

Cloud Native BNG User Plane Overview

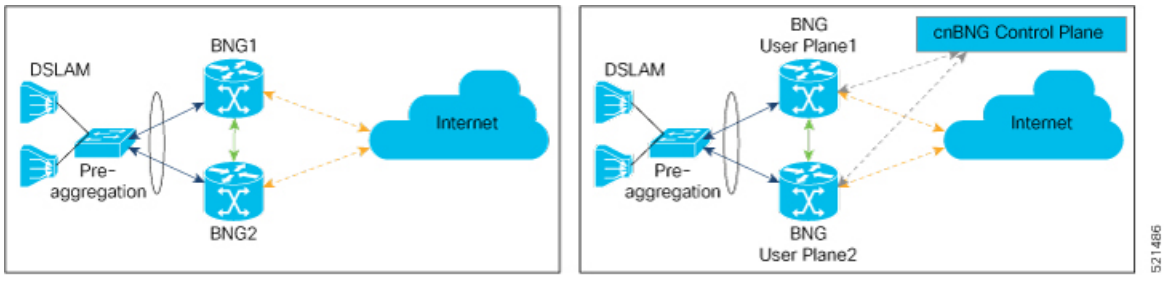
In the cnBNG architecture, which is based on Control and User Plane Separation (CUPS), the CP handles the subscriber management functionality and the UP handles the forwarding functionality of the entire BNG solution. This chapter focuses on the functionality and architecture of the cnBNG user plane.

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 1
- cnBNG User Plane Overview, on page 3
- cnBNG User Plane Architecture, on page 4
- cnBNG User Plane and Control Plane Reachability, on page 5
- Software and Hardware Requirements, on page 6
- Access Types and Subscriber Types, on page 7
- Subscriber Features, on page 8
- High Availability, on page 11
- Usage Guidelines, on page 12
- Restrictions, on page 12

Control and User Plane Separation

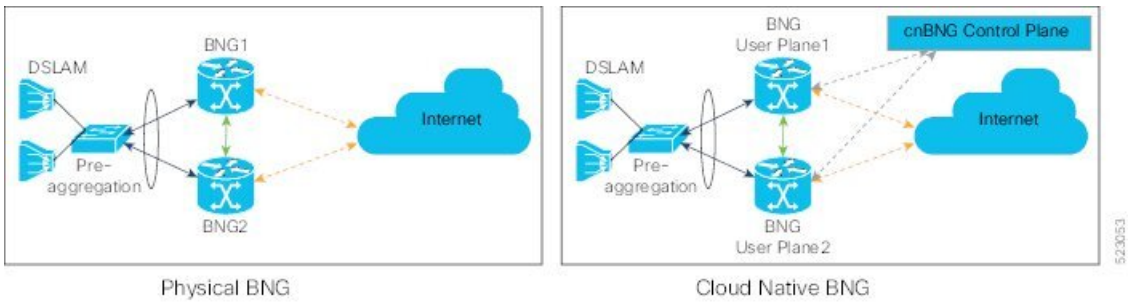
cnBNG is an architectural evolution that is based on Control and User Plane Separation (CUPS), where the CP and UP run in distinct and independent environments. cnBNG redefines the traditional physical BNG by decoupling the BNG CP and UP functions to give better flexibility and scalability for the service 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 1: 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.

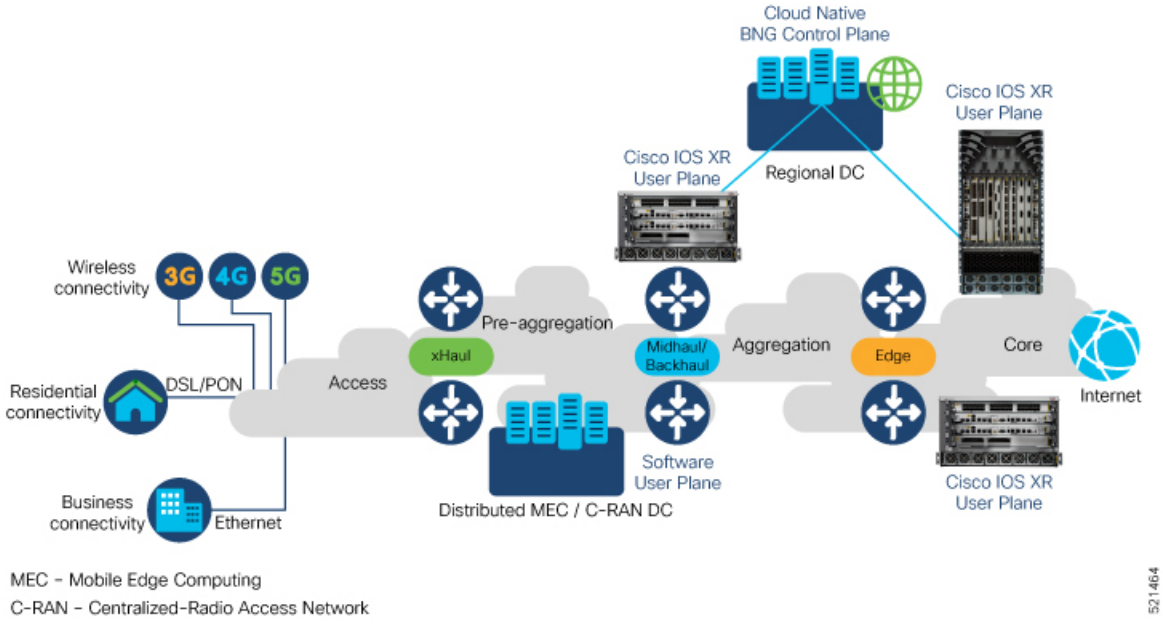
Why CUPS?

