group detail up-instance-id 1

Tue Jun  2 08:25:55.846 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                           : 0x00000fa1
UP Instance ID                   : 1
SRG VRF name                     : **srg_4001
SRG mode                         : HOT
Last SRG role update time        : Jun 01 19:28:26.046033
Virtual mac                      : 0AAA.0BBB.0C01
V4 Table Id                      : 0xe0000012
```

```
V6 Table Id                       : 0xe0800012
V4 Proto Id                       : 0x00000002
V6 Proto Id                       : 0x00000001
Subscriber count                  : 1
Warm Subscriber count             : 0
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            : 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      : FALSE

Access tracking object
----------------------
    Object name                   : track1
    Tracking state                : Up
    Last tracking state update time : Jun 01 19:26:49.537142

Core tracking object
--------------------
    Object name                   : core1
    Tracking state                : Up
    Last tracking state update time : Jun 01 19:26:49.537217

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

IM counters
------------
    Total errors                  : 0

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

Terminate counters
------------------
    IM terminate                  : 0
    RIB terminate                 : 0
    Stats terminate               : 0

Flags
-----
Value: [0x0000007C]

 [VRF_CREATE | V4_RIB_CREATE | V6_RIB_CREATE |
  V4_RIB_COMPLETE | V6_RIB_COMPLETE]

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)   : 1.060 secs
    Last stats interaction time(S->A)   : 0.060 secs
    Max stats interaction time(A->S)    : 1.060 secs (Jun 01 19:25:40.854170)
    Max stats interaction time(S->A)    : 0.060 secs (NA)

Group role switchover RIB Info
------------------------------
    Last rib interaction time(A->S)     : 0.06103 secs
    Last rib interaction time(S->A)     : 0.060 secs
    Max rib interaction time(A->S)      : 0.06103 secs (Jun 01 19:25:40.853550)
    Max rib interaction time(S->A)      : 0.060 secs (NA)

    Active role change count            : 9
    Standby role change count           : 8

[SRG-GROUP-EVENTS]

| Event Name                              | Time Stamp              | User Data
| Group create                            | Jun  1 19:22:58.093098  | -- NA --
| V4 backup vrf create                    | Jun  1 19:22:58.093283  | -- NA --
| V6 backup vrf create                    | Jun  1 19:22:58.093311  | -- NA --
| V4 backup vrf create rsp                | Jun  1 19:22:58.109535  | -- NA --
| V4 backup rib tbl create                | Jun  1 19:22:58.109582  | -- NA --
| V6 backup vrf create rsp                | Jun  1 19:22:58.109585  | -- NA --
| V6 backup rib tbl create                | Jun  1 19:22:58.109593  | -- NA --
| V4 backup rib tbl create rsp            | Jun  1 19:22:58.112594  | -- NA --
| V6 backup rib tbl create rsp            | Jun  1 19:22:58.211073  | -- NA --
| State update held, transport down       | Jun  1 19:24:07.182998  | -- NA --
| Role none                               | Jun  1 19:24:15.954512  | -- NA --
| CP action add                           | Jun  1 19:24:29.121120  | -- NA --
| Subs rib route deletion started         | Jun  1 19:25:40.853550  | -- NA --
| Subs rib route deletion completed       | Jun  1 19:25:40.957261  | -- NA --
| Fast swo timer expiry, Role Active      | Jun  1 19:25:56.235268  | -- NA --
| Notify: State Down                      | Jun  1 19:26:39.621205  |
in_prog:F,comp:T,role:stdby,state:down
| Role standby                            | Jun  1 19:28:04.552720  | -- NA --
| Role standby start                      | Jun  1 19:28:04.552762  | -- NA --
| Role standby end                        | Jun  1 19:28:04.553567  | -- NA --
| Role active                             | Jun  1 19:28:26.046033  | -- NA --
| Role active start                       | Jun  1 19:28:26.046107  | -- NA --
| Role active end                         | Jun  1 19:28:26.046878  | -- NA --
| Notify: State Up                        | Jun  1 19:28:31.046323  |
in_prog:F,comp:T,role:act,state:up
| State change ack'ed                     | Jun  1 19:28:31.046495  | -- NA --
| Sent srg role response to CP            | Jun  1 19:28:37.642414  |
in_prog:F,comp:T,role:act,state:up

==========================================================================

Total Entries : 1
```

```
Router# show cnbng-nal srg-group detail up-instance-id all

Tue Jun  2 08:25:58.913 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                            : 0x00000fa1
UP Instance ID                    : 1
SRG VRF name                      : **srg_4001
SRG mode                          : HOT
Last SRG role update time         : Jun 01 19:28:26.046033
Virtual mac                       : 0AAA.0BBB.0C01
V4 Table Id                       : 0xe0000012
V6 Table Id                       : 0xe0800012
V4 Proto Id                       : 0x00000002
V6 Proto Id                       : 0x00000001
Subscriber count                  : 1
Warm Subscriber count             : 0
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            : 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      : FALSE

Access tracking object
----------------------
    Object name                   : track1
    Tracking state                : Up
    Last tracking state update time : Jun 01 19:26:49.537142

Core tracking object
--------------------
    Object name                   : core1
    Tracking state                : Up
    Last tracking state update time : Jun 01 19:26:49.537217

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

IM counters
------------
    Total errors                  : 0

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

Terminate counters
```

```
------------------
    IM terminate                    : 0
    RIB terminate                   : 0
    Stats terminate                 : 0

Flags
-----
Value: [0x0000007C]

 [VRF_CREATE | V4_RIB_CREATE | V6_RIB_CREATE |
  V4_RIB_COMPLETE | V6_RIB_COMPLETE]

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)   : 1.060 secs
    Last stats interaction time(S->A)   : 0.060 secs
    Max stats interaction time(A->S)    : 1.060 secs (Jun 01 19:25:40.854170)
    Max stats interaction time(S->A)    : 0.060 secs (NA)

Group role switchover RIB Info
------------------------------
    Last rib interaction time(A->S)     : 0.06103 secs
    Last rib interaction time(S->A)     : 0.060 secs
    Max rib interaction time(A->S)      : 0.06103 secs (Jun 01 19:25:40.853550)
    Max rib interaction time(S->A)      : 0.060 secs (NA)

    Active role change count            : 9
    Standby role change count           : 8

[SRG-GROUP-EVENTS]

| Event Name                              | Time Stamp              | User Data
| Group create                            | Jun  1 19:22:58.093098  | -- NA --
| V4 backup vrf create                    | Jun  1 19:22:58.093283  | -- NA --
| V6 backup vrf create                    | Jun  1 19:22:58.093311  | -- NA --
| V4 backup vrf create rsp                | Jun  1 19:22:58.109535  | -- NA --
| V4 backup rib tbl create                | Jun  1 19:22:58.109582  | -- NA --
| V6 backup vrf create rsp                | Jun  1 19:22:58.109585  | -- NA --
| V6 backup rib tbl create                | Jun  1 19:22:58.109593  | -- NA --
| V4 backup rib tbl create rsp            | Jun  1 19:22:58.112594  | -- NA --
| V6 backup rib tbl create rsp            | Jun  1 19:22:58.211073  | -- NA --
| State update held, transport down       | Jun  1 19:24:07.182998  | -- NA --
| Role none                               | Jun  1 19:24:15.954512  | -- NA --
| CP action add                           | Jun  1 19:24:29.121120  | -- NA --
| Subs rib route deletion started         | Jun  1 19:25:40.853550  | -- NA --
| Subs rib route deletion completed       | Jun  1 19:25:40.957261  | -- NA --
```

```
| Fast swo timer expiry, Role Active       | Jun  1 19:25:56.235268  | -- NA --
| Notify: State Down                       | Jun  1 19:26:39.621205  |
in_prog:F,comp:T,role:stdby,state:down
| Role standby                             | Jun  1 19:28:04.552720  | -- NA --
| Role standby start                       | Jun  1 19:28:04.552762  | -- NA --
| Role standby end                         | Jun  1 19:28:04.553567  | -- NA --
| Role active                              | Jun  1 19:28:26.046033  | -- NA --
| Role active start                        | Jun  1 19:28:26.046107  | -- NA --
| Role active end                          | Jun  1 19:28:26.046878  | -- NA --
| Notify: State Up                         | Jun  1 19:28:31.046323  |
in_prog:F,comp:T,role:act,state:up
| State change ack'ed                      | Jun  1 19:28:31.046495  | -- NA --
| Sent srg role response to CP             | Jun  1 19:28:37.642414  |
in_prog:F,comp:T,role:act,state:up

======================================================================

SRG group name                    : group2
SRG group admin state             : UP_Configured
SRG group state                   : Up
SRG role                          : None
SRG ID                            : 0x00000fa2
UP Instance ID                    : 2
SRG VRF name                      : -NA- (fast-switchover disabled)
SRG mode                          : HOT
Last SRG role update time         : Jun 01 19:24:15.954725
Virtual mac                       : 0AAA.0BBB.0A02
V4 Table Id                       : 0x00000000
V6 Table Id                       : 0x00000000
V4 Proto Id                       : 0x0000ffff
V6 Proto Id                       : 0x0000ffff
Subscriber count                  : 0
Warm Subscriber count             : 0
IPV4 route count                  : 0
IPV6 route count                  : 0
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
Stats state                       : IDLE
Sub disconnect resp pend          : NA
Sub count for keep alive start    : 0
Group for routed subscribers      : FALSE

Access tracking object
----------------------
    Object name                     : track1
    Tracking state                  : Up
    Last tracking state update time : Jun 01 19:26:49.537187

Core tracking object
--------------------
    Object name                     : core1
    Tracking state                  : Up
    Last tracking state update time : Jun 01 19:26:49.537261

Access Interfaces
-----------------
    Bundle-Ether1.2

IM counters
```

```
------------
    Total errors                    : 0

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

Terminate counters
------------------
    IM terminate                    : 0
    RIB terminate                   : 0
    Stats terminate                 : 0

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

Checkpoint Flags
----------------
Value: [0x00000048]

 [STATE_RESP_PENDING | ROLE_CHG_IN_PROG]

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)   : 0.060 secs
    Last stats interaction time(S->A)   : 0.060 secs
    Max stats interaction time(A->S)    : 0.060 secs (NA)
    Max stats interaction time(S->A)    : 0.060 secs (NA)

Group role switchover RIB Info
------------------------------
    Last rib interaction time(A->S)     : 0.060 secs
    Last rib interaction time(S->A)     : 0.060 secs
    Max rib interaction time(A->S)      : 0.060 secs (NA)
    Max rib interaction time(S->A)      : 0.060 secs (NA)

    Active role change count            : 0
    Standby role change count           : 1

[SRG-GROUP-EVENTS]

| Event Name                              | Time Stamp              | User Data
| Group create                            | Jun  1 19:22:58.095173  | -- NA --
| V4 backup vrf create                    | Jun  1 19:22:58.095296  | -- NA --
| V6 backup vrf create                    | Jun  1 19:22:58.095324  | -- NA --
| Role standby start                      | Jun  1 19:24:07.754079  | -- NA --
| Role standby                            | Jun  1 19:24:07.754079  | -- NA --
```

```
| Role standby end                          | Jun  1 19:24:07.755454  | -- NA --
| CP action add                             | Jun  1 19:24:07.755455  | -- NA --
| Sent srg role response to CP              | Jun  1 19:24:07.755510  |
in_prog:T,comp:F,role:stdby,state:up
| State update held, transport down         | Jun  1 19:24:07.756603  | -- NA --
| Role none                                 | Jun  1 19:24:15.954725  | -- NA --

======================================================================

Total Entries : 2
```

Check for correct mapping to UP instances and proper Active/Standby roles.

Step 3 Verify dynamic route distribution across UP instances.

Example:

```
Router# show cnbng-nal dynamic-routes summary up-instance-id all

Tue Jun  2 08:26:15.612 IST

Location: 0/0/CPU0
-----------------------

Counter Name              Value
---------------------------------
V4 OC Entries             2
V6 OC Entries             2
V4 Primary Entries        0
V4 Secondary Entries      2
V4 RIB Entries            2
V6 RIB Entries            2
OC replay entry count     0
V6 OC replay entry count  0
V6 Enabled Entries        0


Router# show cnbng-nal dynamic-routes summary up-instance-id 1

Tue Jun  2 08:26:23.611 IST
Location: 0/0/CPU0
-----------------------
Counter Name              Value
---------------------------------
V4 OC Entries             2
V6 OC Entries             2
V4 Primary Entries        0
V4 Secondary Entries      2
V4 RIB Entries            2
V6 RIB Entries            2
OC replay entry count     0
V6 OC replay entry count  0
V6 Enabled Entries        0
```

Routes must be correctly distributed per instance. Look for overlaps or missing entries.

Step 4 Verify subscriber session distribution and status.

Example:

```
Router# show cnbng-nal subscriber all up-instance-id all

Tue Jun  2 08:26:39.880 IST

=====================
```

```
Location: 0/0/CPU0
======================
Codes: CN - Connecting, CD - Connected, AC - Activated,
       ID - Idle, DN - Disconnecting, IN - Initializing
       UN - Unknown


CPID(hex)  Interface               State  Mac Address     Subscriber IP Addr / Prefix (Vrf) Ifhandle
-------------------------------------------------------------------------------------------------
d2         BE1.1.ip3221227008      AC     001b.0000.00d2  100.0.0.210 (default) 0x2b0
                                                          200.0.210.0/24 (Framed IPv4)
                                                          201.0.210.0/24 (Framed IPv4)
                                                          202.0.210.0/24 (Framed IPv4)
                                                          203.0.210.0/24 (Framed IPv4)
                                                          cafe:feed::d2 (IANA)
                                                          cafe:bead:0:d2::/128 (IAPD)
                                                          1:cee:bad:c0de:0:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:100:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:200:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:300:d2::/96 (Framed IPv6)
Session-count: 1
```

```
Router# show cnbng-nal subscriber all up-instance-id 1

Tue Jun  2 08:26:42.928 IST
======================
Location: 0/0/CPU0
======================
Codes: CN - Connecting, CD - Connected, AC - Activated,
       ID - Idle, DN - Disconnecting, IN - Initializing
       UN - Unknown


CPID(hex)  Interface               State  Mac Address     Subscriber IP Addr / Prefix (Vrf) Ifhandle
-------------------------------------------------------------------------------------------------
d2         BE1.1.ip3221227008      AC     001b.0000.00d2  100.0.0.210 (default) 0x2b0
                                                          200.0.210.0/24 (Framed IPv4)
                                                          201.0.210.0/24 (Framed IPv4)
                                                          202.0.210.0/24 (Framed IPv4)
                                                          203.0.210.0/24 (Framed IPv4)
                                                          cafe:feed::d2 (IANA)
                                                          cafe:bead:0:d2::/128 (IAPD)
                                                          1:cee:bad:c0de:0:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:100:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:200:d2::/96 (Framed IPv6)
                                                          1:cee:bad:c0de:300:d2::/96 (Framed IPv6)
Session-count: 1
```

```
Router# show cnbng-nal subscriber all summary up-instance-id all

Tue Jun  2 08:26:55.300 IST
======================
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  1  0  0  0
```

```
                idle  0  0  0  0  0
       disconnecting  0  0  0  0  0
              Total:  0  1  0  0  0


Session Counts by Address-Family:
                none  0  0  0  0  0
                ipv4  0  0  0  0  0
                ipv6  0  0  0  0  0
                dual  0  1  0  0  0
                 lac  0  0  0  0  0
              Total:  0  1  0  0  0

Router# show cnbng-nal subscriber all summary up-instance-id 1

Tue Jun  2 08:26:58.062 IST
=====================
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  1  0  0  0
                idle  0  0  0  0  0
       disconnecting  0  0  0  0  0
              Total:  0  1  0  0  0


Session Counts by Address-Family:
                none  0  0  0  0  0
                ipv4  0  0  0  0  0
                ipv6  0  0  0  0  0
                dual  0  1  0  0  0
                 lac  0  0  0  0  0
              Total:  0  1  0  0  0

Router# show cnbng-nal subscriber all detail up-instance-id 1

Tue Jun  2 08:27:21.590 IST
=====================
Location: 0/0/CPU0
=====================
Interface:                 Bundle-Ether1.1.ip3221227008
UPID:                      0xc0000600
CPID:                      0x000000d2
UP Instance ID:            1
Type:                      IPoE
  Sub-type:                L2 Connected
IPv4 Address:              100.0.0.210
IPv4 Framed Route:
  Prefix:                  200.0.210.0/24
  Next Hop:                3.1.1.1
  Tag:                     0
  Metric:                  0
  Distance:                0
  Next Hop VRF:
```

```
  Prefix:                 201.0.210.0/24
  Next Hop:               3.1.1.1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
  Prefix:                 202.0.210.0/24
  Next Hop:               3.1.1.1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
  Prefix:                 203.0.210.0/24
  Next Hop:               3.1.1.1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
IPv6 IANA Address:        cafe:feed::d2
IPv6 IAPD Prefix:         cafe:bead:0:d2::/128
IPv6 Slaac Prefix:        ::/0
CPE link local Address:   ::
IPv6 Framed Route:
  Prefix:                 1:cee:bad:c0de:0:d2::/96
  Next Hop:               192::1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
  Prefix:                 1:cee:bad:c0de:100:d2::/96
  Next Hop:               192::1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
  Prefix:                 1:cee:bad:c0de:200:d2::/96
  Next Hop:               192::1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
  Prefix:                 1:cee:bad:c0de:300:d2::/96
  Next Hop:               192::1
  Tag:                    0
  Metric:                 0
  Distance:               0
  Next Hop VRF:
IPv4 State:               UP, Mon Jun  1 19:28:40 2026
IPv6 State:               UP, Mon Jun  1 19:28:40 2026
Mac Address:              001b.0000.00d2
Inner VLAN ID:            Not Set
Outer VLAN ID:            1
Outer VLAN Cos:           0
Outer VLAN DEI:           0
Created:                  Mon Jun  1 19:28:40 2026
State:                    Activated
Ifhandle:                 0x000002b0
VRF:                      default
Access-interface:         Bundle-Ether1.1
SRG role:                 Active
SRG group name:           group1
IPOE-AFI-Cap:             SUB_AFI_IPV4_IPV6
Loopback-name:            Loopback0
IPV4-GW-addr:             100.0.0.1
```

```
 Attribute List: 0xdade20
1:  ipv4-unnumbered len=  9  value= Loopback0
2:  ipv6-enable     len=  4  value= 1(1)
3:  sub-ipv4-gateway len= 12  value= 100.0.0.1/32

Session Accounting:       disabled
Interim Interval:          0 secs
Interim fail count: None
Last interim failed reason: NA
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0

Activated Services: NA
Feature IDs activated : NA

[NAL-SUBSCRIBER-EVENTS]

| Event Name                     | Time Stamp              | User Data
| Transaction start              | Jun  1 19:28:40.024840  | txn id: 0x33, src_port: 0x33
| Interface create req           | Jun  1 19:28:40.024913  | -- NA --
| Interface create resp          | Jun  1 19:28:40.087462  | rc: OK
| SVM create req                 | Jun  1 19:28:40.087533  | -- NA --
| Sub GW addr add                | Jun  1 19:28:40.087547  | addr: 100.0.0.1/32
| UP Install req                 | Jun  1 19:28:40.087598  | -- NA --
| UP Install resp                | Jun  1 19:28:40.213931  | rc: OK
| UP Assoc req                   | Jun  1 19:28:40.213958  | -- NA --
| UP Assoc resp                  | Jun  1 19:28:40.215440  | rc: OK
| Produce done req               | Jun  1 19:28:40.215456  | req id: 0x3
| Produce done - Cfg gen         | Jun  1 19:28:40.268242  | rc: OK, req id: 0x3
| Produce done resp              | Jun  1 19:28:40.327444  | rc: OK, req id: 0x3
| SVM create resp                | Jun  1 19:28:40.327446  | rc: OK
| IPv6 Caps Up                   | Jun  1 19:28:40.328977  | vrf table id: 0xe0800000
| IPv4 caps wait                 | Jun  1 19:28:40.328978  | -- NA --
| IPv4 Caps Up                   | Jun  1 19:28:40.329475  | vrf table id: 0xe0000000
| Init data req                  | Jun  1 19:28:40.329947  | -- NA --
| Init data resp                 | Jun  1 19:28:40.383178  | rc: OK
| IPv4 route req                 | Jun  1 19:28:40.383188  | addr: 100.0.0.210/32
| IPv4 framed route req          | Jun  1 19:28:40.383250  | addr: 203.0.210.0/24
| IPv6 route req                 | Jun  1 19:28:40.383262  | addr: cafe:feed::d2/128
| IPv6 PD route req              | Jun  1 19:28:40.383336  | addr: cafe:bead:0:d2::/128
| IPv6 framed route req          | Jun  1 19:28:40.383361  | addr: 1:cee:bad:c0de:300:d2::/96
| IPv4 route resp                | Jun  1 19:28:40.484099  | rc: OK
| IPv4 framed route resp         | Jun  1 19:28:40.484182  | rc: OK, addr: 203.0.210.0/24
| IPv6 route resp                | Jun  1 19:28:40.484195  | rc: OK
| IPv6 PD route resp             | Jun  1 19:28:40.484215  | rc: OK
| IPv6 framed route resp        | Jun  1 19:28:40.484241  | rc: OK, addr: 1:cee:bad:c0de:300:d2::/96
| IPv4 Up                        | Jun  1 19:28:40.484247  | -- NA --
| IPv6 ND session req            | Jun  1 19:28:40.484300  | -- NA --
| IPv6 ND session resp           | Jun  1 19:28:40.484617  | rc: OK
| IPv6 Up                        | Jun  1 19:28:40.484618  | -- NA --
| Session up notified            | Jun  1 19:28:40.484620  | -- NA --
| Transaction stop               | Jun  1 19:28:40.484869  | rc: OK, txn id: 0x33, src_port: 0x33
==================__#__======================


Session-count: 1
```

Sessions should align with expected UP instance and SRG group associations.

Step 5 Review PFCP message, reconciliation, and notification history.

Example:

```
Router# show cnbng-nal spa pfcp-api structure dump all up-instance-id 1

Tue Jun  2 08:27:51.777 IST
Entry               : 1 (NAL Response for session create request),
Timestamp           : Jun 01 19:28:40.484898,
UPID                : 0xc0000600
CPID                : 0xd2
Session Txn id      : 261
Return value        : SPA_RC_SUCCESS
Access Interface    : Bundle-Ether1.1
Router name         :
Error message       :

Up-Instance-ID      : 1
================================================================================
Entry               : 2 (SPA Create Request),
Timestamp           : Jun 01 19:28:40.024568,
UPID                : 0x00000000
CPID                : 0x000000d2
Msg Txn id          : 51
Source port id      : 51
Session Txn id      : 261
Oper code           : SPA_OPER_TYPE_CREATE
IPOE AFI Cap        : SPA_SUB_AFI_IPV4_IPV6

Feature mask        : 0x1ee00
SRG report id       : 0
Feature flags
-------------
[SPA_FEATURE_IPv4_ADDR | SPA_FEATURE_IANA_ADDR | SPA_FEATURE_IAPD_ADDR | SPA_FEATURE_V4_ROUTE |
 SPA_FEATURE_V6_ROUTE | SPA_FEATURE_SRG | SPA_FEATURE_OUTER_VLAN]

IPv4 address Info
-----------------
Oper code           : SPA_OPER_TYPE_CREATE
IPv4 address        : 100.0.0.210

IPv6 address Info
-----------------
Oper code           : SPA_OPER_TYPE_CREATE
IPv6 address(iana)  : cafe:feed::d2
IPv6 address(iapd)  : cafe:bead:0:d2::
Link local          : ::
Prefix length       : 128

Session Data
------------
Access interface    : Bundle-Ether1.1
Mac address         : 001b.0000.00d2
VRF name            :
Session type        : SPA_SESSION_TYPE_IPOE
SRG group name      : group1

    Vlan Info
    ---------
    Outer vlan          : 1

IPv4 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv4 address        : 200.0.210.0
Prefix length       : 24
```

```
Next hop            : 3.1.1.1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv4 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv4 address        : 201.0.210.0
Prefix length       : 24
Next hop            : 3.1.1.1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv4 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv4 address        : 202.0.210.0
Prefix length       : 24
Next hop            : 3.1.1.1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv4 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv4 address        : 203.0.210.0
Prefix length       : 24
Next hop            : 3.1.1.1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv6 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv6 address        : 1:cee:bad:c0de:0:d2::
Prefix length       : 96
Next hop            : 192::1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv6 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv6 address        : 1:cee:bad:c0de:100:d2::
Prefix length       : 96
Next hop            : 192::1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
```

```
Distance            : 0

IPv6 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv6 address        : 1:cee:bad:c0de:200:d2::
Prefix length       : 96
Next hop            : 192::1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0

