



CnBNG Control Plane Quality of Service

Quality of Service (QoS) is a technique for prioritizing network traffic. It ensures that time-sensitive and mission-critical applications receive sufficient bandwidth, low latency, and minimal data loss. QoS forwards high-priority packets by identifying, classifying, and prioritizing them throughout the network.

To achieve these QoS objectives, the cnBNG utilizes Differentiated Services Code Point (DSCP) marking as the primary mechanism for traffic prioritization.

Understanding DSCP marking

Differentiated Services Code Point (DSCP) marking is a field in the IP header used to classify and manage network traffic. By assigning a specific DSCP value to a packet, the cnBNG Control Plane (CP) tags the packet so routers along the data-forwarding path prioritize it. This ensures that traffic receives preferential treatment, such as lower latency and reduced packet loss, during periods of network congestion.

This chapter details the deployment of QoS for the cnBNG CP.

- DSCP marking for replication traffic (CP-CP): This feature prioritizes inter-instance state-synchronization data. By applying consistent DSCP values, the cnBNG CP prevents synchronization delays between redundant CP clusters. This process maintains geo-redundancy and failover readiness.
 - DSCP marking for CP-UP control packets (CP-UP): This feature provides granular, message-level classification for PFCP and GTP-U signaling. It allows the cnBNG CP to distinguish between high-priority node-level operations (for example associations) and standard session-level updates, preventing signaling failures during congestion.
 - DSCP marking for RADIUS traffic: This feature enables QoS prioritization for AAA signaling exchanged between the cnBNG CP and external RADIUS servers. By prioritizing authentication and accounting messages, the system ensures reliable subscriber lifecycle management even during high network load.
 - CP-UP liveliness optimization: This feature complements QoS by replacing static PFCP heartbeat intervals with user-configurable, dynamic timers. By employing an adaptive backoff algorithm, the cnBNG CP accelerates link failure detection and reduces overall system convergence time. This approach minimizes subscriber session impact during CP-UP link failures.
-
- [Feature Summary and Revision History, on page 2](#)
 - [DSCP marking for replication traffic, on page 2](#)
 - [DSCP marking for CP-UP control packets, on page 5](#)
 - [DSCP marking for RADIUS traffic, on page 8](#)
 - [CP-UP liveliness optimization, on page 9](#)

Feature Summary and Revision History

Summary Data

Table 1: Summary Data

Applicable Products or Functional Area	cnBNG
Applicable Platform(s)	SMI
Feature Default Setting	Disabled - Configuration Required
Related Changes in this Release	Not Applicable
Related Documentation	Cloud Native BNG Control Plane Command Reference Guide

Revision History

Table 2: Revision History

Revision Details	Release
First introduced.	2026.03.0

DSCP marking for replication traffic

The DSCP marking for replication traffic is a state-synchronization feature that

- applies configurable DSCP values to inter-instance state data
- ensures geo-redundancy and failover readiness by protecting synchronization messages from congestion, and
- prevents synchronization delays between CP clusters to maintain state consistency.

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
DSCP marking for replication traffic	2026.03.0	Enables DSCP marking for internal state-synchronization traffic between redundant control plane clusters to ensure deterministic synchronization. By applying consistent DSCP values, the cnBNG protects critical replication flows from congestion and eliminates data lag, facilitating faster and more reliable failover convergence for geo-redundancy.

Inter-instance state data

Inter-instance state data is the critical operational information that must be synchronized between redundant CP instances to ensure the system functions as a single, cohesive entity. In a distributed cnBNG architecture, this data is the source of truth that allows the system to maintain high availability and seamless failover capabilities.

This data is essential for maintaining consistency across distributed clusters and typically includes:

- Cluster state (etcd): Information regarding the overall health and configuration of the CP cluster, including SRG caches and UP manager data.
- Cache synchronization: Data related to IP Address Management (IPAM) and system uptime, ensuring all CP instances have an identical view of network resources.

Need for DSCP marking for replication traffic

To ensure reliable failover and geo-redundancy, CP instances require real-time state synchronization. Currently, this critical traffic competes for bandwidth within the data center network, making it susceptible to congestion-related delays. Without traffic prioritization, these delays can compromise system stability and increase the risk of state inconsistency during failover events.

The DSCP marking for replication traffic feature mitigates these risks by implementing a deterministic QoS policy for internal traffic flows. By assigning a specific DSCP value to state-synchronization data at the geo-replication endpoint, the system ensures that this traffic receives preferential treatment across the network infrastructure.

Benefits of DSCP marking for replication traffic

- Deterministic sync: Guarantees the timely delivery of state information even during periods of high network load, ensuring that synchronization data is forwarded ahead of non-critical flows.
- Failover readiness: Ensures that redundant CP nodes remain perfectly synchronized, eliminating data lag and facilitating faster, more reliable failover convergence.
- Congestion resilience: Protects critical administrative and control messages from being dropped during congestion, minimizing queuing delays and preventing packet loss.

- Reduced latency and jitter: Ensures that cluster state, cache information, and geo-redundancy data reach their destination with minimal jitter and maximum reliability.