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 2: 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:

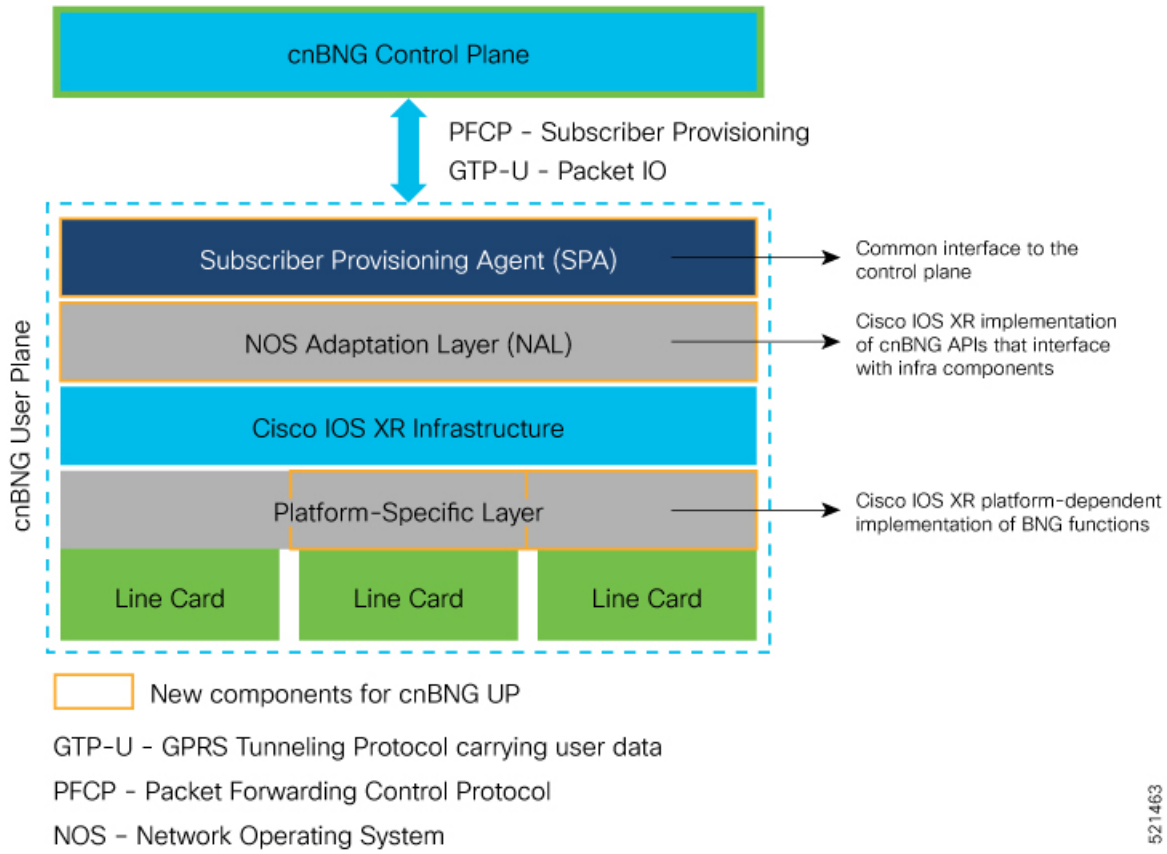
- **Distributed**: With reduced operational complexity and minimal integration efforts with centralized CP, you can distribute the UPs closer to end users. This feature helps to offload the traffic to the nearest peering points and content delivery networks (CDNs) and reduces the core transport costs.

- **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 3: 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