IPv6 route Info
---------------
Oper code           : SPA_OPER_TYPE_CREATE
VRF name            : default
IPv6 address        : 1:cee:bad:c0de:300:d2::
Prefix length       : 96
Next hop            : 192::1
Tag                 : 0
Metric              : 0
Next hop VRF name   :
Distance            : 0
....
```

```
Router# show cnbng-nal spa notif-history type keep-alive-failure up-instance-id 1

Tue Jun  2 08:28:19.816 IST
---- Total entries: 0 ----
```

```
Router# show cnbng-nal spa notif-history type subscriber-event up-instance-id 1

Tue Jun  2 08:28:29.198 IST
---- Total entries: 0 ----
```

Identify unusual message logs or recon reports that could indicate communication or failover issues.

Step 6 Check UP instance process information and flow control status.

Example:

```
Router# show cnbng-nal process-info detail

Tue Jun  2 08:28:38.242 IST
Location: 0/0/CPU0

 HA Pre_Init Role    : PRIMARY
 HA Role     : PRIMARY
 Restart-flag     : FALSE
 Node-Id     : 0
 Disc-Hist File-logging    : Enable
 Test-server config-enabled: FALSE

 Proc-flags : 87FFFF9F

    OT Connection Status: UP
    IM Connection Status: UP
    IPv4 RIB Connection Status: UP
    IPv6 RIB Connection Status: UP
    SUBDB Connection Status: UP
    Stats Connection Status: UP
    L2TP FIB Connection Status: DOWN
    L2TP Netio Connection Status: DOWN
    SPIO PPPoE Connection Status: UP
```

```
   SPIO PPP Connection Status: UP

SVM SUBDB PDone Req-ID     : 3
SVM Reconcilation Last UPID    : 0x0
SVM Reconcilation Pending-Req count      : 0
SVM Reconcilation Request Max    : 3072
RSI File-Descriptor     : 158
Bulk-Disconnect Start-time    : Tue Jun  2 08:28:38 2026
Clear-session Stats Timer-Running    : FALSE
Clear-feature Stats Timer-Running    : FALSE
Statistics Flags     : 0x00000005
SPA message Local seq-no    : 0
SPA message Global seq-no    : 0
Stats Hold time-out     : 200 secs
Stats Hold Request count    : 0
Stats Hold Request max     : 1000
Stats Stop Request count    : 0
Stats Interim Request count    : 0
Stats Clear Session Request count   : 0
Stats Clear Feature Request count   : 0
Interim Inflight Quota     : 4000
Interim Inflight LWM     : 4000
Pending Interim Queue Depth    : 0
Pending Interim Enqueue count    : 0
Pending Interim Dequeue count    : 0
Pending Interim Drop count    : 0
Pending Interim Dup-Skip count    : 0
SRG Stats Request Q-size    : 0
SRG Stats Retry Q-size     : 0
Sub-context Replay-In-Progess    : 0
Sub-context Cache Q-size    : 1
Sub-context Stale Q-size    : 0

Session stats:
===============
Create in progress: 0      [Peak: 1      Timestamp: Jun 01 19:25:10.448945]
Delete in progress: 0      [Peak: 1      Timestamp: Jun 01 19:25:14.858014]
Update in progress: 0      [Peak: 0      Timestamp: -NA-]
Total in progress:  0      [Peak: 1      Timestamp: Jun 01 19:25:10.448945]

CP Recon data:
===============
Duration                  : 300 secs
Replay reqs in progress   : 0
Replay subs in progress   : 0
Flags                     : 0x6

Notification flow control data:
==============================
Sub disconnect notification count       : 0
LCP timeout notification count          : 0
Sub disconnect notification drop count  : 0
LCP timeout notification drop count     : 0
Timer expiry count                      : 0
Timer running                           : No
LCP outstanding ka notification count   : 0
Notification global rate limit          : 1200

Stats flow control data:
========================
Interim stats message count      : 0
SRG stats message count          : 0
Stats oustanding message count   : 0
Stats global rate limit          : 1250
```

```
SRG Stats flow control data
===========================
Max outstanding requests  : 600
In-flight requests        : 0
Peak requests             : 0
Peak timestamp            : -NA-

Inflight flow configurations
============================
Maximum creates in-flight      : 600
Maximum updates in-flight      : 1000
Maximum deletes in-flight      : 3000
Maximum warm create in-flight  : 100
Delete batch timer             : 400
Delete batch size              : 800

Warm role change inflight data
==============================
Warm create in-progress        : 0

Features enabled on platform:
============================
IPoE L2 Terminated        : ENABLED
PPPoE-PTA                 : ENABLED
L2TP-LAC                  : ENABLED
L2TP-LNS                  : ENABLED

Standby group cached SPA request statistics
===========================================
Max cache requests        : 200
Inflight count            : 0
Create cache count        : 0
Update cache count        : 0

UP Instance specific statistics:
================================================

UP Instance ID                          : 1
Notification rate limit                 : 20
Stats rate limit                        : 20
Interim stats message count             : 0
SRG stats message count                 : 0
LCP outstanding ka notification count   : 0
Stats outstanding message count         : 0
Create outstanding message count        : 0
Update outstanding message count        : 0
Delete outstanding message count        : 0
Active session count                    : 1

UP Instance ID                          : 2
Notification rate limit                 : 20
Stats rate limit                        : 20
Interim stats message count             : 0
SRG stats message count                 : 0
LCP outstanding ka notification count   : 0
Stats outstanding message count         : 0
Create outstanding message count        : 0
Update outstanding message count        : 0
Delete outstanding message count        : 0
Active session count                    : 0

================================================
```

Investigate abnormal states or threshold limits that could signal resource or configuration issues.

Step 7 If any UP association is not in **ESTABLISHED** state, validate critical configuration parameters:

- Confirm `hostidentifier` values match on both UP and CP for each instance.
- Validate that `cp-server primary` address, `up-server IP`, and SRG mapping are unique and correct for each instance.
- Check for port or IP conflicts, missing exit commands, or duplicate values.

Expected result: The instance re-associates after correcting discrepancies.

Step 8 If association or SRG mapping fails repeatedly, verify CP connectivity and logs.

- Ping `cp-server` and `up-server` addresses to confirm reachability.
- Check system and debug logs for errors (for example, PFCP handshake failures).
- Ensure CP node process health and end-to-end connectivity.

Expected result: Connectivity and process issues are identified and resolved, restoring association.

Adding a new CP-UP association without disrupting existing connections

When you need to establish a new CP-UP association while an existing association is already active, this feature allows each CP GR instance to form its own connection to the User Plane.

If an association from CP GR instance 1 is already running, simply configure and bring up the new association for another CP GR instance. There is no need to change or interrupt the existing connection.

The new association can onboard additional SRG groups to provide greater redundancy or scalability. You can introduce the second association at any time so the active association remains unaffected.

Percentage-based QoS allocation for shared policy instances

A shared policy instance (SPI) is a policy-driven QoS mechanism that

- distributes a single pool of QoS resources across multiple groups of BNG subscriber sessions using percentage-based allocation
- allows these groups to share the allocated QoS resources collectively, and
- supports aggregate shaping or policing of all included subscriber sessions to a unified bandwidth or rate limit.

Table 8: Feature History Table

Feature Name	Release Information	Feature Description
Percentage-based QoS allocation for shared policy instances	Release 25.4.1	You can now streamline QoS policy management and ensure consistent rate enforcement across subscriber sessions with different port speeds by configuring shaping or policing rates as percentages in QoS policies for both ingress and egress directions. This enhancement replaces the previous requirement of using absolute values and allows a single policy map to be applied across multiple interfaces.

Overcoming challenges with percentage-based QoS allocation for SPI

Prior to Release 25.4.1, Shared Policy Instance (SPI) configurations were limited to absolute shaper or policer values. This limitation posed a significant challenge for customers with diverse access interface speeds, such as 1G, 10G, and 100G, where SPI parameters were tied to the parent subscriber session's fixed capacity.

With Release 25.4.1, this challenge is addressed by enabling the configuration of shaping or policing rates as percentages in QoS policies for both ingress and egress directions.

Benefits of percentage-based QoS allocation for SPI

Percentage-based QoS allocation for SPI offers several key benefits.

- Streamlined QoS policy management: Allows a single policy map across multiple interfaces and subscriber sessions, removing the need for separate configurations based on absolute values or interface speeds.
- Consistent rate enforcement: Ensures uniform rate application across subscriber sessions, adapting dynamically to different port speeds.
- Universal policy application: Allows a single policy map to be applied universally across multiple interfaces and subscriber sessions, eliminating the need for separate configurations based on absolute values or specific interface speeds.

Suppose a BNG manages 100 subscriber sessions and has a shared policy instance allocating 1 Gbps of bandwidth. If Group A is allocated 60% and Group B 40%, then Group A can use up to 600 Mbps collectively, and Group B up to 400 Mbps, regardless of the number of sessions in each group.

Restrictions for percentage-based QoS allocation for shared policy instances

- Configure percentage-based shaping and policing rates for QoS only if the subscriber has SPI enabled.
- The SPI feature is supported only for bundle subscribers.

How SPI use percentage-based QoS allocation

Summary

The key components involved in the process are:

- Network administrator: Defines QoS policies and shaping/policing rates.
- Network device: Receives and applies QoS policies, calculates bandwidth rates, enforces policies on hardware.
- Subscriber session: Receives the applied QoS policy and operates under the enforced bandwidth rules.

Network administrators configure QoS policies with percentage-based shaping and policing rates, which are then applied to SPI-enabled subscriber sessions. The network device dynamically calculates the actual absolute bandwidth by applying these percentages to the parent interface's capacity, and subsequently programs and enforces these determined values on the underlying hardware.

Workflow

These stages describe how shared policy instance works.

1. Policy configuration
   - The network administrator configures QoS policies.
   - The network administrator specifies shaping or policing rates as a percentage (for example, 10%, 50%) of bandwidth, rather than fixed absolute values.
2. Policy deployment
   - The network administrator applies the percentage-based QoS policy.
   - The policy is applied to subscriber sessions that are configured to use SPI.
3. Dynamic rate determination
   - The network device receives the activated policy on a subscriber session.
   - The network device automatically calculates the actual, absolute bandwidth rate.
   - This calculation applies the configured percentage to the current or configured bandwidth of the physical parent session to which the subscriber is connected.
4. Hardware implementation
   - The network device programs the dynamically determined absolute bandwidth values.
   - The network device enforces these values on the underlying network hardware.

Enable percentage-based QoS allocation for SPI

Procedure

Step 1 Define a feature-template on cnBNG CP to apply QoS policies using the shared policy instance mechanism. Associate them with specific SPI groups

This step defines the QoS policies and associates them with specific SPI groups (GRP1 and GRP2). Specify shaping or policing rates as percentages in these policies.

Example:

```
cnbng-cp(config)#profile feature-template default-qos
cnbng-cp(config)#qos
cnbng-cp(config)#in-policy policer-policy-in shared-policy-instance GRP1
cnbng-cp(config)#out-policy shaper-policy-out shared-policy-instance GRP2
cnbng-cp(config)#exit
cnbng-cp#exit
```

Step 2 Configure percentage-based rates in QoS policies on cnBNG UP.

Example:

```
Router-cnBNG-up#configure
Router-cnBNG-up(config-pmap)#policy-map policer-policy-in
Router-cnBNG-up(config-pmap)#class class-default
Router-cnBNG-up(config-pmap-c)#police rate percent 20
Router-cnBNG-up(config-pmap-c-police)#exit
Router-cnBNG-up(config-pmap)#end-policy-map

Router-cnBNG-up#configure
Router-cnBNG-up(config)#policy-map shaper-policy-out
Router-cnBNG-up(config-pmap)#class class-default
Router-cnBNG-up(config-pmap-c)#shape average percent 30
Router-cnBNG-up(config-pmap-c)#exit
Router-cnBNG-up(config-pmap)#end-policy-map
```

Step 3 Verify the configured rates for each subscriber to ensure the policies are applied successfully.

Example:

```
Router-cnBNG-up#show cnbng-nal subscriber all


Tue Jun 17 05:04:03.274 UTC

Location: 0/RSP0/CPU0
Codes: CN - Connecting, CD - Connected, AC - Activated,
       ID - Idle, DN - Disconnecting, IN - Initializing
       UN - Unknown


CPID(hex)  Interface               State  Mac Address     Subscriber IP Addr / Prefix (Vrf) Ifhandle
-------------------------------------------------------------------------------------------------
1          BE1.1.pppoe2147531664   AC     1234.1234.aabb  100.0.0.1 (default) 0x2f060
Session-count: 1



Router-cnBNG-up#show qos-ea interface BE1.1.pppoe2147531664 input member tenGigE 0/1/0/0
Interface: TenGigE0_1_0_0 input policy: policer-policy-in
Total number of classes:   1
Total number of UBRL classes:  0
Total number of CAC classes:   0
------------------------------------------------------
Policy name: policer-policy-in
Hierarchical depth 1
Interface type unknown
Interface rate 10000000 kbps
Port Shaper rate 0 kbps
Interface handle 0x0002F060
ul_ifh 0x060000C0, ul_id  0x00000000
uidb index 0xFFBB
```

```
qos_ifh 0x10002000ffbb
Local port 0, NP 0
Policy map id 0x1020, format 8, uidb index 0xFFBB
------------------------------------------------------
 Index 0 Level 0 Class name class-default service_id 0x0 Policy name policer-policy-in
 Node flags: LEAF DEFAULT DEFAULT-ALL
 Stats flags: policer-policy-in type 1 Max category 0
 Node Config:
 Police Color aware 0 Type 1 CIR/CBS/PIR/PBS: 2000000kbps/25000000B/0kbps/0B
 Node Result: Class-based stats:Stat ID 0x00C68DCC
 Queue: N/A Stat ID(Commit/Excess/Drop): 0x00000000/0x00000000/0x00000000
 Police ID (Token/Conform/Exceed/Violate): 0x00200001/0x00C68DCC/0x00C68DCD/0x00C68DCE
------------------------------------------------------


Router-cnBNG-up#show qos-ea interface BE1.1.pppoe2147531664 output member tenGigE 0/1/0/0
Tue Jun 17 05:04:35.338 UTC
Interface: TenGigE0_1_0_0 output policy: shaper-policy-out
Total number of classes:   1
Total number of UBRL classes:  0
Total number of CAC classes:   0
------------------------------------------------------
Policy name: shaper-policy-out
Hierarchical depth 1
Interface type unknown
Interface rate 10000000 kbps
Port Shaper rate 0 kbps
Interface handle 0x0002F060
ul_ifh 0x060000C0, ul_id  0x00000000
uidb index 0xFFBB
qos_ifh 0x10802000ffbb
Local port 0, NP 0
Policy map id 0x1420, format 8, uidb index 0xFFBB
------------------------------------------------------
 Index 0 Level 0 Class name class-default service_id 0x0 Policy name shaper-policy-out
 Node flags: LEAF Q_LEAF DEFAULT DEFAULT-ALL
 Stats flags: Queuing enabled
 Node Config:
 Shape: CIR/CBS/PIR/PBS: 0kbps/37500000B/3000000kbps/37500000B
 WFQ: BW/Sum of BW/Excess ratio: 0kbps/0kbps/1
 Queue limit 37500000 Guarantee 0
 Node Result: Class-based stats:Stat ID 0x00C68DCF
 Queue: Q-ID 0x0005e012 Stat ID(Commit/Excess/Drop): 0x00165222/0x00000000/0x009E0848
------------------------------------------------------
```

Verify cnBNG User Plane Configuration

This section describes the show commands to be executed on the router to verify cloud native BNG user plane configuration.

For details on cnBNG commands, see the *Cloud Native BNG Command Reference for Cisco ASR 9000 Series Routers*.

Verify cnBNG NAL Process Information

You can use the following commands to verify the NAL process information on cnBNG user plane.

- ```
  Router#show cnbng-nal process-info location 0/RSP0/CPU0
  Mon Aug  3 00:12:42.080 UTC

  Location: 0/RSP0/CPU0

   HA Pre_Init Role          : PRIMARY
   HA Role                   : PRIMARY
   Restart-flag              : FALSE
   card_type                 : 0
   Node-Id                   : 0
   Disc-Hist File-logging    : FALSE
   Test-server config-enabled: FALSE

   Proc-flags     : 8000FFBF

      OT Connection Status: UP
      IM Connection Status: UP
      IPv4 RIB Connection Status: UP
      IPv6 RIB Connection Status: UP
      SUBDB Connection Status: UP
  ```

- ```
  Router#show cnbng-nal process-readiness
  Mon Aug  3 00:12:00.778 UTC

  Location: 0/RSP1/CPU0

  NAL resync pending flags:
          Service Resync Pending
          Interface Resync Pending
          IPv4 Route Resync Pending
          IPv6 Route Resync Pending

  SIR status: not ready

  Location: 0/RSP0/CPU0
  NAL resync pending flags:
          NONE

  SIR status: ready
  ```

- ```
  Router#show processes cnbng_nal
  Fri Sep 11 09:22:45.139 UTC
                   Job Id: 456
                      PID: 1543

  Router#show processes memory 1543
  Fri Sep 11 09:24:12.398 UTC
  JID      Text(KB)   Data(KB)  Stack(KB) Dynamic(KB) Process
  ------ ---------- ---------- ---------- ----------- ------------------------------
  456           992    1700604        200       19999 cnbng_nal
  ```

Verify Control Plane Connection Status

You can use the following command to verify the connection status of cnBNG control plane.
  ```

```
Router#show cnbng-nal cp connection status
Fri Feb 19 11:27:31.178 UTC

Location: 0/RSP0/CPU0


User-Plane configurations:
--------------------------
IP             : 10.105.227.96
 GTP Port       : 2152
PFCP Port      : 8805
VRF            : default


Control-Plane configurations:
-----------------------------
PRIMARY IP     : 10.84.102.235
 GTP Port       : 2152
PFCP Port      : 8805

 Association retry count: 10

Connection Status: Up
Connection Status time stamp: Thu Feb 11 12:46:19 2021

Connection Prev Status : Down
Connection Prev Status time stamp: Thu Feb 11 12:44:55 2021

Association status: Active
Association status time stamp: Thu Feb 11 12:46:18 2021
```