Limitations of DSCP marking for replication traffic

DSCP marking for CDL is not supported.

How DSCP marking for replication traffic works

Summary

The geo-replication endpoint and the data center network coordinate to prioritize internal state-synchronization data. This process ensures the high-priority delivery of critical replication flows, such as `etcd` and `cache-pod` synchronization, to maintain cluster state consistency during periods of network congestion.

Workflow

These are the stages for DSCP marking for replication traffic.

1. Traffic classification: The cnBNG CP applies the configured DSCP value to all outgoing replication traffic at the geo-replication endpoint.
2. Header marking: The geo-replication endpoint embeds the DSCP marking into the IP header of the replication packets.
3. Preferential forwarding: The data center network identifies the marked packets and applies preferential forwarding treatment based on the assigned DSCP value.
4. Reliable delivery: The replication flows reach the destination CP instance with reduced latency and minimal packet loss.

Configure DSCP marking for replication traffic

Configure DSCP marking for replication traffic to ensure the prioritized delivery of inter-instance state synchronization data.



Note You must apply this configuration to all redundant cnBNG CP clusters before they are deployed, as this is a static setting that cannot be modified on the fly.

Procedure

Step 1

Set the DSCP value for the geo-replication endpoint.

Ensure that all instances on each cluster must have the same DSCP value.

Example:

```
instance instance-id 1
  endpoint geo
```

```
dscp 63
end
```

Step 2 Verify the configuration using the show running configuration.

Example:

```
cnbng# show running-config instance instance-id 1 endpoint geo
instance instance-id 1
endpoint geo
replicas 1
dscp 63
interface geo-internal
vip-ip 10.1.35.1 vip-port 7001
exit
interface geo-external
vip-ip 10.1.35.1 vip-port 7002
exit
exit
exit
```

Step 3 Verify that the traffic packets contain the configured DSCP value for replication traffic.

Example:

```
Frame 3: Packet, 100 bytes on wire (800 bits), 100 bytes captured (800 bits)
Ethernet II, Src: Intel_18:99:f8 (40:a6:b7:18:99:f8), Dst: Cisco_b2:2d:c5 (04:76:b0:b2:2d:c5)
Internet Protocol Version 4, Src: 10.1.35.1, Dst: 10.2.35.1
  0100 .... = Version: 4
  .... 0101 = Header Length: 20 bytes (5)
  Differentiated Services Field: 0xfc (DSCP: Unknown, ECN: Not-ECT)
    1111 11.. = Differentiated Services Codepoint: Unknown (63)
    .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
  Total Length: 86
```

DSCP marking for CP-UP control packets

The CP-UP control traffic prioritization is a granular QoS feature that

- enables message-level DSCP marking for PFCP and GTP-U control signaling
- distinguishes between critical node-level operations and standard session-level updates, and
- ensures reliable signaling between the Control Plane (CP) and User Plane (UP) to prevent session instability.

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
DSCP marking for CP-UP Control Packets	2026.03.0	You can improve CP-UP signaling reliability during congestion by enabling granular, message-level QoS classification for N4 signaling protocols, such as PFCP and GTP-U, to prioritize critical node-level operations over session-level updates.

Need for DSCP marking for CP-UP control packets

Existing cnBNG devices treat all CP-UP signaling traffic equally. Critical node-level operations, such as associations and heartbeat exchanges, are essential for maintaining system health, yet they currently compete for bandwidth with standard session-level updates. During periods of network congestion, this lack of traffic differentiation can lead to packet loss, signaling failures, and subscriber session instability.

The DSCP marking for CP-UP control packets feature introduces granular QoS capabilities to the N4 interface. By enabling message-level DSCP classification for PFCP and GTP-U control protocols, CP can distinguish between high-priority node-level operations and standard session-level updates. This feature allows operators to apply specific DSCP tags to different message types, ensuring that vital signaling packets receive preferential forwarding treatment across the data center network.

Benefits of DSCP marking for CP-UP control packets

- **Message-Level granularity:** Enables operators to assign higher priority to critical node-level messages compared to standard session-level updates.
- **Signaling reliability:** Reduces CP packet loss, ensuring that essential association and heartbeat messages are delivered reliably even during periods of network congestion.

How DSCP marking for CP-UP control packets works

Summary

The CP and UP utilize the N4 interface to classify and prioritize signaling traffic. This process ensures that critical node-level operations, such as associations and heartbeats, receive higher priority than standard session-level updates to maintain system stability.

Workflow

These are the stages of DSCP marking for CP-UP control packets.

1. The cnBNG CP classifies outgoing signaling packets based on their message type, for example node-level or session-level, at the control plane.
2. The CP applies the corresponding DSCP marking to the classified signaling packets before transmission.
3. The N4 interface transmits the prioritized packets across the network to the UP.

4. The UP processes the incoming packets, and response packets inherit the DSCP marking from the original request to ensure consistent priority treatment.

Configure DSCP marking for CP-UP control packets

Set DSCP values for N4 protocol endpoints to prioritize CP-UP control packet traffic across clusters.

DSCP marking enables network devices to identify and prioritize CP-UP control traffic. You can configure DSCP values globally or for specific message types and must apply the configuration consistently across all instances in a cluster.