Verify Subscriber Information

You can use the following commands to verify subscriber information on the cnBNG user plane.

```
Router#show cnbng-nal subscriber access-interface bundle-Ether 1.1
Mon Aug  3 00:04:42.558 UTC
======================
Location: 0/RSP0/CPU0
======================

               Type            PPPoE           IPoE
               ====            =====           ====

Session Counts by State:
        initializing           0               0
          connecting           0               0
           connected           0               0
           activated           0               8000
                idle           0               0
       disconnecting           0               0
               Total:          0               8000


Session Counts by Address-Family:
               none            0               0
               ipv4            0               0
               ipv6            0               8000
               dual            0               0
```

```
              Total:            0               8000


=====================
Location: 0/RSP1/CPU0
=====================

               Type            PPPoE           IPoE
               ====            =====           ====

Session Counts by State:
        initializing            0               0
          connecting            0               0
           connected            0               0
           activated            0               8000
                idle            0               0
       disconnecting            0               0
               Total:           0               8000


Session Counts by Address-Family:
               none             0               0
               ipv4             0               0
               ipv6             0               8000
               dual             0               0
              Total:            0               8000
```

-
```
Router#show cnbng-nal subscriber all
Fri Sep 11 06:07:52.343 UTC
   Codes: CN - Connecting, CD - Connected, AC - Activated,

         ID - Idle, DN - Disconnecting, IN - Initializing



CPID(hex)Interface            State Mac Address    Subscriber IP Addr / Prefix (Vrf)
Ifhandle
-------------------------------------------------------------------------------------------

1005ca0  BE2.500.ip2149474448 AC    0010.942e.3b00 13.0.92.160 (default) 0x225e60

                                                   1:4::5c9f (IANA)

                                                   2003:db0:0:5c9e::/64 (IAPD)

10053b2  BE2.500.ip2149466000 AC    0010.942e.3689 13.0.83.175 (default) 0xfdfe0

                                                   1:4::53b1 (IANA)

                                                   2003:db0:0:53b0::/64 (IAPD)

1004c81  BE2.600.ip2149013936 AC    0010.942e.5230 13.0.76.129 (default) 0x4079a0

                                                   1:4::4c80 (IANA)

                                                   2003:db0:0:4c7f::/64 (IAPD)

1004aaa  BE2.500.ip2149353232 AC    0010.942e.3205 13.0.74.169 (default) 0x5192e0

                                                   1:4::4aa9 (IANA)
```

```
                                                  2003:db0:0:4aa8::/64 (IAPD)
1004927  BE2.600.ip2149518576 AC    0010.942e.50b1 13.0.73.116 (default) 0x219ba0
                                                  1:4::4926 (IANA)
                                                  2003:db0:0:4925::/64 (IAPD)
10047e4  BE2.800.ip2149422928 AC    0010.9431.a7c7 13.0.71.228 (default) 0x41ff60
                                                  1:4::47e4 (IANA)
                                                  2003:db0:0:47e2::/64 (IAPD)
1004777  BE2.600.ip2149520224 AC    0010.942e.5021 13.0.71.115 (default) 0x41420
                                                  1:4::4776 (IANA)
                                                  2003:db0:0:4775::/64 (IAPD)
1003a6d  BE2.800.ip2149369728 AC    0010.9431.a3a1 13.0.58.105 (default) 0x141360
                                                  1:4::3a6d (IANA)
                                                  2003:db0:0:3a6a::/64 (IAPD)
10038b7  BE2.600.ip2149362240 AC    0010.942e.4bb2 13.0.56.178 (default) 0x259aa0
                                                  1:4::38b6 (IANA)
                                                  2003:db0:0:38b5::/64 (IAPD)
10028ba  BE2.500.ip2149210768 AC    0010.942e.2873 13.0.40.185 (default) 0x129620
                                                  1:4::28b9 (IANA)
                                                  2003:db0:0:28b8::/64 (IAPD)
100247b  BE2.600.ip2149396320 AC    0010.942e.46a3 13.0.36.113 (default) 0x4b8e0
                                                  1:4::2471 (IANA)
                                                  2003:db0:0:2470::/64 (IAPD)
100207a  BE2.500.ip2149356496 AC    0010.942e.2663 13.0.32.117 (default) 0x1a9460
                                                  1:4::2079 (IANA)
                                                  2003:db0:0:2078::/64 (IAPD)
1001d3f  BE2.600.ip2149251360 AC    0010.942e.44d4 13.0.29.61 (default) 0xcc760
```

-
```
Router#show cnbng-nal subscriber all summary
Sun Aug  2 16:26:44.281 UTC
=====================
Location: 0/RSP0/CPU0
=====================

               Type            PPPoE           IPoE
               ====            =====           ====
```

```
Session Counts by State:
        initializing            0               0
          connecting            0               0
           connected            0               0
           activated            0               130
                idle            0               0
       disconnecting            0               0
              Total:            0               130

Session Counts by Address-Family:
               none             0               0
               ipv4             0               130
               ipv6             0               0
               dual             0               0
              Total:            0               130

=====================

   Location: 0/RSP0/CPU0

   =====================


                  Type         PPPoE       IPoE

                  ====         =====       ====


   Session Counts by State:

          initializing         0           0

            connecting         0           0

             connected         226         0

             activated         31774       0

                  idle         0           0

         disconnecting         0           0

                Total:         32000       0


   Session Counts by Address-Family:

                  none         226         0

                  ipv4         7774        0

                  ipv6         0           0

                  dual         24000       0

                 Total:        32000       0
```

-
```
Router#show cnbng-nal subscriber all detail
Mon Aug  3 00:00:14.624 UTC
Location: 0/2/CPU0
====================
Location: 0/RSP1/CPU0
====================
Interface:                Bundle-Ether1.1.ip2148413040
UPID:                     0x800e2e70
CPID:                     0x0100918f
PPPOE Session Id:         0x0000
Type:                     IPoE
IPv4 Address:             0.0.0.0
IPv4 Framed Route:
  Prefix:                 0.0.0.0/0
  Next Hop:               0.0.0.0
  Tag:                    0
IPv6 IANA Address:        1:5::345c
IPv6 IAPD Prefix:         2004:cd0:0:188d::/64
CPE link local Address:   ::
IPv6 Framed Route:
  Prefix:                 ::/0
  Next Hop:               ::
  Tag:                    0
IPv6 State:               UP, Sat Jul 25 02:09:55 2020
Mac Address:              5065.aaab.d864
Inner VLAN ID:            Not Set
Outer VLAN ID:            100
Outer VLAN Cos:           0
Outer VLAN DEI:           1
Created:                  Sat Jul 25 02:09:54 2020
State:                    Activated
Ifhandle:                 0x000b75a0
VRF:                      default
Access-interface:         Bundle-Ether1.1
 Attribute List: 0x5556aed3f878
1:  ipv6-enable     len=  4  value= 1(1)
2:  ipv4-unnumbered len=  9  value= Loopback1
3:  strict-rpf      len=  4  value= 1(1)
4:  ipv6-strict-rpf len=  4  value= 1(1)
5:  ipv4-icmp-unreachable len=  4  value= 1(1)
6:  ipv6-unreachable len=  4  value= 1(1)
7:  ipv4-mtu        len=  4  value= 1500(5dc)
8:  ipv6-mtu        len=  4  value= 1500(5dc)
Session Accounting:        enabled
Interim Interval:          1800 secs
Last interim timestamp:    Sun Aug  2 23:39:46 2020
Interim fail count: None
Last interim failed reason: NA
Last stats:
  BytesIn: 0
  BytesOut: 384570
  BytesInGiga: 0
  BytesOutGiga: 0
Feature IDs activated :
  0x800e2e71
  0x800e2e72
```

-
```
Router#show cnbng-nal subscriber type ipoe summary
Mon Aug  3 00:06:15.032 UTC
=====================
Location: 0/RSP0/CPU0
```

```
=====================

                Type            PPPoE           IPoE
                ====            =====           ====

Session Counts by State:
        initializing            0               0
          connecting            0               0
           connected            0               0
           activated            0               8000
                idle            0               0
       disconnecting            0               0
              Total:            0               8000

Session Counts by Address-Family:
                none            0               0
                ipv4            0               0
                ipv6            0               8000
                dual            0               0
              Total:            0               8000

=====================
Location: 0/RSP1/CPU0
=====================

                Type            PPPoE           IPoE
                ====            =====           ====

Session Counts by State:
        initializing            0               0
          connecting            0               0
           connected            0               0
           activated            0               8000
                idle            0               0
       disconnecting            0               0
              Total:            0               8000

Session Counts by Address-Family:
                none            0               0
                ipv4            0               0
                ipv6            0               8000
                dual            0               0
              Total:            0               8000

Router#
```

-
```
Router#show cnbng-nal subscriber type pppoe summary
Mon Aug  3 00:06:15.032 UTC
=====================
Location: 0/RSP0/CPU0
=====================

                Type            PPPoE           IPoE
                ====            =====           ====

Session Counts by State:
        initializing            0               0
          connecting            0               0
           connected            0               0
           activated            31031           0
                idle            0               0
       disconnecting            0               0
              Total:            31031           0
```

```
Session Counts by Address-Family:
               none            0              0
               ipv4            31031          0
               ipv6            0              0
               dual            0              0
              Total:           31031          0

Router#
```

-
```
Router#show cnbng-nal subscriber disconnect-history unique
Mon Aug  3 00:07:22.716 UTC

Location: 0/RSP1/CPU0

| Count|       Last Interface        | Disconnected Reason |            Last Time
                                                                       Disconnected
 Location: 0/1/CPU0
 Location: 0/RSP0/CPU0

| Count|       Last Interface        | Disconnected Reason |            Last Time
                                                                       Disconnected
35494    Bundle-Ether1.1.ip2148328848  Disconnect by CP                 Sat Jul 25
                                                                       02:04:55 2020

14154    Bundle-Ether1.1.ip2148324096  Disconnect by clear CLI          Sat Jul 25
                                                                       02:05:48 2020

 2777    Bundle-Ether1.1.ip2148194512  Disconnect due to create failure Sat Jul 25
                                                                       01:38:29 2020
```

-
```
Router#show cnbng-nal subscriber disconnect-history last location
Mon Aug  3 00:08:42.655 UTC

Disconnect-reason:         Disconnect by clear CLI
Disconnect-timestamp:      Sat Jul 25 02:05:48 2020
  Message Txn ID: 55663
  Session Txn ID: 1
  Failed at: Sat Jul 25 01:57:03 2020
  Feature Mask: 0x0
  SVM State: 0
  IPSUB flags: 0x600a200
  Pending callback: 0x2
  Data:

Interface:                 Bundle-Ether1.1.ip2148324096
UPID:                      0x800cd300
CPID:                      0x01007bd8
PPPOE Session Id:          0x0000
Type:                      IPoE
IPv4 Address:              0.0.0.0
IPv4 Framed Route:
  Prefix:                  0.0.0.0/0
  Next Hop:                0.0.0.0
  Tag:                     0
IPv6 IANA Address:         1:5::3de5
IPv6 IAPD Prefix:          2004:cd0:0:616::/64
CPE link local Address:    ::
IPv6 Framed Route:
  Prefix:                  ::/0
  Next Hop:                ::
  Tag:                     0
```

```
IPv6 State:                UP, Sat Jul 25 01:57:03 2020
Mac Address:               5065.aaab.cfbb
Inner VLAN ID:             Not Set
Outer VLAN ID:             100
Outer VLAN Cos:            0
Outer VLAN DEI:            1
Created:                   Sat Jul 25 02:05:48 2020
State:                     Init
Ifhandle:                  0x000323a0
VRF:                       default
Access-interface:          Bundle-Ether1.1
 Attribute List: 0x559125764408
1:  ipv6-enable     len=  4  value= 1(1)
2:  ipv4-unnumbered len=  9  value= Loopback1
3:  strict-rpf      len=  4  value= 1(1)
4:  ipv6-strict-rpf len=  4  value= 1(1)
5:  ipv4-icmp-unreachable len=  4  value= 1(1)
6:  ipv6-unreachable len=  4  value= 1(1)
7:  ipv4-mtu        len=  4  value= 1500(5dc)
8:  ipv6-mtu        len=  4  value= 1500(5dc)
Session Accounting:        enabled
Interim Interval:          1800 secs
Last interim timestamp:    Sat Jul 25 02:05:47 2020
Interim fail count: None
Last interim failed reason: NA
Last stats:
  BytesIn: 0
  BytesOut: 540
  BytesInGiga: 0
  BytesOutGiga: 0
Feature IDs activated :
  0x800cd301
  0x800cd302

[Event History]
UPID:  0x800cd300

| Event Name               | Time Stamp              | S, M
| Create                   | Jul 25 01:57:02.999679 | 0, 0
| New Session Request      | Jul 25 01:57:02.999686 | 0, 0
| Interface create         | Jul 25 01:57:02.999823 | 0, 0
| SVM create               | Jul 25 01:57:03.018268 | 0, 0
| UP Install(req)          | Jul 25 01:57:03.018321 | 0, 0
| UP Install(CB)           | Jul 25 01:57:03.019220 | 0, 0
| Last Assoc(req)          | Jul 25 01:57:03.019232 | 0, 0
| Last Assoc(CB)           | Jul 25 01:57:03.020160 | 0, 1
| Produce done(req)        | Jul 25 01:57:03.020233 | 0, 0
| IPv4 Caps Up             | Jul 25 01:57:03.188034 | 0, 0
| IPv6 Caps Up             | Jul 25 01:57:03.233210 | 0, 0
| Init data req            | Jul 25 01:57:03.254482 | 0, 1
| Init data cb             | Jul 25 01:57:03.369027 | 0, 1
| Client Session up        | Jul 25 01:57:03.379152 | 0, 0
| Produce done             | Jul 25 01:57:03.977629 | 0, 0
| IPv6 Up                  | Jul 25 01:57:03.977643 | 0, 0
| Session up notified      | Jul 25 01:57:03.977650 | 0, 0
| Stats start              | Jul 25 01:57:03.977841 | 0, 0
| Disconnect notified      | Jul 25 02:05:47.548202 | 0, 0
| Disconnect ack           | Jul 25 02:05:47.550293 | 0, 0
| IPv4 Caps Down           | Jul 25 02:05:47.652232 | 0, 0
| IPv6 Caps Down           | Jul 25 02:05:47.652333 | 0, 0
| Final stats              | Jul 25 02:05:47.753805 | 0, 0
| SVM delete               | Jul 25 02:05:47.780713 | 0, 0
| SVM cleanup              | Jul 25 02:05:48.283050 | 0, 0
```

```
Help: S - Sticky Event, M - Multiple Occurrence
```

-

```
Router#show cnbng-nal subscriber fadb
Mon Aug  3 00:03:12.858 UTC

Location: 0/RSP1/CPU0
==================

UPID:       0x800ec810
Service-ID: 0x04000003  Service-Name: JHV_VOICE
Feature-ID: 0x800ec812
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 805306413(3000002d)
Accounting:               enabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0

UPID:       0x800e9470
Service-ID: 0x04000003  Service-Name: JHV_VOICE
Feature-ID: 0x800e9472
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 805306413(3000002d)
Accounting:               enabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0

UPID:       0x800e7ee0
Service-ID: 0x04000003  Service-Name: JHV_VOICE
Feature-ID: 0x800e7ee2
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 805306413(3000002d)
Accounting:               enabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0

UPID:       0x800e16e0
Service-ID: 0x04000004  Service-Name: LIVE_TV
Feature-ID: 0x800e16e1
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 0(0)
Accounting:               disabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
```

```
  BytesOutGiga: 0

UPID:       0x800dda90
Service-ID: 0x04000003  Service-Name: JHV_VOICE
Feature-ID: 0x800dda91
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 805306413(3000002d)
Accounting:               enabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0

UPID:       0x800dd4e0
Service-ID: 0x04000004  Service-Name: LIVE_TV
Feature-ID: 0x800dd4e1
 Attribute List: 0x559cba6d0008
1:  feature-acct-bitmask len=  4  value= 0(0)
Accounting:               disabled
Interim fail count: None
Last interim failed reason: None
Last stats:
  BytesIn: 0
  BytesOut: 0
  BytesInGiga: 0
  BytesOutGiga: 0
```

Verify cnBNG NAL Counters

You can use the following commands to verify various NAL counters on the cnBNG user plane:

-

```
Router#show cnbng-nal counters type all
Sun Aug  2 20:42:49.548 UTC

Location: 0/RSP0/CPU0

Subscriber Counters
-------------------

Counter name                      Value
============                      =====
INTF Delete                       500
IPv4 caps down                    500
IPv6 caps down                    500
IPv4 Rou del                      500
IPv6 Rou del                      500
Blkdis q empty                    1
DB cache hit                      17113

Error Counters
--------------

Counter name                      Value
============                      =====

Accounting Counters
-------------------
```

```
Counter name                    Value
============                    =====
Sess Stop req                   500
Feat Stop req                   500
Stop req                        3000
Stop cb                         3000
Final cb                        3000
Feat Final cb                   500
Sess Final cb                   2500

SVM Counters
------------

Counter name                    Value
============                    =====
Sess deleted                    500
Delete CB                       500
Feat deleted                    1000
Cleanup                         500
Sess stats, before svm          500
Feat stats, before svm          500

SPA Counters
------------

Counter name                    Value
=============                   =====
SPA Delete Req                  500
SPA Update Req                  500
Sub Delete Res                  500
Sub Update Res                  500
Blkdic adm more                 39
GTPu pkt sent                   1000
PFCP pkt sent                   1463
GTPu pkt punt                   500
PFCP pkt punt                   1463
DHCPv4 pkt punt                 500
DHCPv6 pkt punt                 500
DHCPv6 pkt inj                  500
Alloc count                     3463
Free count                      3463
Mutex lock                      6741
Mutex unlock                    6741
Timer start                     463
Timer expiry                    463
Sub Update IPOE OK              500
Sub Delete IPOE OK              500

CP Recon Counters
-----------------

Counter name                    Value
=============                   =====

Histogram/API Performance Stats
-------------------------------

API name       1ms       10ms  100ms  1s  5s  10s  20s  50s  100s
========       ===       ====  =====  ==  ==  ===  ===  ===  ====
Per trans      410       90    0      500 0   0    0    0    0
Sub Create     0         0     0      0   0   0    0    0    0
Sub Update     445       55    0      0   0   0    0    0    0
Sub Delete     0         0     0      500 0   0    0    0    0
```

```
IPOE Int Crt   0         0     0      0   0   0    0    0    0
IPOE Int Upd   0         0     0      0   0   0    0    0    0
IPOE Int Del   0         0     0      500 0   0    0    0    0
PPPOE Int Crt  0         0     0      0   0   0    0    0    0
PPPOE Int Upd  0         0     0      0   0   0    0    0    0
PPPOE Int Del  0         0     0      0   0   0    0    0    0
Sess Create    0         0     0      0   0   0    0    0    0
Sess Update    0         0     0      0   0   0    0    0    0
Sess Delete    0         0     10     490 0   0    0    0    0
V4 RT Inst     0         0     0      0   0   0    0    0    0
V4 RT Del      0         6     320    174 0   0    0    0    0
V4 FR Inst     0         0     0      0   0   0    0    0    0
V4 FR Del      0         0     0      0   0   0    0    0    0
V6 RT Inst     0         0     0      0   0   0    0    0    0
V6 RT Del      0         6     310    184 0   0    0    0    0
V6 PD RT Inst  0         0     0      0   0   0    0    0    0
V6 PD RT Del   0         0     0      0   0   0    0    0    0
V6 FR Inst     0         0     0      0   0   0    0    0    0
V6 FR Del      0         0     0      0   0   0    0    0    0
CDM Lookup     0         0     0      0   0   0    0    0    0
CDM Insert     0         0     0      0   0   0    0    0    0
CDM Update     1469      31    0      0   0   0    0    0    0
Eval Lookup    0         0     0      0   0   0    0    0    0
```

-
```
Router#show cnbng-nal counters type all | beg SPA LIB
Sun Aug  2 20:44:07.902 UTC
SPA LIB Counters
----------------

Counter name                      Value
=============                     =====
pfcp_rx_counter                   6899
pfcp_tx_counter                   6900
gtpu_tx_counter                   9048
gtpu_rx_counter                   7510
pfcp_keepalive_tx_counter         891
pfcp_keepalive_rx_counter         890

SPA API counters
----------------------
```

-
```
Router#show cnbng-nal counters type spa
Sun Aug  2 20:42:13.703 UTC

Location: 0/RSP0/CPU0

SPA Counters
------------

Counter name                      Value
=============                     =====
SPA Delete Req                    500
SPA Update Req                    500
Sub Delete Res                    500
Sub Update Res                    500
Blkdic adm more                   39
GTPu pkt sent                     1000
PFCP pkt sent                     1461
GTPu pkt punt                     500
PFCP pkt punt                     1461
```

```
DHCPv4 pkt punt                  500
DHCPv6 pkt punt                  500
DHCPv6 pkt inj                   500
Alloc count                      3461
Free count                       3461
Mutex lock                       6727
Mutex unlock                     6727
Timer start                      461
Timer expiry                     461
Sub Update IPOE OK               500
Sub Delete IPOE OK               500
```

P-BNG and cnBNG coexistence

P-BNG and cnBNG coexistence is a BNG capability that

- supports the simultaneous operation of physical BNG (P-BNG) and Cloud-Native BNG (cnBNG) on a single router
- ensures each interface uses only one BNG type at a time to prevent configuration conflicts, and
- allows for a gradual transition from traditional p-BNG to cnBNG architectures with minimal service disruption.

Table 9: Feature History Table

Feature name	Release information	Feature description
P-BNG and cnBNG coexistence	Release 26.2.1	Enhances network agility and migration efficiency by enabling the simultaneous operation of physical BNG and Cloud-Native BNG (cnBNG) on a single router. This feature provides a flexible mechanism for a phased transition with minimal disruption to cloud-native architectures, supporting interface-level configuration exclusivity, service continuity, and unified subscriber management across hybrid deployments.

Need for P-BNG and cnBNG coexistence

Service providers are increasingly adopting cnBNG architectures to improve network agility, scalability, and resource efficiency. Migration from traditional physical BNG architectures to cloud-native environments requires a strategy that minimizes service impact.

Phased migration: You can migrate subscriber groups gradually, which eliminates the need for high-risk migrations and helps you maintain service continuity.

Operational flexibility: During long-term transition phases, your network may require a mix of legacy and cloud-native architectures. This feature allows both systems to operate simultaneously on a single platform, providing the flexibility to manage hybrid environments.

Benefits

- Seamless migration: Supports gradual migration from P-BNG to cnBNG at the SRG group level with minimal service disruption.
- Operational clarity: Clear identification of interface types prevents configuration conflicts.
- Backward compatibility: Existing P-BNG configurations continue to function unchanged.
- High availability: Maintains service continuity through SRG failover mechanisms during upgrades and migrations.
- Support for mixed subscribers: Enables concurrent handling of IPoE subscribers for both P-BNG and cnBNG, as well as PPPoE PTA, LAC, and LNS specifically for cnBNG subscribers.

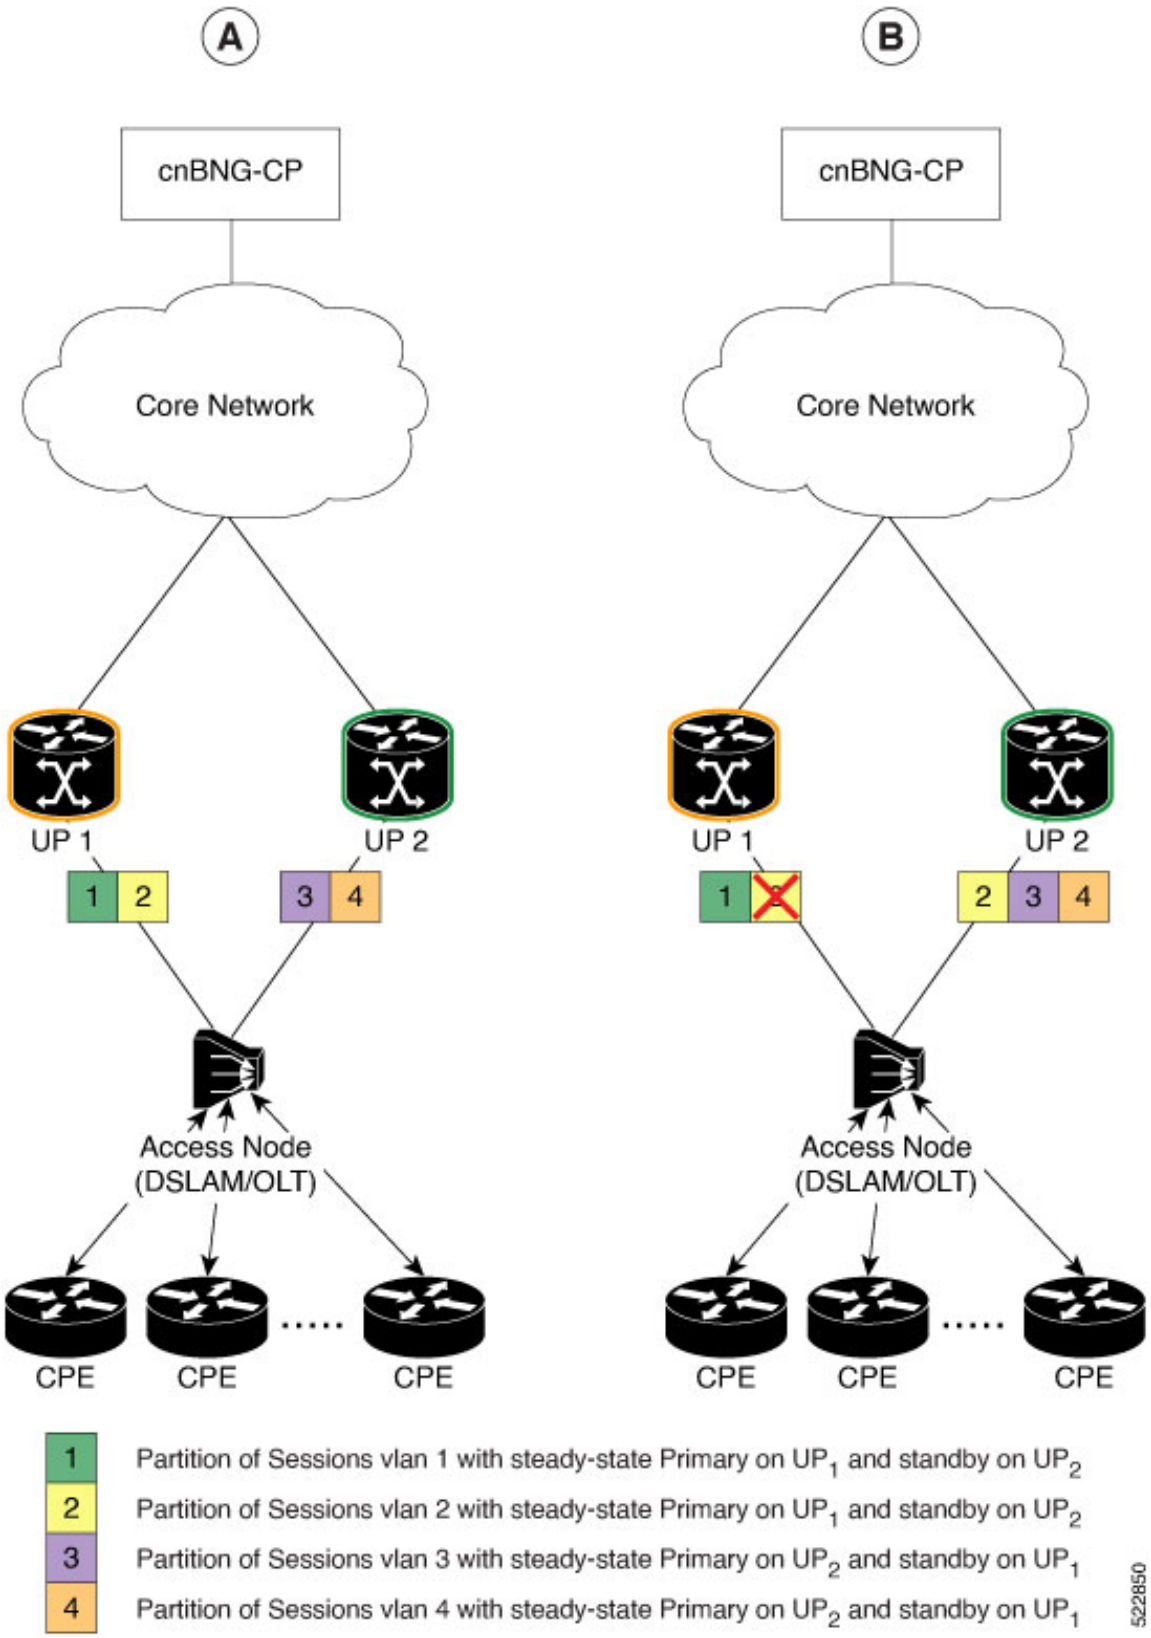
Limitations of P-BNG and cnBNG coexistence

- The coexistence solution supports IPoE subscribers for both P-BNG and cnBNG, PPPoE PTA for cnBNG, and LAC/LNS for cnBNG. PPPoE and LAC support for P-BNG subscribers relies on existing P-BNG RPMs.
- Mixed subscriber counts of P-BNG and cnBNG must not exceed the maximum scale limits defined for each profile.
- Certain dynamic configuration changes are restricted when subscribers are active:
- Subscriber Redundancy Group (SRG) cannot be deleted or modified (including virtual MAC, access tracking, damping timer) while subscribers are present in those groups.
- Interfaces cannot be added or deleted from SRG groups while active subscribers are present on those interfaces.
- Route tags cannot be modified when dynamic routes are present; dynamic routes must be cleared first.

Enable P-BNG and cnBNG coexistence

Use this task to migrate subscriber groups from P-BNG to cnBNG at the SRG group-level gradually with minimal service interruption while existing P-BNG configurations continue to function.

This task migrates subscribers from PBNG SRG group 1 on UP1 to cnBNG SRG group 17. During the migration, UP2 temporarily handles subscriber traffic for the group. After the cnBNG configuration is added and the control plane creates the new group, subscriber sessions are restarted and recreated as cnBNG sessions. Then you configure the cnBNG peer group on UP2 and run SRG synchronization from the control plane.


A
B
cnBNG-CP
cnBNG-CP
Core Network
Core Network
UP 1
UP 2
UP 1
UP 2
1
2
3
4
1
2
3
4
Access Node
(DSLAM/OLT)
Access Node
(DSLAM/OLT)
CPE
CPE
CPE
CPE
CPE
CPE
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
522850

Before you begin

Before you start the migration, ensure that you check the status of these

- cnBNG control plane operational and connected to both BNGs
- CP-UP PFCP session established
- Existing cnBNG groups (if any) operational
- Maintenance window scheduled
- Configuration backup taken for both BNGs

Ensure that you have these information for each group to be migrated:

Table 10:

Attribute	Values for the example in this section
p-BNG SRG Group ID	1
New cnBNG SRG Group ID	17
Access Interface	Bundle-Ether20.1
VMAC	0002.0100.1015
Route Tag	117
VRF	default / RJIL-VRF-OLT-MGMT
Loopback	Loopback1
Current p-BNG Primary	BNG1 / BNG2
Access Tracking Object	access-17
Core Tracking Object	CORE
DHCP v4 Proxy Profile	ipoe_proxy1
DHCP v6 Proxy Profile	ipoev6_proxy1001015
CP Subscriber Profile	DUAL_STACK_SUB_GR_SET
Resource ID (per LC)	0, 1, 2...

Follow these steps to configure P-BNG and cnBNG coexistence.

Procedure

Step 1 Use the **subscriber redundancy switchover** command to move the subscriber traffic for SRG group 1 from UP1 to UP2.

Example:

```
pbng-UP1# subscriber redundancy switchover group 1
```

Step 2 Verify that UP1 is no longer handling subscriber traffic for SRG group 1.

Confirm that UP1 changes to standby and UP2 becomes active for the group before you modify the configuration on UP1.

Example:

```
pbng-UP1#show subscriber redundancy group

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

   Node Name   | Group ID | Role | Flags  |         Peer Address        | P/S | I/F Count |  SS Count
  | Sync Pending
--------------------------------------------------------------------------------------------------------------------

 0/RSP0/CPU0            1  Slave   ESHT           10.1.1.2                E            1
4000              0
--------------------------------------------------------------------------------------------------------------------
```

Step 3 Delete the configuration for P-BNG SRG group 1 on UP1

a) Disable the subscriber redundancy group on UP1.

Example:

```
pbng-UP1# #configure
pbng-UP1(config)# subscriber redundancy group 1
pbng-UP1(config-subscr-red)# disable
```

b) Clear p-BNG sessions on the primary BNG.

Example:

```
pbng-UP1#clear subscriber session identifier access-interface Bundle-Ether20.1 force
```

c) Verify that there are zero sessions on the access interface after you clear the sessions using the show **subscriber redundancy group** command.

Note
When you delete the P-BNG configuration for the group on UP1, the sessions created under group 1 on UP1 are also deleted.

Step 4 Configure cnBNG SRG group 17 on UP1 and control plane.

Configure the user plane first. Then, add the same group on the control plane.

a) Add the SRG configuration on user plane for SRG group 17.

Example:

```
cnbng-UP1#configure
cnbng-UP1(config)#cnbng-nal location 0/0/CPU0
cnbng-UP1(config-cnbng-nal-local)#subscriber-redundancy
cnbng-UP1(config-cnbng-nal-sub-red)#group group17
cnbng-UP1(config-cnbng-nal-srg-grp)#virtual-mac 0xxx.xxxx.0c01
cnbng-UP1(config-cnbng-nal-srg-grp)# core-tracking core1
cnbng-UP1(config-cnbng-nal-srg-grp)#access-tracking track1
cnbng-UP1(config-cnbng-nal-srg-grp)#access-interface-list
cnbng-UP1(config-cfg-srg-grp-intf)#interface Bundle-Ether20.1
cnbng-UP1(config-cfg-srg-grp-intf)# exit
```

a) Create the SRG configuration on the control plane.

Example:

```
cnbng-Up1#configure
user-plane instance UP1
 user-plane CNBNG-UP1
  port-id Bundle-Ether20.1
   subscriber-profile DUAL_STACK_SUB_GR_SET
  exit
  subscriber-redundancy
   group 17
    preferred-role-active
    revertive-timer 720
    peer-identifier Group_17
    port-id-map port-name Bundle-Ether20.1 17
   exit
```

b) Synchronize the routes between the control plane and the user plane.

Example:

```
cnbng-CP#subscriber route-synchronize upf CNBNG-BNG1
```

The control plane sends a create message to the user plane to create the SRG group.

Step 5 Shut down the access interface on UP2 so that the SRG group1 becomes standby.

Step 6 Delete the SRG group configuration on UP2.

a) Disable the subscriber redundancy group on both BNGs.

Example:

```
pbng-UP2#configure
pbng-UP2(config)#subscriber redundancy group 1
pbng-UP2(config-subscr-red)# disable
```

b) Clear p-BNG sessions on the primary BNG.

Example:

```
pbng-UP2#clear subscriber session identifier access-interface Bundle-Ether20.1 force
```

c) Verify that there are zero sessions on the access interface after you clear the sessions.

Step 7 Restart the CPE to recreate subscriber sessions as cnBNG sessions on UP1.

Step 8 Configure cnBNG SRG group 17 on UP1 and control plane.

Configure the user plane first, then add the same group on the control plane.

a) Add the SRG configuration on user plane for the SRG group 17.

Example:

```
cnbng-UP2#configure
cnbng-UP2(config)#cnbng-nal location 0/0/CPU0
cnbng-UP2(config-cnbng-nal-local)#subscriber-redundancy
cnbng-UP2(config-cnbng-nal-sub-red)#group group17
cnbng-UP2(config-cnbng-nal-srg-grp)#virtual-mac 0xxx.xxxx.0c01
cnbng-UP2(config-cnbng-nal-srg-grp)# core-tracking core1
cnbng-UP2(config-cnbng-nal-srg-grp)#access-tracking track1
cnbng-UP2(config-cnbng-nal-srg-grp)#access-interface-list
cnbng-UP2(config-cfg-srg-grp-intf)#interface Bundle-Ether20.1
cnbng-UP2(config-cfg-srg-grp-intf)# exit
```

a) Create SRG configuration on control plane.

Example:

```
Cnbng-Up1#configure
user-plane instance UP2
 user-plane CNBNG-UP2
  port-id Bundle-Ether20.1
   subscriber-profile DUAL_STACK_SUB_GR_SET
  exit
  subscriber-redundancy
   group 17
    preferred-role-active
    revertive-timer 720
    peer-identifier Group_17
    port-id-map port-name Bundle-Ether20.1 17
   exit
```

Step 9 Synchronize the routes between the control plane and the user plane.

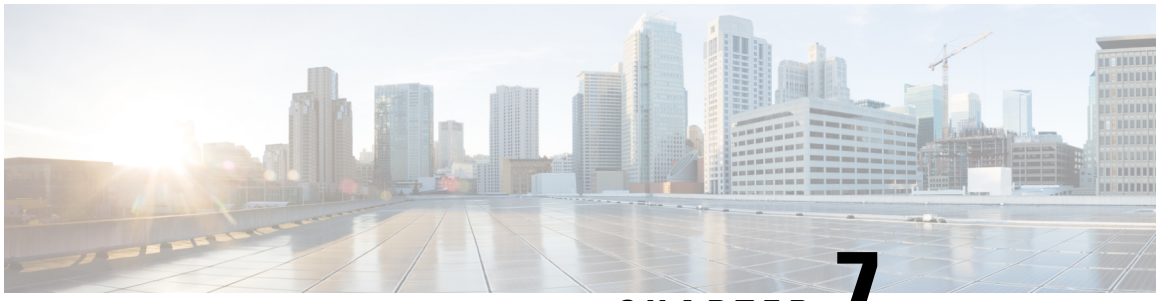
Example:

```
cnbng-CP#subscriber route-synchronize upf CNBNG-BNG2
```

Verify the status of route synchronization.

```
cnbng-CP#show subscriber route-synchronize upf upf-name status
```

After migration, subscriber sessions are recreated as cnBNG sessions for SRG group 17.



CHAPTER 7

Subscriber Management

This chapter provides information about various types of subscriber sessions, namely IPoE and PPPoE, and IP addressing by DHCP. Also, on how the point-point frames are tunnelled across the network using the Layer 2 Tunneling Protocol.

- Subscriber Session Overview, on page 91
- IPoE Session, on page 92
- PPP over Ethernet (PPPoE), on page 93
- Enable SLAAC for PPPoE Subscriber Sessions, on page 102
- RADIUS-Based Policing - QoS shape-rate parameterization, on page 106
- Shared Policy Instance, on page 110
- Enhanced subscriber routing with Framed-Route Tag and Preference attributes, on page 117
- Geographical redundancy for L2TP sessions, on page 122

Subscriber Session Overview

To enable subscribers to access the network resources, the network has to establish a session with the subscriber. A subscriber session represents the logical connection between the customer premise equipment (CPE) and the network resource. Each session establishment comprises the following phases:

- Establishing a connection—in this phase CPE finds the cnBNG with which to communicate.
- Authenticating and authorizing the subscriber—in this phase, cnBNG authenticates the subscribers and authorizes them to use the network. This phase is performed with the help of the RADIUS server.
- Giving the subscriber an identity—in this phase, the subscriber is assigned an identity, the IP address.
- Monitoring the session—in this phase, cnBNG ascertains that the session is up and running.

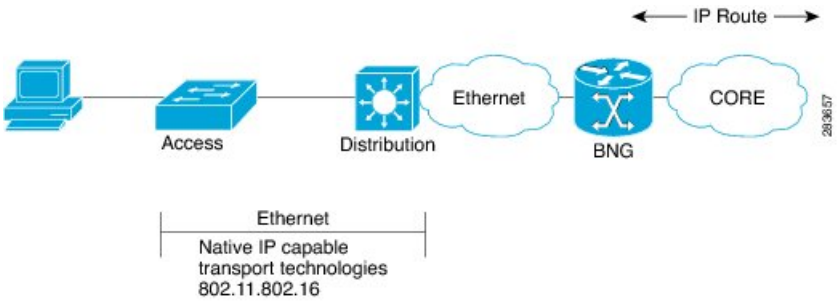
The subscriber sessions are established over the subscriber interfaces, which are virtual interfaces. It's possible to create only one interface for each subscriber session. A port can contain multiple VLANs, each of which can support multiple subscribers. cnBNG creates subscriber interfaces for each kind of session. These interfaces are named based on the parent interface, such as bundle-ether 2.100.pppoe312. The subscribers on bundle interfaces (or bundle-VLANs) allow redundancy and are managed on the cnBNG route processor (RP).

There are two mechanisms to establish a subscriber session, namely, IPoE and PPPoE.

IPoE Session

In an Internet over Ethernet (IPoE) subscriber session, subscribers run IPv4 or IPv6 on the CPE device and connect to the cnBNG through a Layer-2 aggregation. IP subscriber sessions that connect through a Layer-2 aggregation network are called L2-connected. IPoE subscriber sessions are always terminated on cnBNG and then routed into the service provider network. IPoE relies on DHCP to assign the IP address.

Figure 8: IPoE Session



cnBNG supports both DHCP v4 and DHCP v6 subscriber sessions.

Limitations

The following are the limitations:

- L3 routed subscribers are not supported.
- Geo redundancy or subscriber redundancy is not supported.
- Line card or physical port termination-based subscribers are not supported.

Configuration Example

```
Router#configure
Router(config)#interface Bundle-Ether1.1
Router(config-subif)#ipv4 point-to-point
Router(config-subif)#ipv4 unnumbered Loopback1
Router(config-subif)#ipv6 enable
Router(config-subif)#encapsulation dot1q 1
Router(config-subif)#ipsubscriber
Router(config-cnbng-nal-ipsub)#ipv4 l2-connected
Router(config-cnbng-nal-ipsub-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-l2conn)#exit
Router(config-cnbng-nal-ipsub)#ipv6 l2-connected
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#initiator dhcp
Router(config-cnbng-nal-ipsub-ipv6-l2conn)#commit
```

Running Configuration

```
Router#show running-config interface be1.1
interface Bundle-Ether1.1
 ipv4 point-to-point
 ipv4 unnumbered Loopback1
 ipv6 enable
```

```
encapsulation dot1q 1
ipsubscriber
 ipv4 l2-connected
  initiator dhcp
 !
 ipv6 l2-connected
  initiator dhcp
 !
!
```

PPP over Ethernet (PPPoE)

The Point-to-Point Protocol (PPP) is used for communications between two nodes, like a client and a server. The PPP provides a standard method for transporting multiprotocol datagrams over point-to-point links. It defines an encapsulation scheme, a link layer control protocol (LCP), and a set of network control protocols (NCPs) for different network protocols that can be transmitted over the PPP link.

One of the methods to establish PPP connection is by the use of PPPoE. In a PPPoE session, the PPP protocol runs between the CPE and cnBNG. The Home Gateway (which is part of the CPE) adds a PPP header (encapsulation) that is terminated at the cnBNG.

PPPoE Discovery

The PPPoE discovery-stage protocol consists of basic packet exchange between the subscriber and server (cnBNG). The following is the list of the various PPPoE Active Discovery (PAD) messages:

- PPPoE Active Discovery Initiation (PADI)—The CPE broadcasts to initiate the process to discover cnBNG.
- PPPoE Active Discovery Offer (PADO)—The cnBNG responds with an offer.
- PPPoE Active Discovery Request (PADR)—The CPE requests to establish a connection.
- PPPoE Active Discovery Session confirmation (PADS)—cnBNG accepts the request and responds by assigning a session identifier (Session-ID).
- PPPoE Active Discovery Termination (PADT)—Either CPE or cnBNG terminates the session.

PPoE Sessions

The PPPoE sessions are of the following types:

- PPPoE PPP Terminated sessions Terminated (PTA)
- PPPoE L2TP Access Concentrator Sessions (LAC)
- L2TP Network Server Sessions (LNS)

Majority of the digital subscriber line (DSL) broadband deployments use Point-to-Point Protocol over Ethernet (PPPoE) sessions to provide subscriber services. These sessions terminate the Point-to-Point Protocol (PPP) link and provide all the features, service, and billing on the same node. These sessions are called PPP Terminated (PTA) sessions. See PPPoE PPP Terminated and Aggregation Sessions (PPPoE-PTA), on page 94.

There are some wireline subscriber deployments in the wholesale retail model where ISPs work with others to provide the access and core services separately. In such cases, the subscribers are tunneled between wholesale

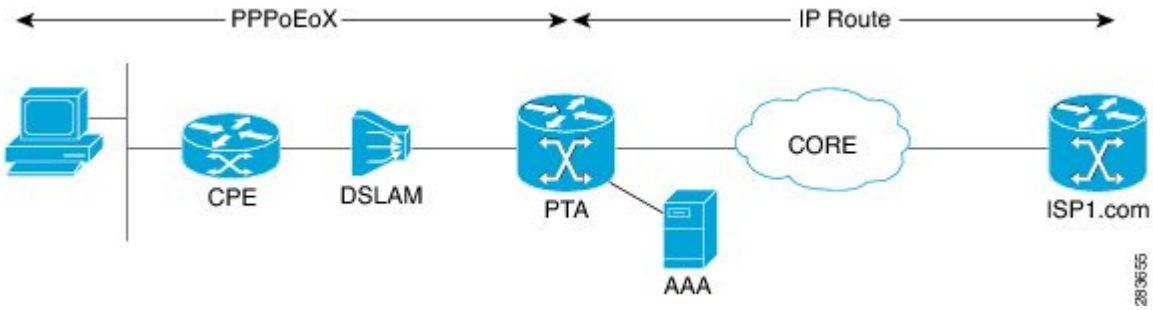
and retail ISPs using the Layer 2 Tunneling Protocol (L2TP), a client-server protocol. See L2TP Access Concentrator Sessions (LAC), on page 95 and L2TP Network Server Sessions (LNS), on page 99.

Note For the functioning of PPP PTA and PPP LAC session, the RADIUS server must be set up to authenticate and forward sessions as necessary. There's no local authentication available on cnBNG.

PPPoE PPP Terminated and Aggregation Sessions (PPPoE-PTA)

In a PPPoE-PPP Termination and Aggregation (PTA) session, the PPP encapsulation is terminated on cnBNG. After it's terminated, cNBNG routes the traffic to the service provider using IP routing. A typical PTA session is depicted in this figure.

Figure 9: PPPoE-PTA Session



PPPoE session configuration information is contained in PPPoE profiles. After a profile is defined, it's assigned to an access interface. Multiple PPPoE profiles can be created and assigned to multiple interfaces. A global PPPoE profile can also be created; the global profile serves as the default profile for any interface that has not been assigned a specific PPPoE profile.

The PPP PTA session is typically used in the Network Service Provider (retail) model where the same service operator provides the broadband connection to the subscriber and also manages the network services.

Limitations

The following are the limitations:

- L3 routed subscribers are not supported.
- Geo redundancy or subscriber redundancy is not supported.
- Line card or physical port termination-based subscribers aren't supported.

Configure PPPoE-PTA Session

The following section describes the steps to configure PPPoE-PTA sessions:

- Configure the access-interface
- Enable PPPoE

Configuration Example

```
Router#configure
Router(config)#interface Bundle-Ether1.1
Router(config-subif)#ipv4 point-to-point
Router(config-subif)#ipv4 unnumbered Loopback1
Router(config-subif)#ipv6 enable
Router(config-subif)#encapsulation dot1q 1

/* Enable PPPoE */
Router(config-subif)#pppoe enable
Router(config-subif)#commit
```

Running Configuration

```
Router#show running-config interface be1.1
interface Bundle-Ether1.1
 ipv4 point-to-point
 ipv4 unnumbered Loopback1
 ipv6 enable
 encapsulation dot1q 1

 pppoe enable
!
```

L2TP Access Concentrator Sessions (LAC)

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
Enable LAC on Cloud Native BNG	Release 25.2.1	You can now enable the cloud-native BNG user plane to act as an L2TP Access Concentrator (LAC) on these Cisco ASR 9000 5th Generation High-Density Ethernet line cards and fixed routers, facilitating the tunneling of point-to-point frames between a remote system or LAC client and an LNS. • A99-32X100GE-X-SE • A9K-20HG-FLEX-SE • A9K-8HG-FLEX-SE • A9K-4HG-FLEX-SE • Cisco ASR 9902 Router • Cisco ASR 9903 Router

Feature Name	Release Information	Feature Description
Enable LAC on Cloud Native BNG	Release 7.4.2	This feature enables the cloud native BNG user plane to become an L2TP access concentrator (LAC), allowing you to tunnel point-to-point frames between the remote system or LAC client and an LNS located at a wholesaler. This functionality provides highly flexible deployments options to suit different customer use-cases and needs. To enable this feature, use the **l2tp enable** command.

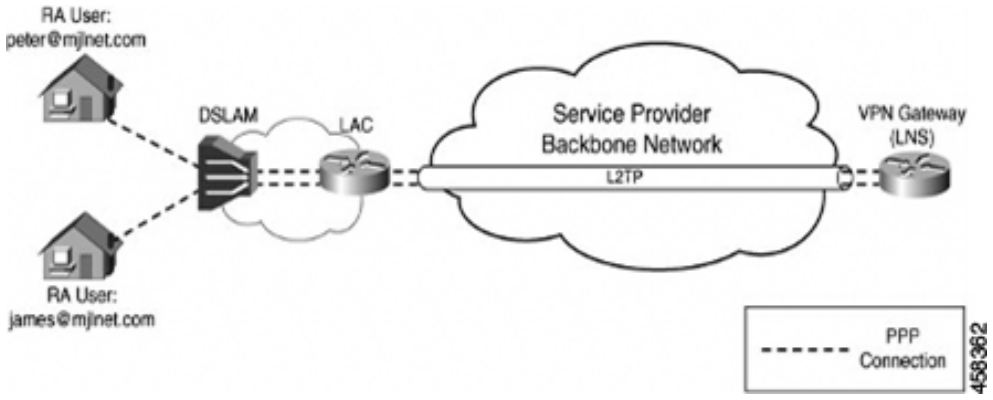
L2TP encapsulates and tunnels the PPP Layer 2 frames through a Layer 3 network. With L2TP, you can have a layer 2 connection to an access concentrator. The concentrator then tunnels individual PPP frames to the Network Access Server (NAS). This allows the processing of PPP packets on different devices. L2TP can be used to make all multilink channels terminate at a single NAS. Thus-allowing multilink operation even when the calls are spread across distinct physical NASs.

In cnBNG, L2TP uses the following two components to perform the hand-off task of the subscriber traffic to the Internet service provider (ISP).

- L2TP Access Concentrator (LAC)—The L2TP enables subscribers to dial into the LAC, which extends the PPP session to the LNS. cnBNG provides LAC.
- L2TP Network Server (LNS)—The L2TP extends PPP sessions over an arbitrary network to a remote network server that is, the LNS. The ISP provides LNS.

The following image depicts the overall topology of LAC and LNS:

Figure 10: Topology of LAC and LNS



The remote user initiates a PPP connection across the cloud to a LAC. The LAC acts as a client and then tunnels the PPP connection across the Internet to an LNS that acts as a server.

L2TP utilizes two types of messages, control messages and data messages. Control messages are used in the establishment, maintenance, and clearing of tunnels and calls. Data messages are used to encapsulate PPP frames over the tunnel.

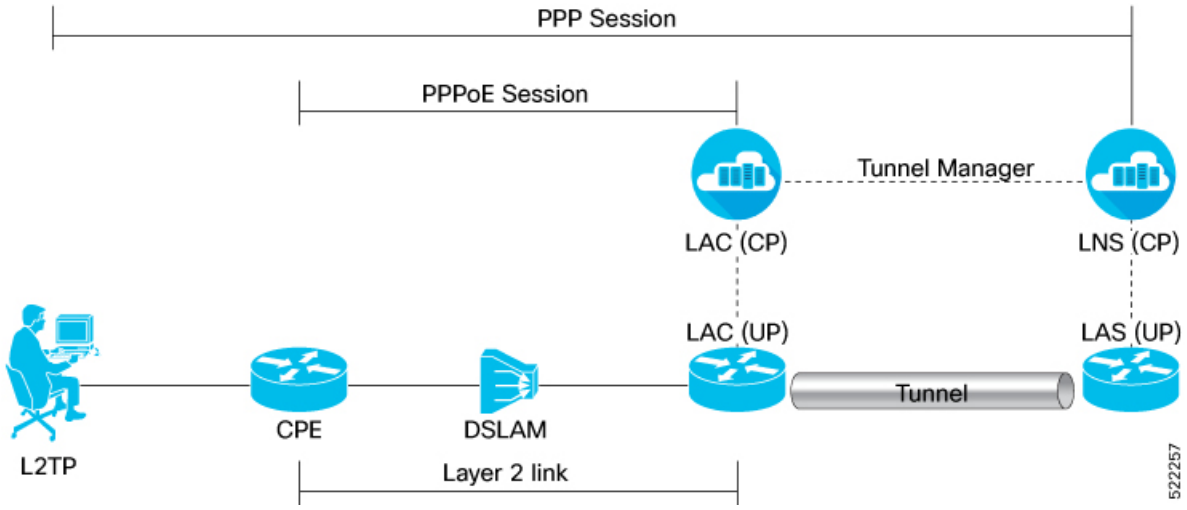
```
+--------------------+
| PPP Frames         |
```

```
+-------------------+    +-----------------------+
| L2TP Data Messages|    | L2TP Control Messages |
+-------------------+    +-----------------------+
| L2TP Data Channel |    | L2TP Control Channel  |
| (unreliable)      |    | (reliable)            |
+---------------------------------------------+
|      Packet Transport (UDP, FR, ATM)      |
+---------------------------------------------+
```

PPP frames are passed over an unreliable data channel that is encapsulated first by an L2TP header. Then a Packet Transport such as UDP. Control messages are sent over a reliable L2TP Control Channel, which transmits packets in-band over the same Packet Transport.

During a PPP LAC session, the PPPoE encapsulation terminates on cnBNG; however, the PPP packets travel beyond cnBNG to LNS through the L2TP tunnel. A typical LAC session is depicted in the following figure.

Figure 11: LAC Session



Both LAC and LNS sessions use L2TP protocol for negotiation and creation of L2TP sessions.

For more information on the LAC high-level work flow, see the *L2TP Subscriber Management* chapter in the *Cloud Native BNG Control Plane Configuration Guide*.

The PPP LAC session is used in the wholesaler model, where the network service provider is a separate entity from the local access network provider. In this kind of setup, the access network provider owns the LAC and the network service provider owns the LNS.

- Network service provider performs access authentication, manage and provide IP addresses to subscribers, and are responsible for overall service.
- The access network prover is responsible for providing the last-mile digital connectivity to the customer, and for passing on the subscriber traffic to the service provider.

Limitations for LAC Sessions

The following are the limitations for the LAC sessions:

- Tunnel specific statistics are not supported.

- LAC and LNS cannot coexist on the same node.
- IPv6 L2TP tunnel is not supported.
- L2TP tunnel keep alive or hello packet offload is not supported.
- Setting of type of service is not supported.
- Multicast group is not supported.
- L2TP packet segmentation or reassemble is not supported.
- The following features aren't supported:
  - Access Control List (ACL)
  - Quality of Service (QoS)
  - Policy-based Routing (PBR)
  - Unicast Reverse Path Forwarding (uRPF)
  - ICMP unreachable

Configure LAC Sessions

This section describes how to configure the LAC session on the cnBNG user plane.

- Enable L2TP
- Establish PPPoE connection

Configuration Example

Enable L2TP:

```
Router#configure
Router(config)#cnbng-nal location 0/1/CPU0

Router(config-cnbng-nal-local)#hostidentifier RTR1

Router(config-cnbng-nal-local)#up-server ipv4 192.0.2.1 gtp-port 15002 pfcp-port 15003
vrf default
Router(config-cnbng-nal-local)#cp-server primary ipv4 198.51.100.1

Router(config-cnbng-nal-local)#enable-test-server

Router(config-cnbng-nal-local)#disconnect-history file-logging-enable

Router(config-cnbng-nal-local)#cp-association retry-count 5

Router(config-cnbng-nal-local)#l2tp enable

Router(config-cnbng-nal-local)#l2tp-tcp-mss-adjust 1400
```

Establish PPPoE connection:

```
Router(config-cnbng-nal-local)#interface Bundle-Ether1.1
```

```
Router(config-subif)#ipv4 address 192.11.1.1 255.255.255.0

Router(config-subif)#ipv6 enable

Router(config-subif)#encapsulation dot1q 1

Router(config-subif)#ppoe enable
Router(config-subif)#commit
Router(config-subif)#exit
Router(config)#exit
```

Running Configuration

```
Router#show running-config

cnbng-nal location preconfigure 0/1/CPU0
 l2tp-tcp-mss-adjust 1400
 hostidentifier RTR1
 up-server ipv4 192.0.2.1 gtp-port 15002 pfcp-port 15003 vrf default
 cp-server primary ipv4 198.51.100.1
 disconnect-history file-logging-enable
 cp-association retry-count 5
 l2tp enable
 enable-test-server
!
interface Bundle-Ether1
!
interface Bundle-Ether1.1
 ipv4 address 192.11.1.1 255.255.255.0
 ipv6 enable
 encapsulation dot1q 1
 pppoe enable
!
```

L2TP Network Server Sessions (LNS)

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
Shared policy database (SPD) support on bundle main interfaces	Release 25.3.1	You can now support a larger number of shaper-based subscribers on bundle main interfaces across all Cisco ASR 9000 5th Generation High-Density Ethernet line cards.

Feature Name	Release Information	Feature Description
Enable LNS on Cloud Native BNG	Release 25.2.1	You can now enable the cloud-native BNG (cnBNG) to act as an L2TP Network Server (LNS), allowing the termination of tunnels or subscriber sessions initiated by the LAC client on the following Cisco ASR 9000 5th Generation High-Density Ethernet line cards: • A99-32X100GE-X-SE • A9K-20HG-FLEX-SE • A9K-8HG-FLEX-SE • A9K-4HG-FLEX-SE This feature is also supported on these fixed routers: • Cisco ASR 9902 Router • Cisco ASR 9903 Router
Enable LNS on Cloud Native BNG	Release 7.4.2	This feature enables cloud native BNG (cnBNG) to act as an L2TP Network Server (LNS) located at the wholesaler and allows you to terminate the tunnel or the subscriber sessions initiated by the LAC client. The cnBNG LNS solution offers control and user plane separation (CUPS) and cloud-native advantages for next-generation subscriber services in operator networks where subscribers connect directly to a retailer. To enable this feature, use the **lns enable** command.

L2TP Network Server (LNS) resides at one end of an L2TP tunnel and acts as a peer to the LAC. An LNS acts like an L2TP server that terminates the incoming tunnel from the L2TP LAC. An LNS is the logical termination point of the PPP session that is being tunneled from the client by the LAC.

LNS sessions are similar to PTA sessions in the overall functionality. Instead of the PPPoE protocol, here the First-Sign-Of-Life (FSOL) packets are the L2TP Incoming-Call-Request (ICRQ) messages.

For more information on the LNS high-level workflow, see the *L2TP Subscriber Management* chapter in the *Cloud Native BNG Control Plane Configuration Guide*.

Limitations for LNS Sessions

The following are the limitations for the LNS sessions:

- IPv6 L2TP tunnel is not supported.
- L2TP tunnel keep alive or hello packet offload is not supported.
- Tunnel statistics are not supported.
- Termination on non bundle-ether is not supported (for example, PWHE, physical interface).
- Termination of the VLAN interface is not supported.
- Supports parent interface only and not subinterface.
- L2TP packet segmentation or reassemble is not supported.
- Parent interface SVLAN policy must be different for other interfaces on the chassis.
- The following features are not supported:
  - Unicast Reverse Path Forwarding (uRPF)
  - Lawful Intercept (LI)
- If more than one bundle main interface (for example, BE1 and BE2) is configured with the same resource-id and their bundle members are from the same NP, the maximum number of subscribers that can be supported on that NP is limited to 1,500.

Configure LNS Sessions

This section describes how to configure the LNS session on the cnBNG user plane.

Configuration Example

To enable L2TP:

```
Router#configure
Router(config)#cnbng-nal location 0/0/CPU0
Router(config-cnbng-nal-local)#hostidentifier RTR1
Router(config-cnbng-nal-local)#up-server ipv4 192.0.2.1 gtp-port 15002 pfcp-port 15003
vrf default
Router(config-cnbng-nal-local)#cp-server primary ipv4 198.51.100.1
Router(config-cnbng-nal-local)#enable-test-server
Router(config-cnbng-nal-local)#disconnect-history file-logging-enable
Router(config-cnbng-nal-local)#cp-association retry-count 5
Router(config-cnbng-nal-local)#l2tp enable << Enble L2TP
Router(config-cnbng-nal-local)#commit
Router(config-cnbng-nal-local)#exit
Router(config)#
```

To establish the LNS session:

```
Router(config)#interface bundle-ether 1.1
Router(config-subif)#service-policy output SVLAN subscriber-parent subscriber-group
resourceid 4 << To allow maximum capacicty on the linecard
Router(config-subif)#ipv4 address 192.5.1.1 255.255.255.0
```

```
Router(config-subif)#ipv6 enable
Router(config-subif)#lns enable << Establish LNS session
Router(config-subif)#commit
Router(config-subif)#exit
```

Note To allow maximum capacity on the linecard, we recommend you to use the **service-policy output SVLAN subscriber-parent subscriber-group resourceid** command in the main interface.

Running Configuration

```
Router#show running-config

cnbng-nal location preconfigure 0/0/CPU0
 hostidentifier RTR1
 up-server ipv4 192.0.2.1 gtp-port 15002 pfcp-port 15003 vrf default
 cp-server primary ipv4 198.51.100.1
 disconnect-history file-logging-enable
 cp-association retry-count 5
 l2tp enable
 enable-test-server
!
interface Bundle-Ether1.1
 service-policy output SVLAN subscriber-parent subscriber-group resourceid 4
 ipv4 address 192.11.1.1 255.255.255.0
 ipv6 enable
 lns enable
!
```

Enable SLAAC for PPPoE Subscriber Sessions

The PPPoE support for SLAAC is a network protocol functionality that

- enables the use of the Point-to-Point Protocol over Ethernet (PPPoE) to establish network connections,
- facilitates Stateless Address Autoconfiguration (SLAAC) for IPv6 addresses,
- allows seamless integration of IPv6 address configuration in PPPoE environments, and
- allows networks to function without requiring a DHCPv6 server, as periodic Router Advertisements (RA) inform hosts of the active prefixes on a link.

Table 13: Feature History Table

Feature Name	Release Information	Feature Description
Enable SLAAC for PPPoE Subscriber Sessions	Release 24.4.1	You can achieve seamless connectivity for Customer Premise Equipment (CPEs) using Stateless Address Auto-Configuration (SLAAC) with PPPoE for IPv6 address assignment. This method enables CPEs to automatically configure IPv6 addresses without relying on a DHCP server. By leveraging SLAAC, devices can self-assign addresses based on the IPv6 prefix from the router, simplifying address configuration and reducing administrative overhead. Previously, PPPoE only supported DHCPv6 for IPv6 address assignment. The feature introduces these changes: **CLI:** • **slaac** **YANG Data Models:** • `Cisco-IOS-XR-ptp-cfg.yang` (see GitHub, YANG Data Models Navigator)

Key Concepts

- Point-to-Point Protocol over Ethernet (PPPoE): The PPPoE is a network protocol that encapsulates PPP frames inside Ethernet frames, enables multiple hosts on an Ethernet LAN to share a common internet connection while supporting IPv6 address configuration methods like DHCPv6 and SLAAC, and is often used by ISPs for authentication, session management, and scalable user connections. For more information on PPPoE, see PPP over Ethernet (PPPoE).
- Stateless Address Autoconfiguration (SLAAC): The SLAAC is an IPv6 stateless autoconfiguration mechanism that enables hosts to configure addresses autonomously without requiring manual intervention or additional servers, allows routers to broadcast prefixes that identify subnets, and permits hosts to create unique interface identifiers to form their addresses when combined with these prefixes.

SLAAC versus DHCPv6

SLAAC and DHCPv6 are two methods for assigning IPv6 IANA addresses to devices on a network. Although DHCPv6 support has been available for PPPoE, starting with Cisco IOS XR Release 24.4.1, we have now extended PPPoE support to include SLAAC.

Attributes	SLAAC (Stateless Address Autoconfiguration)	DHCPv6
Definition	Devices generate their own IPv6 addresses using a combination of locally available information and information advertised by routers. Routers send Router Advertisements (RAs) containing network prefix information.	A DHCPv6 server assigns IPv6 addresses and provides configuration information to devices on the network. It can also handle Prefix Delegation (PD) for distributing network prefixes.
Key Attributes	• No need for a dedicated server. • Minimal configuration required.	• Requires a dedicated DHCPv6 server. • Provides centralized control over IP address assignment.
Where used	• Small to medium-sized networks where ease of configuration is a priority. • Networks without a need for centralized control over IP address assignment.	• Large networks where centralized management of IP addresses is necessary. • Networks that need to use prefix delegation for hierarchical address distribution.

Configure SLAAC for PPPoE Subscriber Sessions

Configure SLAAC as the IPv6 address protocol with PPPoE to allow routers to generate their IPv6 addresses autonomously.

Procedure

Step 1 Enable SLAAC for PPPoE subscriber sessions on the access interface.

Example:

```
Router#configure
Router(config)#interface Bundle-Ether1.1
Router(config-subif)# service-policy output spd subscriber-parent resource-id 0
Router(config-subif)#ipv4 point-to-point
Router(config-subif)#ipv4 unnumbered Loopback1
Router(config-subif)#ipv6 enable
Router(config-subif)#encapsulation dot1q 1
Router(config-subif)# cnbng-nal ipv6 nd
Router(config-cnbng-nal)# ra-initial 0 16
```

```
Router(config-cnbng-nal-ra)# slaac
Router(config-cnbng-nal-ra)#exit
```

Step 2 Enable PPPoE on the access interface.

Example:

```
Router#configure
Router(config)#interface Bundle-Ether1.1
Router(config-subif)#pppoe enable
Router(config-subif)#commit
```

Step 3 Verify the configuration using the show run configuration.

Example:

```
Router#show run inter be1.1
interface Bundle-Ether1.1
 service-policy output spd subscriber-parent resource-id 0
 ipv4 point-to-point
 ipv4 unnumbered Loopback101
 ipv6 enable
 encapsulation dot1q 1
 cnbng-nal ipv6 nd
  ra-initial 0 16
  slaac
 !
 pppoe enable
!
```

Step 4 Once the subscribers are up, verify the IPv6 SLAAC prefix:

Example:

```
Router#show cnbng-nal subscriber all detail internal

=====================
Location: 0/RSP0/CPU0
=====================
Interface:                 Bundle-Ether1.1.pppoe2147483744
UPID:                      0x80000060
CPID:                      0x02239afc
Type:                      PPPoE
PPPOE Session Id:          00047194
PPP Params Info:
  Retry-count:             7
  Local-magic-number:      0xb6adb742
  peer-magic-number:       0xeac13140
  keep-alive-interval:     60
  MTU:                     0x000005dc
  Is-encap-string-ready:   TRUE
  Total KA Req Sent:       1550
  Total KA Resp Recv:      1550
  Total KA Req Recv:       0
  Total KA Resp Sent:      0
  PPP-flags:               0x00000000
IPv4 Address:              206.0.1.23
IPv4 Framed Route:
IPv6 IANA Address:         ::
IPv6 IAPD Prefix:          ::/0
IPv6 Slaac Prefix:         901:0:0:xxxx::/64
CPE link local Address:    ::
IPv6 Framed Route:
IPv4 State:                UP, Tue Oct 20 11:48:14 2024
IPv6 State:                UP, Tue Oct 20 11:48:14 2024
```

```
:
:
:
Attribute List: 0x55a6dfd0bc90
1:  ipv6-enable     len=  4  value= 1(1)
2:  nd-ra-initial   len=  3  value= 0.16
3:  nd-cnbng-ra-info len= 19  value= 901:0:0:xxxx::/64.1
4:  ip-vrf          len= 33  value= RJIL-VRF-OLT-MGMT
5:  inacl           len= 14  value= iACL_BNG_IPv4
6:  outacl          len= 14  value= iACL_BNG_IPv4
7:  ipv6_inacl      len= 14  value= iACL_BNG_IPv6
8:  ipv6_outacl     len= 14  value= iACL_BNG_IPv6
9:  strict-rpf      len=  4  value= 1(1)
10:  ipv6-strict-rpf len=  4  value= 1(1)
11:  ipv4-icmp-unreachable len=  4  value= 1(1)
12:  ipv6-unreachable len=  4  value= 1(1)
13:  ipv4-mtu        len=  4  value= 1492(5d4)
14:  ipv6-mtu        len=  4  value= 1492(5d4)
15:  ipv4-unnumbered len=  9  value= Loopback1
16:  sub-ipv4-gateway len= 12  value= 206.0.0.1/32
Last Transaction Result: SUCCESS
Session Accounting:        enabled
```

RADIUS-Based Policing - QoS shape-rate parameterization

RADIUS-Based Policing (RaBaPol) is a network management method that allows the activation of cnBNG subscriber services using customized parameters rather than default settings.

Table 14: Feature history

Feature Name	Release Information	Description
RADIUS-Based Policing - QoS shape-rate parameterization	Release 25.2.1	You can now dynamically manage your cnBNG subscriber services through RADIUS-based activation. With RADIUS-Based Policing (RaBaPol), you can customize service parameters, such as the QoS shape-rate, according to your requirements, giving you greater control over service management.

Parameterization of QoS shape-rate

RaBaPol supports the customization of the QoS shape-rate parameter. This parameter can be sent to the cnBNG Control Plane (CP) by the RADIUS server either during the initial connection setup as Cisco VSAs in an Access Accept message, or through Change of Authorization (CoA) messages.

Handling service changes and errors

If a service associated with a subscriber needs a change in the variable list, deactivate the current service using CoA Session-Disconnect and activate the updated service using CoA Session-Activate process. If an error occurs during feature activation, the cnBNG UP reverts all features and associated variable lists to their previous states.

Benefits of RADIUS-Based Policing

The RADIUS-Based Policing feature provides these benefits.

- **Dynamic activation**: Enables dynamic and flexible service activation based on RADIUS messages.
- **QoS customization**: Allows for the customization of QoS parameters to meet specific subscriber needs.
- **Policy merging**: Supports the merging of QoS policies from multiple dynamic templates for a subscriber.
- **Error rollback**: Provides rollback capabilities to previous states in case of errors during service activation.

Use case for QoS-based service activation

This use case illustrates how to manage and customize network QoS settings when a subscriber starts a session.

1. **Subscriber session initiation**: A user starts a session with specific credentials and settings, such as a username, password, and protocol type. For example,

```
user-cpe@abc.com         Password="abc"
            Framed-Protocol=PPP,
            Service-Type=Framed-User
              .....
               Cisco-avpair = "subscriber:sa=DEFAULT-QOS(shape-rate=120000)
```

2. **AAA server communication**: The Authentication, Authorization, and Accounting (AAA) server sends an Access-Accept message to the cnBNG. This message specifies the service name, action type, and a list of variables with their values, like the QoS shape-rate.
3. **Policy configuration**: The service name from the AAA message maps to a feature-template on the cnBNG's control plane, and the specified QoS shape-rate is used to override the default settings on the cnBNG's user plane. The policy merges these custom values with default values, retaining defaults where no specific values are provided.
4. **Service activation via CoA**: Alternatively, service activation can be achieved using CoA, which involves removing the old policy and configuring a new, merged policy in the hardware.

Limitations of configuring RADIUS-Based Policing

This limitation applies to the RADIUS-Based Policing feature:

- Service modifications with different RaBaPol configurations are not supported.

Configure QoS shape-rate parameterization

To establish QoS shape-rate parameterization, use the **shape average $var_name = value** command in the policy-map class configuration mode on the cnBNG User Plane (UP). This customization is feature-dependent and requires specific syntax and semantics. For QoS, a dollar sign ($) is added as a prefix to the **shape-rate** variable, and the default value, along with the variables, is configured in the policy-map definition.

Follow these steps to configure QoS shape-rate parameterization.

Procedure

Step 1 Define a feature template with the desired QoS configuration on the cnBNG CP.

Example:

```
config
 profile feature-template feature_template_name
  qos
      in-policy qos_input_policy_name
      out-policy qos_output_policy_name
      merge-level integer
   exit
 exit
```

This is a sample configuration.

```
config
  profile feature-template DEFAULT-QOS
    qos
     in-policy hqos-policy1
     out-policy hqos-policy2
     merge-level 10
   exit
  exit
```

Step 2 Configure the policy map with a shape-rate value, on the cnBNG UP.

Example:

```
config
 policy-map policy_map_name
    class class-default
       shape average $shape-rate =  rate (units)
    exit
    end-policy-map
 exit
```

This is a sample configuration.

```
config
  policy-map hqos-policy2
      class class-default
          shape average $shape-rate = 100000 kbps
      exit
      end-policy-map
  exit
```

This example enables QoS features for DEFAULT-QOS and configures the associated template with outgoing policies. The default value of shape-rate (the rate at which traffic is shaped) is set to 100000 kbps.

Step 3 Add the user profile to the USER file in the RADIUS server.

Example:

```
user-cpe@example.com         Password="abc"
              Framed-Protocol=PPP,
```

```
Service-Type=Framed-User
          .....
Cisco-avpair = "subscriber:sa=DEFAULT-QOS(shape-rate=120000)”
```

This specified QoS shape-rate value (for example, 120000) overrides the default value configured on the cnBNG UP.

Step 4 Use the **show subscriber session detail** command on the Control Plane to verify the configuration.

Example:

```
show subscriber session detail
```

```
subscriber-details
{
  "subResponses": [
    {
      "subLabel": "16777218",
      "mac": "cc11.0000.0001",
      "acct-sess-id": "01000002",
      "upf": "asr9k-1",
      "port-id": "Bundle-Ether1",
      "up-subs-id": "1",
      "sesstype": "ppp",
      "state": "established",
      "subCreateTime": "Fri, 15 Nov 2024 03:34:47 UTC",
      "pppAuditId": 3,
      "transId": "2",
      "subcfgInfo": {
      "activatedServices": [
          {
            "serviceName": "DEFAULT-QOS",
            "serviceAttrs": {
              "attrs": {
                "accounting-list": "automation-aaaprofile",
                "acct-interval": "900",
                "service-acct-enabled": "true",
                "service-parameters": "shape-rate=120000",
                "sub-qos-policy-in": "hqos-policy1",
                "sub-qos-policy-out": "hqos-policy2"
              }
} } ] } } ] }
```

Step 5 Use the **show policy-map applied interface** command on the User Plane to view sessions configured with RaBaPol.

Example:

```
bng# show policy-map applied interface Bundle-Ether1.1.pppoe100

Input policy-map applied to Bundle-Ether1.1.pppoe100:

  policy-map hqos-policy1
   class class-default
    police rate 200 kbps
    !
   !

Output policy-map applied to Bundle-Ether1.1.pppoe100:

  policy-map hqos-policy2
   class class-default
    shape average $shape-rate = 100000 kbps
   !
```

Shared Policy Instance

A shared policy instance (SPI) is a policy-driven QoS mechanism that

- enables the allocation of a single set of QoS resources to groups of BNG subscriber sessions
- allows these groups to share the allocated QoS resources collectively, and
- facilitates efficient resource management for multiple BNG subscriber sessions.

Table 15: Feature History Table

Feature Name	Release Information	Feature Description
Shared Policy Instance	Release 25.4.1	You can now streamline QoS policy management and ensure consistent rate enforcement across subscriber sessions with different port speeds by configuring shaping or policing rates as percentages in QoS policies for both ingress and egress directions. This enhancement replaces the previous requirement of using absolute values and allows a single policy map to be applied across multiple interfaces.
Shared Policy Instance	Release 25.2.1	You can now allocate and share a single set of QoS resources across multiple cnBNG sub-interfaces and bundle sub-interfaces. By using a single QoS policy instance across multiple sub-interfaces, you can enable aggregate shaping to one rate, promoting streamlined bandwidth management.

Efficient QoS policy sharing across sub-interfaces

SPI allows you to share a single QoS policy instance among multiple sub-interfaces to maintain a unified rate through aggregate shaping. Sub-interfaces sharing the QoS policy must belong to the same physical interface. The number of sub-interfaces can range from two to the maximum supported by the port.

Addressing challenges with absolute shaper values for SPI

Before Release 25.4.1, SPI configurations only supported absolute shaper or policer values. This limitation created challenges in environments with varying access interface speeds, such as 1G, 10G, and 100G. In these cases, SPI parameters were tied to the fixed capacity of the parent subscriber session.

Release 25.4.1 introduces the ability to configure shaping and policing rates as percentages within QoS policies for both ingress and egress directions, effectively resolving this limitation. For details on configuring SPI with percentage-based QoS allocation, refer to Enable percentage-based QoS allocation for SPI, on page 68.

Limitations of configuring Shared Policy Instance

Session consistency within S-VLAN interface

Sessions sharing the same SPI must remain within the same S-VLAN interface.

Session consistency within S-VLAN interface

Sessions sharing the same SPI must remain within the same S-VLAN interface.

Service accounting

Service accounting is not supported for services configured with an SPI.

SPI name change requirements

- If you modify the policy-map associated with an SPI, you must also change the SPI name.
- Avoid the following scenarios:
  - Applying a new policy with the same policy-map name but a different SPI name to a subscriber who already has an SPI policy applied. The system will reject this configuration.
  - Applying a new policy with a different policy-map name but the same SPI name. The system will reject this configuration as well.

Supported interfaces

- The SPI feature is supported only for bundle subscribers.

CoA service-update request limitation

When a service policy with a user profile configuration that includes an SPI is enabled, you cannot simultaneously use an SPI in a CoA service-update request.

Percentage-based QoS allocation for shared policy instances

- Configure percentage-based shaping and policing rates for QoS only if the subscriber has SPI enabled.

Configure Shared Policy Instance

To implement SPI, you must configure a complete hierarchical policy-map that includes both parent and child policies. The SPI name can be defined and linked to a feature template or downloaded from a RADIUS server.

There are two main ways to configure these policies:

- Using a feature template
- Using a RADIUS server

Configure a QoS policy with SPI using a feature template

Follow these steps to configure a QoS policy with shared policy instance in the input and output direction using a feature template.

Procedure

Step 1 Define a feature template on the Control Plane (CP) that includes the SPI configuration.

Example:

```
config
 profile feature-template feature_template_name
  qos
      in-policy qos_input_policy_name
        in-shared-policy-instance spi_name
      out-policy qos_output_policy_name
        out-shared-policy-instance spi_name
   exit
 exit
```

This is a sample configuration on the cnBNG CP.

```
config
   profile feature-template DEFAULT-QOS
      qos
         in-policy hqos-policy1
           in-shared-policy-instance spi1
         out-policy hqos-policy2
           out-shared-policy-instance spi2
     exit
  exit
```

Step 2 Configure traffic policing on the cnBNG UP to monitor the traffic rate and apply actions (such as dropping or remarking packets) when the traffic exceeds the allowed limit.

Example:

```
config
 policy-map policy_map_name
    class class-default
       police rate value
    exit
    end-policy-map
 exit
```

This is a sample configuration.

```
policy-map hqos-policy1
  class class-default
    police rate 1024 kbps
  exit
  end-policy-map
exit
```

Step 3 Configure traffic shaping for a specific interface on the cnBNG UP.

Example:

```
config
 policy-map policy_map_name
    class class-default
       shape average value
    exit
    end-policy-map
 exit
```

This is a sample configuration.

```
policy-map hqos-policy2
  class class-default
    shape average 4096 kbps
  exit
  end-policy-map
exit
```

Step 4 Use the **show subscriber session detail** command on the Control Plane to verify the configuration.

Example:

```
bng# show subscriber session detail
subscriber-details
{
  "subResponses": [
    {
      "subLabel": "16777220",
      "mac": "0011.9400.0001",
      "acct-sess-id": "01000004",
      "upf": "asr9k-1",
      "port-id": "Bundle-Ether1.1",
      "up-subs-id": "3",
      "sesstype": "ppp",
      "state": "established",
      "subCreateTime": "Fri, 15 Nov 2024 04:18:51 UTC",
      "pppAuditId": 3,
      "transId": "2",
      "subcfgInfo": {
        "committedAttrs": {
        "activatedServices": [
          {
            "serviceName": "DEFAULT-QOS",
            "serviceAttrs": {
              "attrs": {
                "accounting-list": "aaaprofile",
                "acct-interval": "900",
                "service-acct-enabled": "true",
                "sub-qos-policy-in": "hqos-policy1",
                "sub-qos-policy-out": "hqos-policy2",
                "sub-qos-spi-in": "spi1",
                "sub-qos-spi-out": "spi2"
              }
            } } ] } ] }
```

Configure a QoS policy with SPI using a RADIUS server

Follow these steps to configure a QoS policy with SPI using a RADIUS server.

Procedure

Step 1 Configure a policy map that can be shared to one or more interfaces to specify a service policy, on the cnBNG UP.

Example:

```
Router# config
Router(config)# policy-map hqos-policy1
Router(config-pmap)# class class-default
```

```
Router(config-pmap-c)# police rate 1024 kbps
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# exit

Router(config)# policy-map hqos-policy2
Router(config-pmap)# class class-default
Router(config-pmap-c)# shape average 4096 kbps
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# exit
```

This is a sample configuration.

```
config
  policy-map hqos-policy1
      class class-default
          police rate 1024 kbps
      !
      end-policy-map
  !
  policy-map hqos-policy2
      class class-default
          shape average 4096 kbps
      !
      end-policy-map
  !
```

Step 2 Add the QoS policy with the SPI name to the USER file in the RADIUS server.

Example:

```
abc@example.com Cleartext-Password:= "xyz"
cisco-avpair += "sub-qos-policy-in=hqos-policy1 shared-policy-instance spi1",
cisco-avpair += "sub-qos-policy-out=hqos-policy2 shared-policy-instance spi2",
```

Step 3 Use the **show subscriber session detail** command to verify the configuration of a subscriber with a user-profile that includes both QoS and SPI settings, on the cnBNG CP.

Example:

```
bng# show subscriber session detail
subscriber-details
      {
      "subResponses": [
      {
      "subLabel": "16777221",
      "mac": "cc11.0000.0001",
      "acct-sess-id": "01000005",
      "upf": "asr9k-1",
      "port-id": "Bundle-Ether1",
      "up-subs-id": "4",
      "sesstype": "ppp",
      "state": "established",
      "subCreateTime": "Fri, 15 Nov 2024 04:35:15 UTC",
      "pppAuditId": 3,
      "transId": "2",
      "subcfgInfo": {
        "committedAttrs": {
          "attrs": {
            "accounting-list": "aaaprofile",
            "acct-interval": "900",
            "addr-pool": "pool-ISP",
```

```
            "ppp-authentication": "pap,chap",
            "ppp-ipcp-reneg-ignore": "true",
            "ppp-ipv6cp-reneg-ignore": "true",
            "ppp-lcp-delay-seconds": "1",
            "ppp-lcp-reneg-ignore": "true",
            "service-type": "Framed(2)",
            "session-acct-enabled": "true",
            "sub-qos-policy-in": "hqos-policy1 shared-policy-instance spi1",
            "sub-qos-policy-out": "hqos-policy2 shared-policy-instance spi2",
            "vrf": "default"
          }
        } } } ] }
```

Step 4 Use the **show cnbng-nal subscriber all detail** command to display sessions with user-profile having QoS and SPI, on the cnBNG UP.

Example:

```
show cnbng-nal subscriber all detail
Interface:                  Bundle-Ether1.1.pppoe4
UPID:                       0x00000004
CPID:                       0x01000005
Type:                       PPPoE
PPPOE Session Id:           00000006

Attribute List: 0x175d470
1:  ipv4-unnumbered len=  9  value= Loopback0
2:  sub-qos-policy-in len= 59  value= hqos-policy1 shared-policy-instance spi1
3:  sub-qos-policy-out len= 63  value= hqos-policy2 shared-policy-instance spi2
```

Enable percentage-based QoS allocation for SPI

Procedure

Step 1 Define a feature-template on cnBNG CP to apply QoS policies using the shared policy instance mechanism. Associate them with specific SPI groups

This step defines the QoS policies and associates them with specific SPI groups (GRP1 and GRP2). Specify shaping or policing rates as percentages in these policies.

Example:

```
cnbng-cp(config)#profile feature-template default-qos
cnbng-cp(config)#qos
cnbng-cp(config)#in-policy policer-policy-in shared-policy-instance GRP1
cnbng-cp(config)#out-policy shaper-policy-out shared-policy-instance GRP2
cnbng-cp(config)#exit
cnbng-cp#exit
```

Step 2 Configure percentage-based rates in QoS policies on cnBNG UP.

Example:

```
Router-cnBNG-up#configure
Router-cnBNG-up(config-pmap)#policy-map policer-policy-in
Router-cnBNG-up(config-pmap)#class class-default
```

```
Router-cnBNG-up(config-pmap-c)#police rate percent 20
Router-cnBNG-up(config-pmap-c-police)#exit
Router-cnBNG-up(config-pmap)#end-policy-map

Router-cnBNG-up#configure
Router-cnBNG-up(config)#policy-map shaper-policy-out
Router-cnBNG-up(config-pmap)#class class-default
Router-cnBNG-up(config-pmap-c)#shape average percent 30
Router-cnBNG-up(config-pmap-c)#exit
Router-cnBNG-up(config-pmap)#end-policy-map
```

Step 3 Verify the configured rates for each subscriber to ensure the policies are applied successfully.

Example:

```
Router-cnBNG-up#show cnbng-nal subscriber all


Tue Jun 17 05:04:03.274 UTC

Location: 0/RSP0/CPU0
Codes: CN - Connecting, CD - Connected, AC - Activated,
       ID - Idle, DN - Disconnecting, IN - Initializing
       UN - Unknown


CPID(hex)  Interface               State  Mac Address     Subscriber IP Addr / Prefix (Vrf) Ifhandle
-------------------------------------------------------------------------------------------------
1          BE1.1.pppoe2147531664   AC     1234.1234.aabb  100.0.0.1 (default) 0x2f060
Session-count: 1



Router-cnBNG-up#show qos-ea interface BE1.1.pppoe2147531664 input member tenGigE 0/1/0/0
Interface: TenGigE0_1_0_0 input policy: policer-policy-in
Total number of classes:   1
Total number of UBRL classes:  0
Total number of CAC classes:   0
------------------------------------------------------
Policy name: policer-policy-in
Hierarchical depth 1
Interface type unknown
Interface rate 10000000 kbps
Port Shaper rate 0 kbps
Interface handle 0x0002F060
ul_ifh 0x060000C0, ul_id  0x00000000
uidb index 0xFFBB
qos_ifh 0x10002000ffbb
Local port 0, NP 0
Policy map id 0x1020, format 8, uidb index 0xFFBB
------------------------------------------------------
 Index 0 Level 0 Class name class-default service_id 0x0 Policy name policer-policy-in
 Node flags: LEAF DEFAULT DEFAULT-ALL
 Stats flags: policer-policy-in type 1 Max category 0
 Node Config:
 Police Color aware 0 Type 1 CIR/CBS/PIR/PBS: 2000000kbps/25000000B/0kbps/0B
 Node Result: Class-based stats:Stat ID 0x00C68DCC
 Queue: N/A Stat ID(Commit/Excess/Drop): 0x00000000/0x00000000/0x00000000
 Police ID (Token/Conform/Exceed/Violate): 0x00200001/0x00C68DCC/0x00C68DCD/0x00C68DCE
------------------------------------------------------


Router-cnBNG-up#show qos-ea interface BE1.1.pppoe2147531664 output member tenGigE 0/1/0/0
Tue Jun 17 05:04:35.338 UTC
```

```
Interface: TenGigE0_1_0_0 output policy: shaper-policy-out
Total number of classes:   1
Total number of UBRL classes:  0
Total number of CAC classes:   0
------------------------------------------------------
Policy name: shaper-policy-out
Hierarchical depth 1
Interface type unknown
Interface rate 10000000 kbps
Port Shaper rate 0 kbps
Interface handle 0x0002F060
ul_ifh 0x060000C0, ul_id  0x00000000
uidb index 0xFFBB
qos_ifh 0x10802000ffbb
Local port 0, NP 0
Policy map id 0x1420, format 8, uidb index 0xFFBB
------------------------------------------------------
 Index 0 Level 0 Class name class-default service_id 0x0 Policy name shaper-policy-out
 Node flags: LEAF Q_LEAF DEFAULT DEFAULT-ALL
 Stats flags: Queuing enabled
 Node Config:
 Shape: CIR/CBS/PIR/PBS: 0kbps/37500000B/3000000kbps/37500000B
 WFQ: BW/Sum of BW/Excess ratio: 0kbps/0kbps/1
 Queue limit 37500000 Guarantee 0
 Node Result: Class-based stats:Stat ID 0x00C68DCF
 Queue: Q-ID 0x0005e012 Stat ID(Commit/Excess/Drop): 0x00165222/0x00000000/0x009E0848
------------------------------------------------------
```

Enhanced subscriber routing with Framed-Route Tag and Preference attributes

Framed-Route is a routing attribute that

- allows you to define specific IP prefixes and next hops for subscriber sessions
- enables the use of additional attributes such as administrative distance (preference) and tag values, and
- supports both IPv4 and IPv6 address families for flexible routing.

Table 16: Feature History

Feature Name	Release Information	Description
Enhanced subscriber routing with Framed-Route Tag and Preference attributes	Release 25.4.1	You can now get automatic route prioritization and failover to a backup path, ensuring reliable subscriber connectivity. This is achieved by using RADIUS attributes to assign static IPv4 or IPv6 routes with preference and tag values.

Supported Framed-Route formats

cnBNG supports several variations of the Framed-Route and Framed-IPv6-Route attributes, including the use of tag and preference (admin distance) fields.

Format example	Description
`192.168.100.0/24 0.0.0.0`	Basic static route: Specifies a destination subnet and a next-hop IP address.
`192.168.100.0/24 0.0.0.0 10`	Static route with metric: Adds a route metric for path selection within routing protocols.
`192.168.100.0/24 0.0.0.0 10 tag 123`	Route with metric and tag: Includes a user-defined tag value for policy-based routing decisions.
`192.168.100.0/24 0.0.0.0 10 tag 123 pref 100`	Route with metric, tag, and preference: Specifies a metric, a tag, and an administrative distance (preference) to indicate the route's priority.
`192.168.100.0/24 0.0.0.0 10 pref 100`	Route with metric and preference: Specifies a metric and administrative distance for prioritizing routes.
`192.168.100.0/24 0.0.0.0 pref 100`	Route with preference only: Sets an administrative distance to determine route selection priority.

The same flexibility applies for Framed-IPv6-Route attributes.

Processing multi-homed subscriber routing scenarios

Multi-homed subscriber routing ensures high availability and uninterrupted connectivity by leveraging multiple access links for a single subscriber.

Summary

The key components involved in this process are:

- Subscriber device: A device that connects to the broadband network using multiple access links (for example, fiber for primary and wireless for backup).
- cnBNG: A network gateway that anchors the subscriber device and advertises subscriber routes to the core network with distinct administrative distance (preference) values for different links.
- Core network: A network infrastructure that selects the optimal path for subscriber traffic based on the administrative distance or preference values advertised by BNGs.

Workflow

These stages describe the process:

1. The subscriber device establishes connections with two different BNGs, each associated with a separate access link.
2. Under normal circumstances, the primary link (such as fiber) is used for subscriber traffic, as indicated by a lower administrative distance (higher preference).
3. Both BNGs advertise the subscriber's route to the core network, with their respective preference values.
4. The core network uses the administrative distance to select the preferred route for downstream traffic, favoring the primary link when available.

5. If the primary link fails, the backup link (such as mobile broadband) automatically becomes active, ensuring continued connectivity for the subscriber.
6. The administrative distance or preference is dynamically set using RADIUS attributes provided by the AAA server, enabling flexible path prioritization.

Backward compatibility considerations

This table outlines how different combinations of CP and UP software versions affect Framed-Route processing and subscriber session behavior.

CP version	UP version	Framed-Route attributes	CP behavior	UP behavior
New (2025.03.0 or newer)	Old (25.3.x or lower)	With `pref` and `tag`	CP encodes preference (`pref`) and tag into the framed route attributes as configured.	UP cannot decode the `pref` attribute and ignores it; installs route with default metric (admin distance 1).
New (2025.03.0 or newer)	Old (25.3.x or lower)	Without `pref`/`tag`	CP processes and forwards routes normally without the extra attributes.	UP installs and uses the framed route as expected.
Old (2025.02.0 or lower)	New (25.4.x or newer)	With `pref` and `tag`	CP cannot process unknown `pref` or `tag` attributes; drops the configuration and does not establish session.	UP receives no route installation request because session is not established.
Old (2025.02.0 or lower)	New (25.4.x or newer)	Without `pref`/`tag`	CP processes and forwards routes normally.	UP installs and uses the framed route as expected.

Restrictions for configuring Framed-Routes tag and preference attributes

Follow these requirements when configuring Framed-Route preference and tag attributes on cnBNG:

- Limit each subscriber to a maximum of 4 IPv4 framed routes and 4 IPv6 framed routes per address family.
- Use only the supported attribute formats. Unsupported or invalid formats may cause route installation to fail.
- If you specify the gateway as `0.0.0.0` (IPv4) or `::` (IPv6), the system uses the subscriber's IP address as the next-hop gateway by default.
- If you omit the prefix length in the framed route, the cnBNG infers the length based on the IP class (Class A: /8, Class B: /16, Class C: /24).
- Ensure only the active UP installs routes. The standby UP holds the information and installs routes only after a switchover.
- Confirm that both the CP and UP are running compatible software versions that support framed route preference or tag attributes.
- Do not use Framed-Route attributes in RADIUS accounting messages on cnBNG, as they are not supported.

- When planning mixed deployments, always verify feature support and limits between different BNG platforms such as physical BNG or cnBNG, as they may vary.

Configure Framed-Route preference support

Enable and customize Framed-Route administrative distance (preference) handling on cnBNG UP.

By default, cnBNG uses the **pref** value as the administrative distance for framed routes. However, if you apply this configuration on the UP, the metric value received in the framed-route will be used as the administrative distance instead.

Procedure

Step 1 Use the **framed-route-metric-as-distance** command on the UP to enable the cnBNG use the metric value as the administrative distance.

Important
Apply this configuration when there are no subscriber sessions present on the router.

Example:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# framed-route-metric-as-distance
Router(config-cnbng-nal-local)# exit
```

You do not need to execute any commands on the CP to configure this feature.

Step 2 Use the **show subscriber session detail command** command on the CP to view the framed route.

Example:

```
bng# show subscriber session detail | more
Thu Jun  12 04:00:23.852 UTC+00:00
subscriber-details
{
  "subResponses": [
      <snip>
      "subcfgInfo": {
        "committedAttrs": {
          "attrs": {
            "accounting-list": "automation-aaaprofile",
            "acct-interval": "2000",
            "addr-pool": "automation-poolv4",
            "ipv4-mtu": "1400",
           "ipv6-route": "vrf prefix-vrf 2001:db8:1::/64 vrf gw-vrf 2001:db8:100::1 100 tag 101 pref
 200,2001:db8:2::/64 2001:db8:100::2 101 tag 12 pref 200,2001:db8:3::/64 0:0:0:0:0:0:0:0 pref 300
tag 444 199,2001:db8:4::/64 :: 103 tag 12 pref 400",
            "ppp-ipcp-reneg-ignore": "true",
            "ppp-ipv6cp-reneg-ignore": "true",
            "ppp-lcp-reneg-ignore": "true",
            "route": "vrf prefix-vrf 192.168.1.0/24 vrf gw-vrf 192.168.1.1 100 tag 101 pref
200,192.168.3.0/24 192.168.3.100 101 pref 200 tag 102,172.10.1.0 172.10.1.1 pref 200 tag 102
405,10.10.1.0 10.10.1.1 tag 102 555 pref 399",
            "session-acct-enabled": "true",
            "vrf": "automation-vrf"
          }
```

```
    },
    "activatedServices": [
      <snip>
    ]
  },
  <snip>
  "v4FramedRoute": [
    "vrf prefix-vrf 192.168.1.0/24 vrf gw-vrf 192.168.1.1 100 tag 101 pref 200",
    "192.168.3.0/24 192.168.3.100 101 pref 200 tag 102",
    "172.10.1.0 172.10.1.1 pref 200 tag 102 405",
    "10.10.1.0 10.10.1.1 tag 102 555 pref 399"
  ],
  "v6FramedRoute": [
    "vrf prefix-vrf 2001:db8:1::/64 vrf gw-vrf 2001:db8:100::1 100 tag 101 pref 200",
    "2001:db8:2::/64 2001:db8:100::2 101 tag 12 pref 200",
    "2001:db8:3::/64 0:0:0:0:0:0:0:0 pref 300 tag 444 199",
    "2001:db8:4::/64 :: 103 tag 12 pref 400"
  ],
  <snip>
```

Step 3 Use the **show cnbng-nal subscriber all detail** command on the UP to view the tag and admin distance in a framed-route session.

Example:

```
Router# show cnbng-nal subscriber all detail
IPv4 Framed Route:
  Prefix:                 192.168.1.0/24
  Next Hop:               192.168.1.1
  Tag:                    101
  Metric:                 100
  Distance:               200
  Next Hop VRF:           abc
  Prefix:                 192.168.2.0/24
  Next Hop:               192.168.2.1
  Tag:                    101
  Metric:                 100
  Distance:               200
  Next Hop VRF:           abc
IPv6 IANA Address:        2001:DB8::7002
IPv6 IAPD Prefix:         ::/0
IPv6 Slaac Prefix:        ::/0
CPE link local Address:   ::
IPv6 Framed Route:
  Prefix:                 2001:DB8:4004:800::/64
  Next Hop:               ::
  Tag:                    101
  Metric:                 120
  Distance:               71
  Next Hop VRF:           abc
  Prefix:                 2001:DB8:4700::/64
  Next Hop:               2001:DB8:35::35
  Tag:                    23
  Metric:                 99
  Distance:               0
  Next Hop VRF:           abc
```

References

Framed-Routes align with these industry standards:

- RFC 2865 (RADIUS – Framed-Route attribute for IPv4)
- RFC 3162 (RADIUS – Framed-IPv6-Route attribute for IPv6)
- Broadband Forum TR-459 (BNG requirements and interoperability)

Geographical redundancy for L2TP sessions

Geographical redundancy for L2TP sessions is a redundancy framework for service providers that

- extends existing Control Plane geographical redundancy (CP-GR) to L2TP tunnels and sessions
- helps maintain session continuity and service availability during Control Plane (CP) failover events, and
- enables seamless failover and recovery of L2TP tunnels and sessions without manual intervention.

This feature builds on existing redundancy capabilities for IPoE and PPPoE, now including L2TP tunnels and sessions. For more information, see the *CP Geographical Redundancy* chapter of *Cloud Native BNG Control Plane Configuration Guide*.

Table 17: Feature history table

Feature Name	Release Information	Description
Geographical redundancy for L2TP sessions	Release 25.4.1	You can now ensure continuous service and session availability for L2TP users during cnBNG CP failovers, minimizing disruption for end users. When a failover occurs, the CP automatically synchronizes critical L2TP tunnel and session state with the cnBNG User Plane (UP), keeping tunnels and sessions operational without manual intervention.

Limitations of configuring geographical redundancy for L2TP sessions

- Do not reload or upgrade the cnBNG UP while L2TP tunnels and sessions are active. Reloading or upgrading will disconnect all active tunnels and sessions.
- Use CP-GR for L2TP only on supported CP and UP variants.
- Do not use Subscriber Redundancy Group (SRG) for L2TP LAC sessions, because it is not supported.
- Do not use OpenShift deployment, because it is currently not supported.

Configure CP-GR support for L2TP tunnels and sessions

Enable GR for L2TP tunnels and sessions on the cnBNG UP to support failover protection across remote sites.

Procedure

Step 1 Enable GR for L2TP tunnels and sessions on the specific UP.

Example:

```
Router# configure
Router(config)# cnbng-nal location 0/RSP0/CPU0
Router(config-cnbng-nal-local)# l2tp track-ns-nr
```

Step 2 Use the **show cnbng-nal l2tp-tunnel** command on the UP to display tunnel context details for a specific tunnel or all tunnels.

Example:

```
Router# show cnbng-nal l2tp-tunnel

Tunnel ID              NS              NR
=========             ====            ====
11                    500             501
12                    450             451
```

Step 3 Use the **show l2tp-tunnel detail instance-id** command on the CP to view the configuration and status of a specific L2TP tunnel.

Example:

```
bng# show l2tp-tunnel detail instance-id 1
Wed Jul  2  11:27:52.590 UTC+00:00
tunnel-details
{
  "tunResponses": [
    {
      "state": "established",
      "profileName": "lns-prof1",
      "tunnelType": "lns",
      "sessionCount": 1,
      "IDs Allocated": 1,
      "routerID": "asr9k-lns",
      "srcIP": "41.41.41.1",
      "dstIP": "91.91.91.1",
      "localTunnelID": 39832,
      "remoteTunnelID": 6410,
      "tunnelClientAuthID": "Local-DC",
      "tunnelServerAuthID": "bng-lns"
    }
  ]
}
```

Step 4 Use the **show l2tp-tunnel count instance-id** command on the CP to view the total number of L2TP tunnels.

Example:

```
bng# show l2tp-tunnel count instance-id 1
Thu Jul  3  15:55:02.172 UTC+00:00
tunnel-details
{
  "tunnelCount": 8996
}
```

Step 5 Use the **show subscriber lns count instance-id** command on the CP to view the number of LNS sessions.

Example:

```
bng# show subscriber lns count instance-id 1
Thu Jul  3  15:55:04.665 UTC+00:00
subscriber-details
{
  "sessionCount": 55964
}
```

To view the number of LAC sessions, use the **show subscriber pppoe count instance-id** command.

Step 6 Use the **show subscriber lns detail instance-id** command on the CP to view the detailed information about a specific LNS session.

Example:

```
bng# show subscriber lns detail instance-id 1
Mon Oct  20 06:30:27.240 UTC+00:00
subscriber-details
{
  "subResponses": [
    {
      "state": "complete",
      "key": {
        "routerID": "asr9k-1",
        "portID": "Bundle-Ether10",
        "subLabel": "16777218",
        "upSubID": "1"
      },
      "flags": [
        "SM_START_DONE",
        "SM_ACTIVATE_DONE",
        "SM_UPDATE_DONE",
        "IPCP_UP",
        "IPV6CP_UP"
      ],
      "lcpInfo": {
        "state": "opened",
        "keepAliveInterval": 60,
        "keepAliveRetries": 5,
        "localMru": 1492,
        "peerMru": 1492,
        "localMagic": "0xe2ab84f",
        "peerMagic": "0xe2ab850",
        "authOption": "PAP",
        "authCompleted": true,
        "username": "cnbng"
      },
      "ipcpInfo": {
        "state": "opened",
        "peerIpv4Pool": "pool-ISP",
        "peerIpv4Address": "11.0.32.2",
        "peerIpv4Netmask": 22,
        "localIpv4Address": "11.0.32.1",
        "isIpamPoolIPaddr": true
      },
      "ipv6cpInfo": {
        "state": "opened",
        "localIntfID": "0x1",
        "peerIntfID": "0xcc11000000010001"
      },
      "lnsInfo": {
        "srcIP": "10.1.39.139",
        "dstIP": "10.1.34.52",
        "state": "established",
```

```
        "profileName": "l2tp-prof2",
        "tunnelClientAuthID": "bng-lac",
        "tunnelServerAuthID": "Local-DC",
        "callSerialNumber": 16777408,
        "localTunnelID": 12226,
        "localSessionID": 11428,
        "remoteTunnelID": 1017,
        "remoteSessionID": 24885
      },
      "sessionType": "lns",
      "vrf": "default",
      "AuditId": 4,
    }
  ]
}
```

To view details for a specific LAC session, use the **show subscriber pppoe detail instance-id** command.

Synchronize L2TP tunnels between CP and UP for CP-GR

Procedure

Step 1 Use the **subscriber tunnel-synchronize upf** command on the CP to manually trigger tunnel synchronization to UP after HA events.

Example:

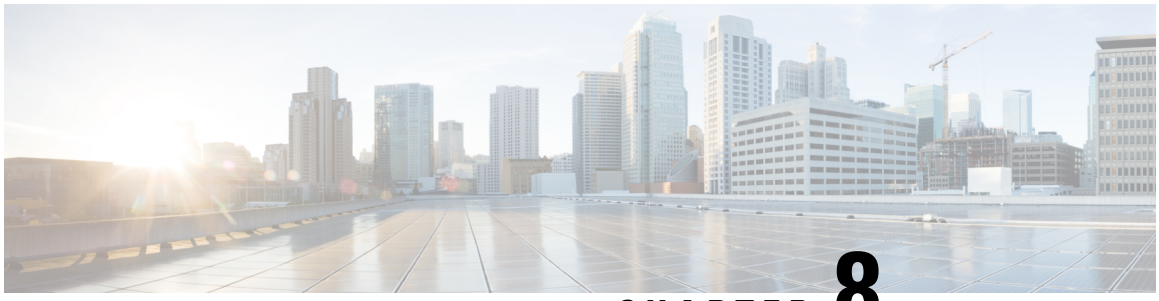
```
bng# subscriber tunnel-synchronize upf upf1
```

Step 2 Use the **show subscriber tunnel-synchronize upf** command on the CP to view the tunnel synchronization status.

Example:

```
bng# show subscriber tunnel-synchronize upf asr9k-lns

Thu Jul  3  04:00:32.799 UTC+00:00
subscriber-details
{
  "upf": "asr9k-lns",
  "total-tunnels": 10000,
  "passed-tunnels": 10000,
  "failed-tunnels": 0,
  "start-time": "2025-07-03 04:00:25.806",
  "end-time": "2025-07-03 04:00:26.098",
  "sync-time": "292 Milliseconds",
  "sync-status": "Sync Completed"
}
```



CHAPTER 8

Geo Redundancy (Subscriber Redundancy Group)

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

Table 18: 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	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. Subscriber Redundancy Group (SRG) provides flexible redundancy pairing on an access link by mirroring the subscriber session to a standby node. 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. This feature introduces these changes: **subscriber-redundancy**

- Overview, on page 128
- SRG Modes, on page 129
- Subscriber Session Set up Call Flow, on page 133
- Benefits of BNG Geo Redundancy, on page 135
- Supported Features in BNG Geo Redundancy, on page 136
- Unsupported Features and Restrictions for cnBNG Geo Redundancy, on page 136
- Guidelines to Configure SRG, on page 137
- Configure SRG, on page 138
- Routed subscriber sessions, on page 144
- SRG warm-standby mode , on page 152

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.

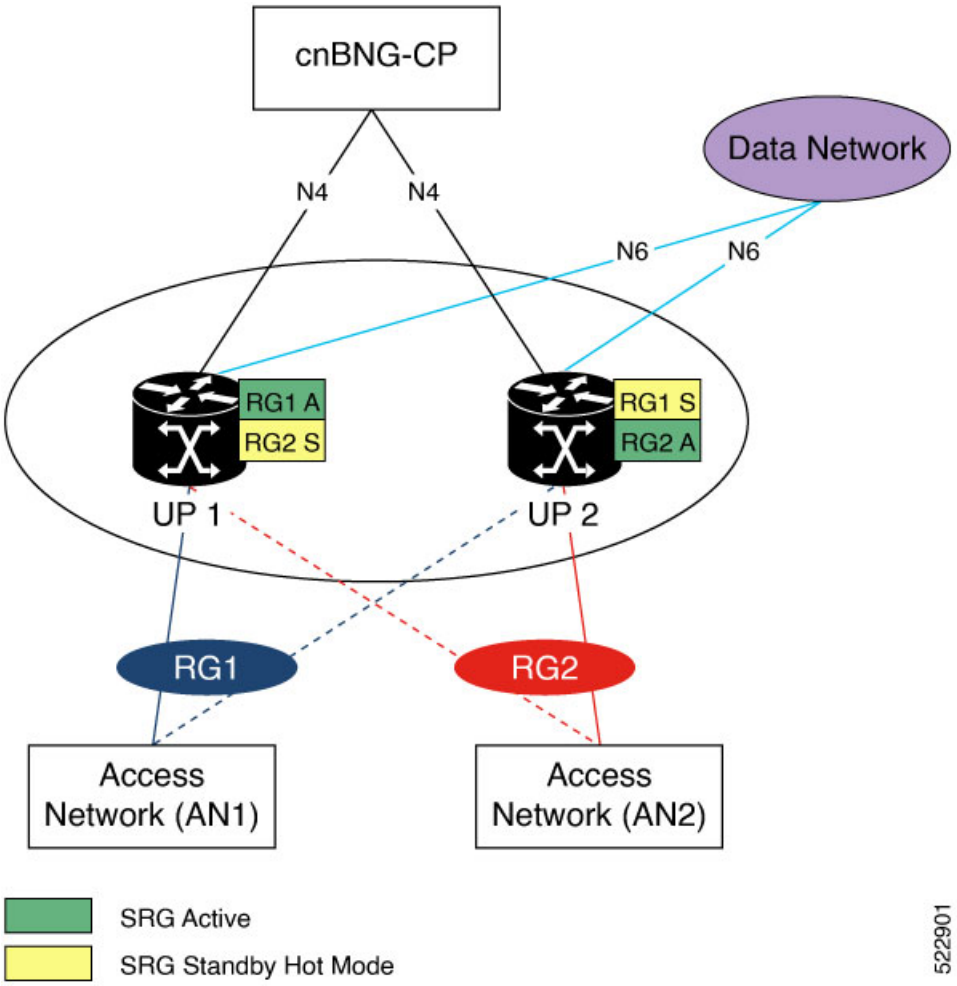
The Hot-standby mode supports 1:1 and M:N submodes.

1:1 Hot-standby

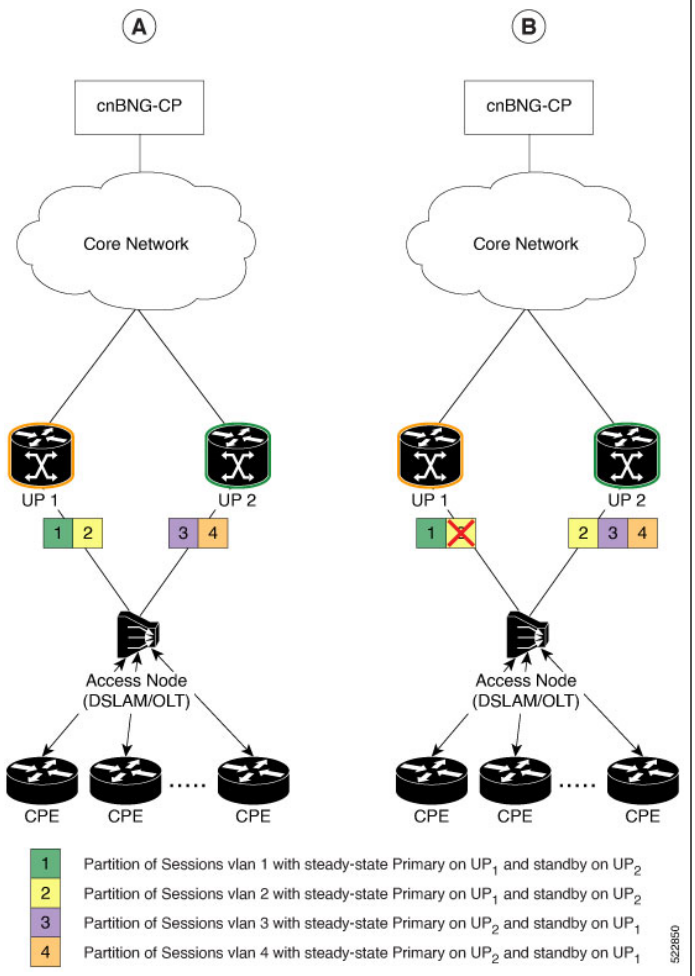
In this mode, 50 percent of the groups are in the active state and 50 percent of the groups are in the standby state.

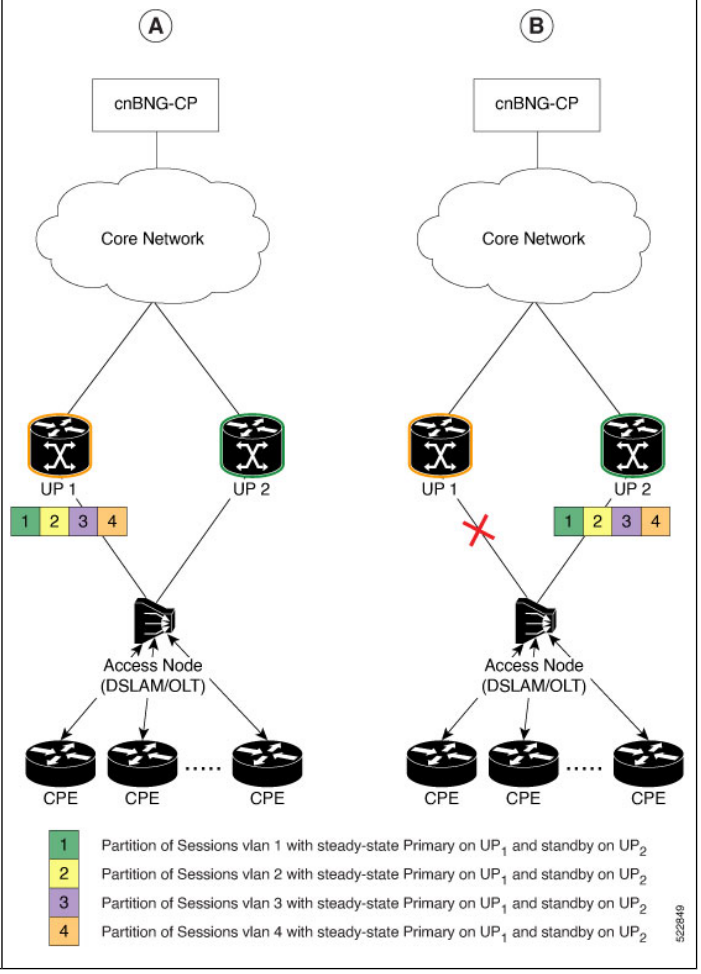
In this topology, access network AN1 is dual homed to UP1 and UP2. All subscribers from AN1 are grouped under the RG1 group. Access network AN2 is dual homed to UP1 and UP2 and all subscribers from AN2 are grouped under the RG2 group.

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.	

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.	

M:N Hot-standby

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

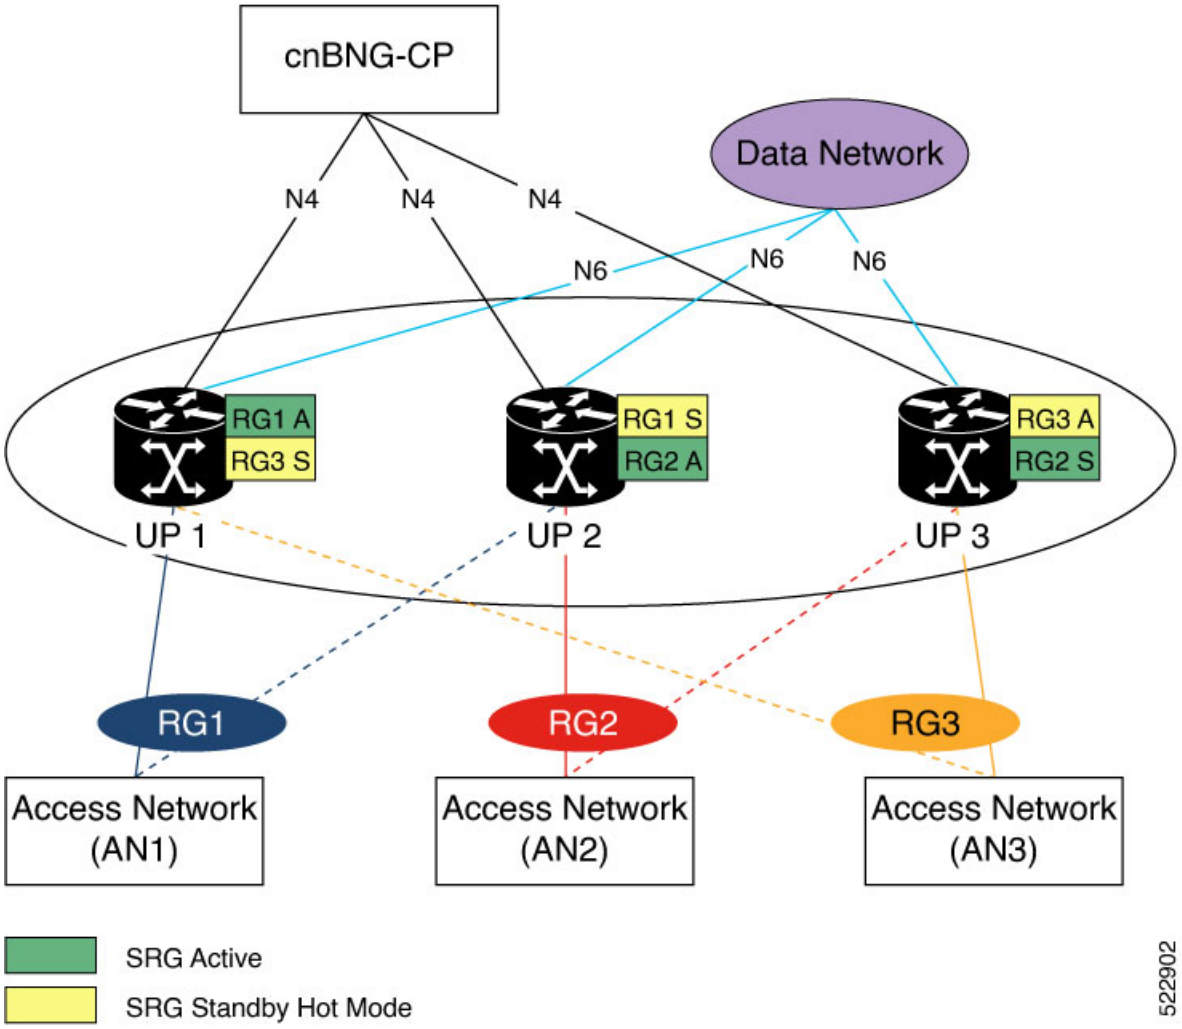
In this topology, access network AN1 is dual homed to UP1 and UP2. All subscribers from AN1 are grouped under the RG1 group. Access network AN2 is dual homed to UP1 and UP2 and all subscribers from AN2 are grouped under the RG2 group. Access network AN3 is dual homed to UP1 and UP3 and all subscribers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



Subscriber Session Set up Call Flow

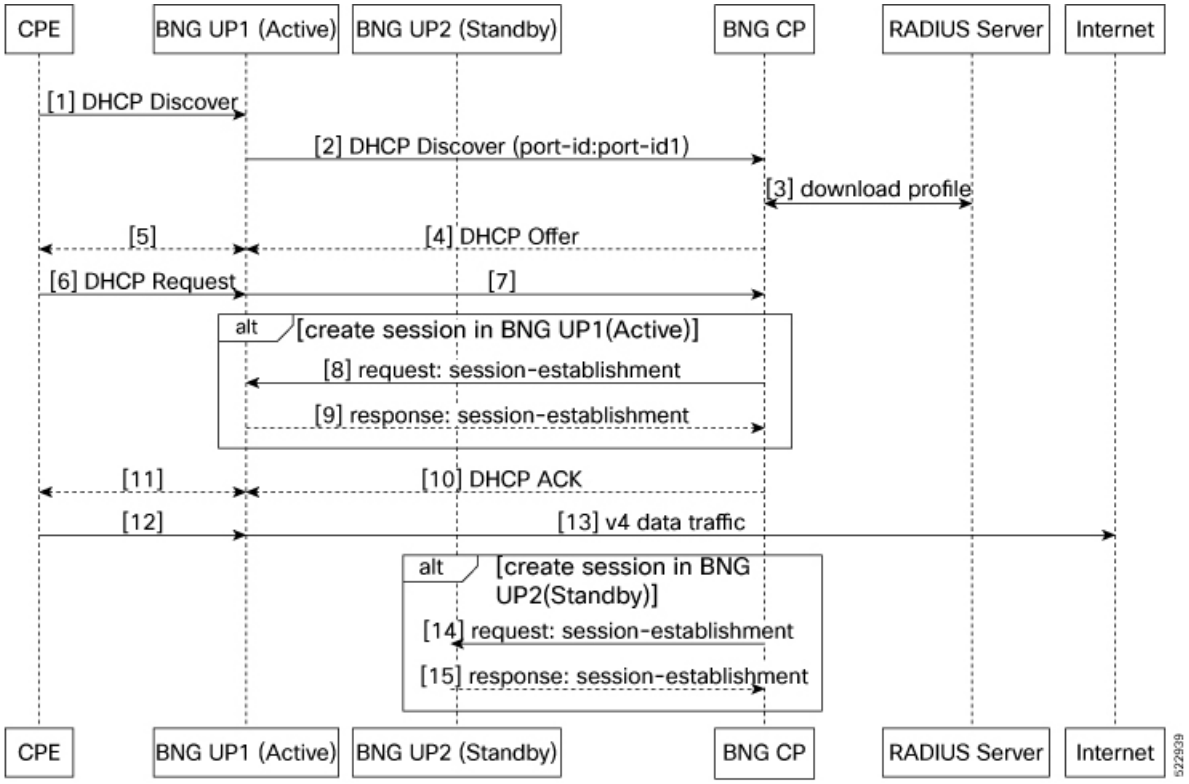
The following section graphs out the call flow and messaging between cnBNG SRG devices and the session.

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 12: Subscriber Session Creation for SRG

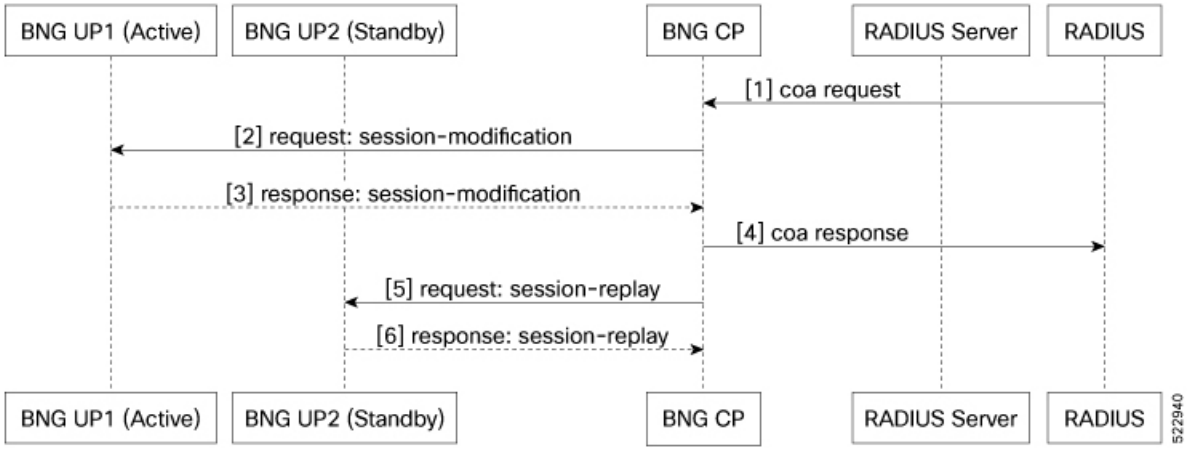


Subscriber Session attributes Modification Call flow for SRG

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:

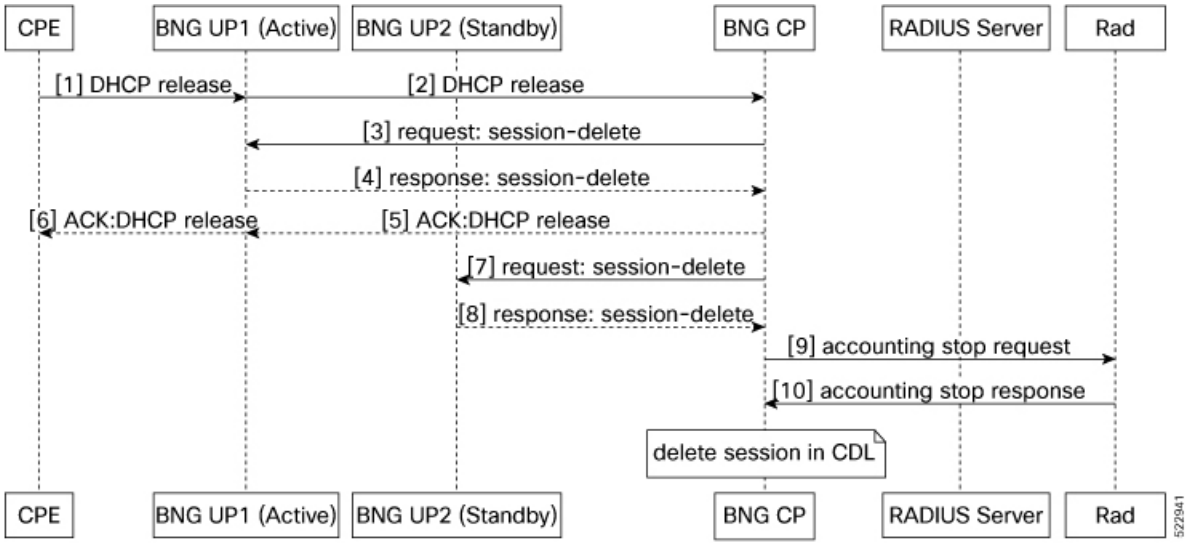
1. When there's a COA request from RADIUS, cnBNG CP triggers UP1 to modify the session attributes.
2. After the subscribers session attributes are modified, UP1 sends a response back to cNBNG CP with the subscriber session attribute details.
3. cnBNG CP now sends the COA response back to Radius and also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

Figure 13: Subscriber Session attributes Modification for SRG



Subscriber Session Deletion Call flow for SRG

Figure 14: 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 for IPoE and PPPoE 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 with out SRG is not supported with PPPoE-PTA, LAC, IPoE Dual stack sessions

- ARP scale mode

Guidelines to Configure SRG

- Atleast 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.

Fast switchover is a mechanism that minimizes traffic disruption during SRG role transitions. It is enabled by default and is typically used with SRG HOT mode to preserve subscriber forwarding continuity during standby to active role change.

When fast switchover is enabled, the system creates a Backup VRF (`srg_<group_id>`, where the group ID is internally generated) for each SRG group. This backup VRF is used to temporarily forward traffic during a role change.

The **fast-switchover-disable** option is not recommended by default. You can enable **fast-switchover-disable** only when traffic convergence is not a critical requirement.

This table summarizes subscriber route installation behavior with fast switchover enabled and disabled:

Configuration	Route Type	Main Table (Active)	Main Table (Standby)	Backup Table (Active)	Backup Table (Standby)
Enabled (Default)	Subscriber Routes	INSTALLED	NOT INSTALLED	INSTALLED	INSTALLED
Disabled	Subscriber Routes	INSTALLED	INSTALLED	NA	NA
N/A (Independent)	Framed/Subnet Routes	INSTALLED	NOT INSTALLED	NA	NA

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 0xxx.xxxx.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-cnbng-nal-sub-red)#group group2
Router(config-cnbng-nal-srg-grp)#virtual-mac 0xxx.0xxx.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-cnbng-nal-srg-grp-red)#group group3
Router(config-cnbng-nal-srg-grp)#virtual-mac 0xxx.0xxx.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#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 pfcp-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 0xxx.0xxx.0c01
   core-tracking core1
   access-tracking track1
   access-interface-list
    interface Bundle-Ether1.1
   !

  group group2
   virtual-mac 0xxx.0xxx.0a02
   core-tracking core1
   access-tracking track1
   access-interface-list
    interface Bundle-Ether1.2
   !
```

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 group1 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                            : 0x00000fa1
UP Instance ID                    : 0
SRG VRF name                      : **srg_4001
SRG mode                          : HOT
Last SRG role update time         : Mar 24 15:36:43.405702
Virtual mac                       : 0xxx.0xxx.0c01
V4 Table Id                       : 0xe0000012
V6 Table Id                       : 0xe0800012
V4 Proto Id                       : 0x00000001
V6 Proto Id                       : 0x00000003
Subscriber count                  : 0
Warm Subscriber count             : 0
IPV4 route count                  : 5
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      : FALSE

Access tracking object
----------------------
    Object name                   : track1
    Tracking state                : Up
    Last tracking state update time : Mar 24 15:36:11.820461

Core tracking object
--------------------
    Object name                   : core1
    Tracking state                : Up
    Last tracking state update time : Mar 24 15:36:11.820429

Access Interfaces
-----------------
    Bundle-Ether1.1
    Bundle-Ether1.2

IM counters
------------
    Total errors                  : 0

RIB counters
------------
    Total entries                 : 10
    Pending                       : 0
    Total errors                  : 0

Terminate counters
------------------
    IM terminate                  : 0
    RIB terminate                 : 0
    Stats terminate               : 0

Flags
-----
```

```
Value: [0x0000007C]

 [VRF_CREATE | V4_RIB_CREATE | V6_RIB_CREATE |
  V4_RIB_COMPLETE | V6_RIB_COMPLETE]


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)   : 0.060 secs
    Last stats interaction time(S->A)   : 0.060 secs
    Max stats interaction time(A->S)    : 0.060 secs (NA)
    Max stats interaction time(S->A)    : 0.060 secs (NA)

Group role switchover RIB Info
------------------------------
    Last rib interaction time(A->S)     : 0.060 secs
    Last rib interaction time(S->A)     : 0.060 secs
    Max rib interaction time(A->S)      : 0.060 secs (NA)
    Max rib interaction time(S->A)      : 0.060 secs (NA)

    Active role change count            : 1
    Standby role change count           : 1

[SRG-GROUP-EVENTS]

| Event Name                              | Time Stamp              | User Data
| Group create                            | Mar 24 15:36:11.818142  | -- NA --
| V4 backup vrf create                    | Mar 24 15:36:11.818355  | -- NA --
| V6 backup vrf create                    | Mar 24 15:36:11.818386  | -- NA --
| V4 backup vrf create rsp                | Mar 24 15:36:12.824602  | -- NA --
| V4 backup rib tbl create                | Mar 24 15:36:12.824629  | -- NA --
| V6 backup vrf create rsp                | Mar 24 15:36:12.824633  | -- NA --
| V6 backup rib tbl create                | Mar 24 15:36:12.824643  | -- NA --
| V4 backup rib tbl create rsp            | Mar 24 15:36:12.829051  | -- NA --
| V6 backup rib tbl create rsp            | Mar 24 15:36:12.838257  | -- NA --
| Role standby start                      | Mar 24 15:36:28.245876  | -- NA --
| Role standby                            | Mar 24 15:36:28.245876  | -- NA --
| Role standby end                        | Mar 24 15:36:28.246024  | -- NA --
| CP action add                           | Mar 24 15:36:28.246025  | -- NA --
| Role active                             | Mar 24 15:36:43.405701  | -- NA --
| Role active start                       | Mar 24 15:36:43.405737  | -- NA --
| Sent srg role response to CP            | Mar 24 15:36:43.405751  |
in_prog:T,comp:F,role:act,state:up
| Role active end                         | Mar 24 15:36:43.414395  | -- NA --
| Notify: State Up                        | Mar 24 15:36:48.405904  |
in_prog:F,comp:T,role:act,state:up
```

```
| State change ack'ed                       | Mar 24 15:36:48.406016  | -- NA --
| Fast swo timer expiry, Role Active        | Mar 24 15:36:53.406337  | -- NA --

=========================================================================
```

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 unreachables 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 unreachables 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 unreachables disable
 ipv6 nd cache-limit 5000
 ipv6 nd other-config-flag
 ipv6 nd managed-config-flag
 ipv6 address 2007::c353/128
 ipv6 enable
 ipv6 unreachables 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 19: Feature History Table

Feature Name	Release Information	Feature Description
Routed subscriber sessions	Release 26.2.1	We have enhanced this feature to support IPv4, and dual-stack routed subscriber sessions, including on bundle-ether VLAN interfaces, which improves scalability and deployment flexibility in Layer 3 access networks. Previously, only IPv6 was supported.

Feature Name	Release Information	Feature Description
Routed subscriber sessions	Release 25.1.1	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. 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. This feature introduces these changes: • 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. Yang Data Models: • `Cisco-IOS-XR-subscriber-nal-cfg.yang` • `Cisco-IOS-XR-cnbng-nal-oper.yang` (see GitHub, YANG Data Models Navigator)

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

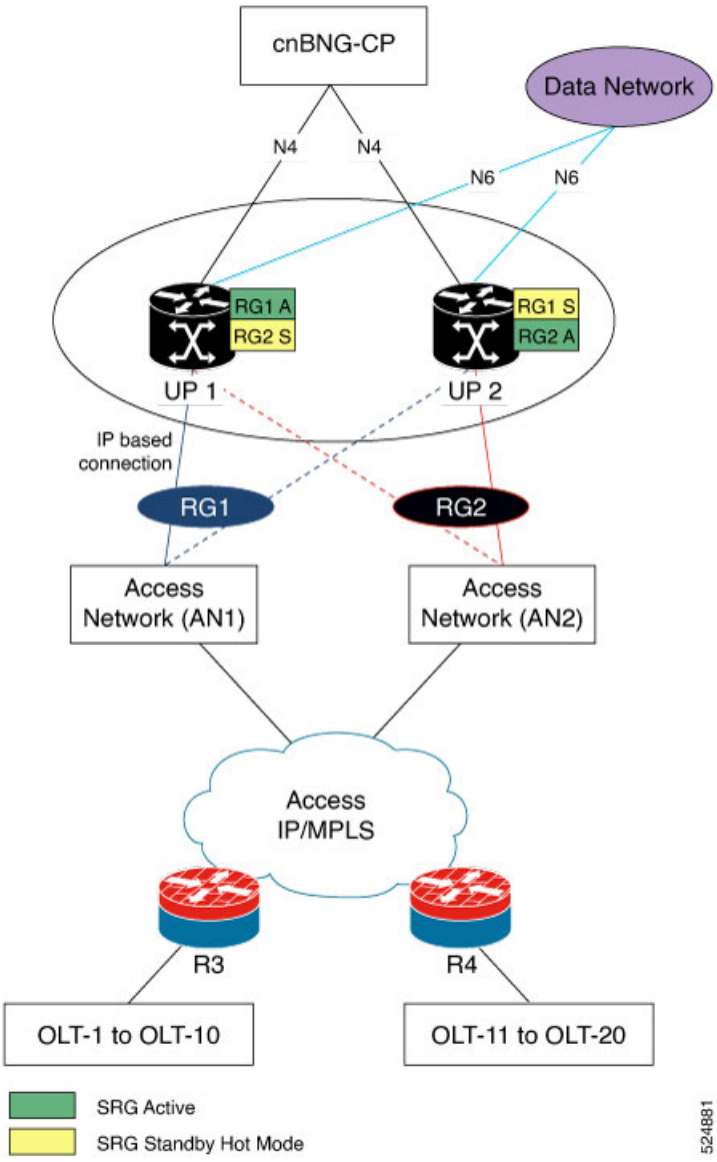
Routed subscriber sessions allow devices to connect to the cnBNG through routed devices. Instead of using MAC addresses, this session identification relies on IP addresses. The Customer Premises Equipment (CPE) uses Network Address Translation (NAT) to hide connected devices. The BNG sees a subscriber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

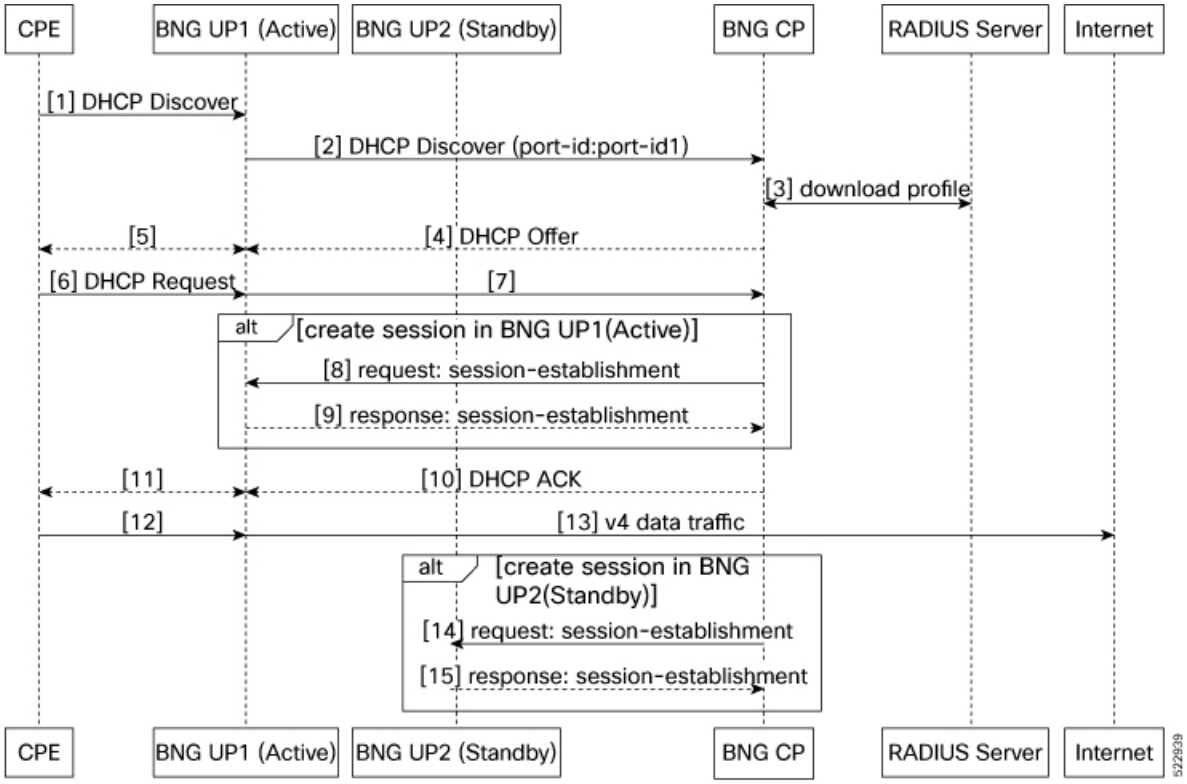
Figure 15: Topology for routed subscriber



Workflow

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

Figure 16: 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:

- For bandwidth capacity expansion, multiple interfaces would be needed in bundle interface and not as ECMP.
- 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.

Example:

```
Router#configure
Router(config)#interface bundle-ether 1
Router(config-subif)#cnbng-nal
Router(config-subif)#ipsubscriber
Router(config-cnbng-nal-ipsub)#ipv4 routed
Router(config-cnbng-nal-ipsub-ipv4-routed)#initiator dhcp
Router(config-cnbng-nal-ipsub-ipv4-routed)#exit
Router(config-cnbng-nal-ipsub)#exit
Router(config-subif)#ipsubscriber
Router(config-cnbng-nal-ipsub)#ipv6 routed
Router(config-cnbng-nal-ipsub-ipv6-routed)#initiator dhcp prefix-len 20 src-ip-dual-lookup
```

Step 2 Configure the access-control route and state-control next-hop IP address under SRG group.

The next-hop IP configured must be reachable by subscriber VRF.

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)#access-tracking track1
Router(config-cnbng-nal-srg-grp)#access-control-route ipv4 10.1.2.2/0 vrf test next-hop-address
10.1.0.3 active-tag 10 standby-tag 20
Router(config-cnbng-nal-srg-grp)#access-control-route ipv6 ::/0 vrf test next-hop-address 2001:db8:4
 active-tag 10 standby-tag 20
Router(config-cnbng-nal-srg-grp)#state-control-next-hop-ip ipv4 10.5.0.1
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:       10.0.2.1
  Routed IPv6 Prefix:       1::2
IPv4 Address:               10.1.2.2
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 20: 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
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

Aspect	Warm standby	Hot standby	Non-redundant deployment
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

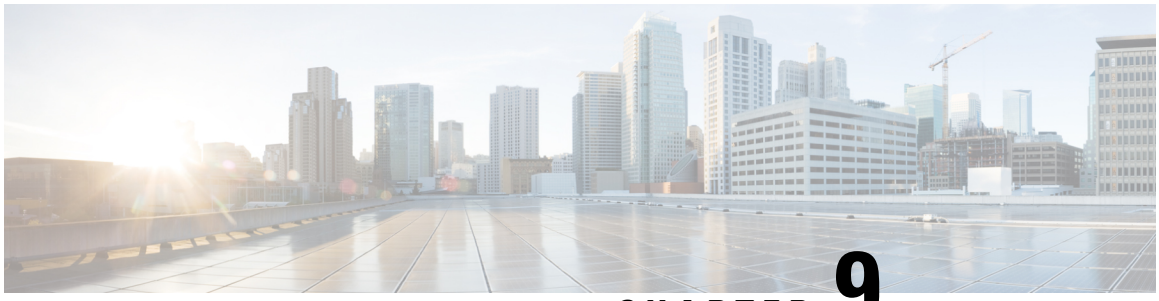
Total Entries : 3

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
====================
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)
```



CHAPTER 9

IPV6 Neighbor Discovery

This chapter provides information about the IPV6 Neighbor Discovery.

Table 21: Feature History Table

Feature Name	Release Information	Feature Description
IPv6 Neighbor Discovery	Release 7.10.1	You can now configure IPv6 Neighbor Discovery (ND) parameters on the access interface that facilitates address resolution, router discovery, and duplicate address detection. The IPv6 ND protocol discovers and establishes communication with neighboring IPv6 nodes within a local network. This feature introduces the following changes: • **CLI**: **cnbng ipv6 nd** commands • **YANG** `Cisco-IOS-XR-um-asr9k-cnbng-nal-cfg` (see GitHub, YANG Data Models Navigator):

• IPv6 Neighbor Discovery, on page 157

IPv6 Neighbor Discovery

IPv6 Neighbor Discovery (ND) is a protocol used to determine the link-layer addresses of neighboring nodes, such as customer routers to forward IPv6 traffic.

IPv6 ND enables cloud native BNG to act as a router and uses IPv6 ND to learn the link-layer addresses of customer routers connected to it. IPv6 ND maintains the information about other devices in the IPV6 network and tracks the presence of neighboring devices, and determines its reachability to those devices.

To communicate with neighboring nodes, IPv6 ND uses the following set of messages:

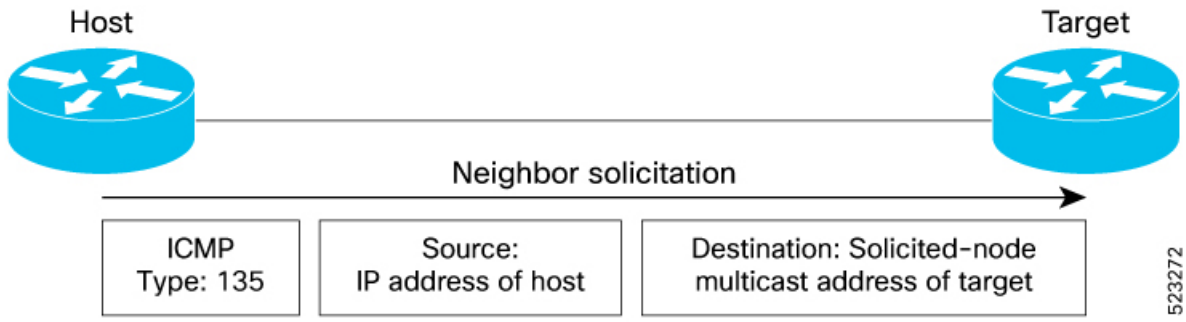
- IPv6 Router Solicitation Message (RS) is an IPv6 message sent by a host to request Router Advertisements (RA) from routers on the network. Router solicitation messages are sent on the local link when a host wants to determine the link-layer address of another node on the same local link. The ICMP packet header has a value of 135 in the Type field to identify the RS message.
- Router Advertisement (RA) is an IPv6 message sent periodically by routers or in response to an RS message sent by a host. The RA message advertises the presence of routers on the network to provide hosts with configuration information. RA contains prefixes used to determine whether another host shares the same link. The ICMP packet header has a value of 134 in the Type field to identify the RA message. An ambiguous VLAN does not have an association with any particular VLAN; therefore, a unicast router advertisement message is sent for ambiguous VLAN interfaces.

  To enable IPv6 unicast router advertisement, use the **cnBNG NAL ipv6 nd ra-unicast** command in the `cnbng-nal` configuration mode.
- Neighbor Solicitation (NS) is an IPv6 message sent by a node to determine the link-layer address of a neighbor, or to verify that a neighbor is still reachable through a cached link-layer address. NS messages also checks the Duplicate Address Detection (DAD) if the IPv6 address configured is already in-use by another node on the same link.
- Neighbor Advertisement (NA) is an IPv6 message sent in response to a Neighbor Solicitation (NS) message to notify its link-layer address neighbors. When a node receives an NS message, it responds with an NA message that includes its link-layer address. The NA message is sent to the source address of the NS message and can be either unicast or multicast.
- Redirect is an IPv6 message that routers use to notify hosts of an optimal first-hop router for a given destination. When a host sends a packet to a destination router, and the router receiving the packet determines that the next hop is not the best one, the router sends a Redirect message to the host. The Redirect message includes the IP address of the destination router and the IP address of the new next-hop router that the host must use instead. The ICMP packet header has a value of 137 in the Type field to identify the redirect message.

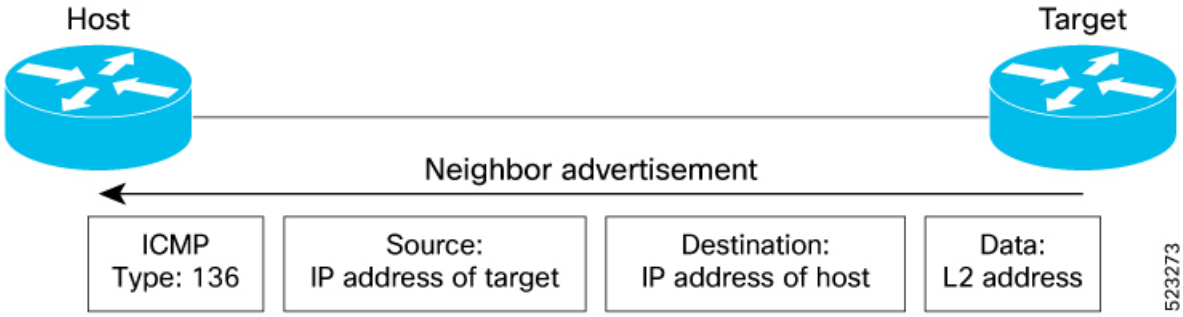
How IPv6 Neighbor Discovery Works?

In an IPv6 network, during the communication with a neighboring device:

- The host (P1) sends an Neighbor Solicitation (NS) message to the link-local multicast address requesting the MAC address of the destination router (P2) with the specified IPv6 address.



- The neighbor responds with a Router Advertisement (RA) message that includes its link-layer address. This exchange allows the routers to establish a mapping between the IPv6 address of the neighbors and link-layer address for communication to occur.



IPv6 neighbor discovery uses Duplicate Address Detection (DAD) to ensure that no two devices on the same network have the same IPv6 address. When a device joins the network or configures a new IPv6 address, the host sends a neighbor solicitation message with its IPv6 address as the target. If the host receives a neighbor advertisement message in response, the host knows that another device on the network, which is already using that IPv6 address, and the host must choose a different IPv6 address.

You have the flexibility to configure the IPv6 ND parameters such as the frequency of RA messages or the interval between NS and NA messages. according to your network design under the access interface.

Configure IPv6 Neighbor Discovery

Based on your requirements, configure the following IPv6 ND parameters:

```
Router#configure
Router(config)#interface Bundle-Ether1.1

/*Set the managed address configuration flag in IPv6 router advertisements*\
Router(config-subif)# cnbng-nal ipv6 nd managed-config-flag

/*Configure the interval between IPv6 neighbor solicitation retransmissions on an interface
 *\
Router(config-subif)# cnbng-nal ipv6 nd ns-interval 1999

/*Enable the IPv6 neighbor un-reachability detection (NUD) *\
Router(config-subif)#cnbng-nal ipv6 nd  nud-enable

/*Set the other stateful configuration flag in IPv6 router advertisements *\
Router(config-subif)# cnbng-nal ipv6 nd other-config-flag

/*Set the IPv6 initial router advertisement count and interval *\
Router(config-subif)# cnbng-nal ipv6 nd ra-initial 9 4

/*Configure the interval between IPv6 router advertisement transmissions on an interface
*\
Router(config-subif)# cnbng-nal ipv6 nd ra-interval 888000 8000

/* Configure the router lifetime value in IPv6 router advertisements on an interface*\
Router(config-subif)# cnbng-nal ipv6 nd ra-lifetime 777

/* Enable the IPv6 unicast router advertisement (RA)*\
Router(config-subif)# cnbng-nal ipv6 nd ra-unicast

/*Configure the amount of time that a remote IPv6 node is considered reachable after some
```

```
reachability confirmation event has occurred*\
Router(config-subif)# cnbng-nal ipv6 nd reachable-time 9000

/*Automatically send IPv6 router advertisements to a subscriber interface after configuring
 IPv6 *\
Router(config-subif)# cnbng-nal ipv6 nd start-ra-on-ipv6-enable

/* Suppress IPv6 router advertisement transmissions on a LAN interface *\
Router(config-subif)# cnbng-nal ipv6 nd suppress-ra

/* Suppress cache learning for IPv6 neighbor discovery*\
Router(config-subif)# cnbng-nal ipv6 nd suppress-cache-learning

/* Set the managed address configuration flag in IPv6 router advertisements*\
Router(config-subif)# cnbng-nal ipv6 nd managed-config-flag

/* configure the IPv6 ND router advertisement hop-limit on the VLAN*\
Router(config-subif)# cnbng-nal ipv6 nd ipv6 nd hop-limit unspecified

/* Set the IPv6 neighbor discovery router preference *\
Router(config-subif)# cnbng-nal ipv6 nd router-preference high

/* Suppress the MTU option in IPv6 Neighbor Discovery (ND) Router Advertisement (RA) header
 *\
Router(config-subif)# cnbng-nal ipv6 nd mtu suppress
```

Verification

Verify the configured parameters appear in the output.

```
Router#%show ipv6 nd idb interface Gi0/2/0/2.1.ip536870944 detail location 0/2/CPU0;

 ifname: Gi0/2/0/2.1.ip536870944, ifh: 0x1000100, iftype: 65, VI-type: 0, Pseudo IDB: FALSE

 vrf-id: 0x60000000, table-id: 0xe0800000
 Mac Addr: xxxx.xxx.xxx, size: 6, VLan tag set: FALSE

 Media Name: ipsub_base, Media Encap: 0xe (IPSUB)
 Mac Length: 1, Media Header Len: 4, Media Proto: 0xdd86
 Current  Encap: 0xe (IPSUB), Mcast  Encap : 0xe (IPSUB)

 IPV6 Interface: Enabled, IPV6: Enabled,  MPLS: Disabled
 Link local address: xxxx::xx:xxxx:xxxx:xxxx, Global Addr count: 0
 Default Prefix Address: ::, Prefix Addr Count: 0,

 RA Specific Route Count: 0,

 RA DNS Servers Addr Count: 0,

 RA DNS Search List Count: 0,

 DAD Attempts: 0, DAD pending 0,

 RA flag: 0x0, Unicast RA send: TRUE, Initial RA count: 9, RA pkts sent count: 0
 Initial RA interval: 999000 msec,
 Time of Last RA sent: N/A, Next Scheduled Periodic RA Time due in: N/A
 RA Managed flag 0x1, RA Other flag 0x1, RA Hop limit 0x1
 RA Suppress MTU: 0x1, RA Lifetime: 777 sec
 RA interval min: 888000, max 8000 msec
 RA Router Preference: Low

 Reachable time: 9000 msec, Reachable delay: 8000 msec
 RA retransmits: 1999 msec,  NS retransmits interval: 1999 msec
```

```
 AIB stats time interval: 1000 msec

 ND Redirects: 0x0, NUD Conform: 0, MTU: 1500, IDB Flags: 0x1024

 Cache entry limit: 1000000000, Last over limit count: 0
 Complete protocol count: 0, Complete glean count: 0
 Incomplete protocol count: 0, Incomplete glean count: 0
 Dropped protocol req count: 0, Dropped glean req count: 0

 IPC notification handle: 0, Config Flags: 0xcfc0fd, Parent if: GigabitEthernet0_2_0_2.1
(0x10000c0)
 Refresh from RP: FALSE,
 IM call for IDB: Success, Mac addr changed:  TRUE
 IM error recover retries count: 0
 Check point Obj ID: 0x2e80, Framed IPv6 prefix pool name:
 Subscriber status flag: 0x0, Supressed cache learning: TRUE
 BNG nud: Enabled, Master Node: (0xddddddddddddddd)
 Global Mac Accounting: Disabled, IDB Mac Accounting : Disabled, Marked: No
 Notfn sent to iedge - Up: No, Down: No
                       Update: No
                       Last notif reason:None
 SRG Stby Role : FALSE , SRG peer route-disable : FALSE, SRG EOMS sync pending : FALSE
 Subscriber Label : 0x0
 Prefix Address from Iedged: ::,
 Input Bytes: 0, Input Bytes Pkts: 0
 Output Bytes: 0, Output Bytes Pkts: 0
.........    IDB Statistics     .........
Service    Attribute   Operation  Success Failure Avg Min Max

IM          MAC         Reg            1     0    0    0    0
IM          MAC         Notfn          1     0    0    0    0
IM          MTU         Notfn          1     0    0    0    0
IM          MPLS        Notfn          1     0    0    0    0

NETIO       NA          Sent           1     0    0    0    0

IPV6-MA     IDB         Add            1     0    0    0    0

CHKPT       IDB         Add            5     0    0    0    0
------------------------------EVT-HISTORY-------------------------------------
    Nov 17 12:20:32.576 idb-calloc-happened
    Nov 17 12:20:32.576 idb-bng-srg-master-or-none
    Nov 17 12:20:32.576 idb-bng-srg-slave-route-enable
    Nov 17 12:20:32.576 idb-bng-srg-eoms-sync-not-pend
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-bng-subdb-strt-ra
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.576 idb-apply-func
    Nov 17 12:20:32.704 idb-im-create-notification
    Nov 17 12:20:32.704 idb-ma-state-enabled
    Nov 17 12:20:32.704 idb-chkpt-save [many]
```