Procedure

- Step 1** On each cluster instance, set the DSCP value for the N4 protocol endpoint.

You must configure every instance with the same configuration.

Example:

```
instance instance-id 1
endpoint n4-protocol
  dscp node-association af12
end
```

- Step 2** Verify the DSCP marking for CP-UP control packets configuration.

Example:

```
cnbng# show running-config instance instance-id endpoint n4-protocol
instance instance-id 1
endpoint n4-protocol
  replicas 2
  dscp node-association af11
  dscp session-message af21
  dscp srg-message af12
  dscp reconcile-message af13
  dscp gtpu-control-packet af43
  retransmission max-retry 1
exit
instance instance-id 2
endpoint n4-protocol
  replicas 2
  dscp node-association af11
  dscp session-message af21
  dscp srg-message af12
  dscp reconcile-message af13
  dscp gtpu-control-packet af43
  retransmission max-retry 1
exit
exit
```

DSCP marking for RADIUS traffic

The DSCP marking for RADIUS traffic is a QoS feature that

- enables configurable DSCP marking for RADIUS protocol traffic
- prioritizes AAA signaling and CoA responses between the cnBNG CP and external AAA servers, and
- ensures reliable signaling and subscriber session stability during network congestion.

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
DSCP marking for RADIUS traffic	2026.03.0	Enables configurable DSCP marking for RADIUS protocol traffic to prioritize AAA signaling and ensure subscriber session stability during network congestion.

Benefits of RADIUS DSCP marking

- Signaling reliability: Protects time-sensitive RADIUS signaling from being dropped or delayed during periods of high network congestion.
- Enhanced session stability: Minimizes authentication timeouts and signaling failures, ensuring that subscriber session establishment and lifecycle management remain consistent.
- Deterministic performance: Guarantees that AAA signaling packets receive preferential forwarding treatment across the data center network, reducing jitter and latency for critical control-plane traffic.
- Operational determinism: Provides a consistent QoS policy for RADIUS interactions, allowing network administrators to maintain predictable control-plane behavior in large-scale deployments.

How RADIUS DSCP marking works

Summary

The cnBNG CP classifies and marks outgoing RADIUS signaling traffic to ensure it receives preferential treatment across the data center network. This process ensures that critical authentication and accounting messages are prioritized, maintaining subscriber session stability during periods of network congestion.

Workflow

These are the stages of RADIUS DSCP marking.

1. Policy configuration: Define a traffic priority level within the RADIUS endpoint configuration.
2. Traffic classification: The cnBNG CP identifies all outgoing RADIUS signaling packets associated with the configured instance.

3. Header marking: The cnBNG CP embeds the configured DSCP value into the IP header of each RADIUS packet at the egress point.
4. Preferential forwarding: The data center network identifies the marked packets and applies preferential forwarding treatment based on the assigned DSCP value.
5. Reliable delivery: The RADIUS signaling packets reach the destination aaa server with reduced latency and minimal packet loss.

Configure RADIUS DSCP marking

Use this task to assign a specific DSCP value to RADIUS signaling traffic, ensuring preferential forwarding treatment for authentication, authorization, and accounting messages across the data center network.

Procedure

Configure the DSCP value for RADIUS traffic.

Apply the same configuration on all instances of each cluster.

Example:

```
instance instance-id 1
  endpoint radius
    dscp 20
  exit
exit
```

CP-UP liveliness optimization

The CP-UP liveliness optimization is a failure-detection cnBNG feature that

- replaces static Packet Forwarding Control Protocol (PFCP) heartbeat intervals with user-configurable, dynamic timers
- employs an adaptive backoff algorithm to accelerate link failure detection, and
- reduces overall system convergence time to minimize subscriber session instability during CP-UP link failures.

Table 6: Feature History Table

Feature Name	Release Information	Feature Description
CP-UP liveliness optimization	2026.03.0	You can now reduce CP-UP failure detection time and improve SRG convergence by optimizing CP-UP liveliness detection with configurable heartbeat timers and an adaptive backoff mechanism.

Key terms

- Packet Forwarding Control Protocol (PFCP): The control plane protocol used between the CP and UP to manage sessions and exchange heartbeat messages for liveliness detection.
- Liveliness detection: The process of verifying the operational status of a network link or connection by periodically exchanging keep-alive or heartbeat messages.
- Convergence time: The duration required for the network to detect a failure and successfully switch traffic to a redundant path; reducing this is a primary goal of this feature.
- Heartbeat interval: The configurable time period between consecutive heartbeat messages sent between the CP and UP.

Benefits of CP-UP liveliness optimization

- Faster failure detection: Reduces the detection time from a fixed 90 seconds to a configurable range, with a minimum of 10 seconds for a heartbeat configuration using a five-second interval and a miss-count of two.
- Configurable convergence: Provides the flexibility to tune heartbeat intervals (5 to 60 seconds) to meet specific SRG convergence targets.
- Adaptive efficiency: Automatically reduces heartbeat intervals during failure scenarios to speed up detection, while maintaining a stable, bounded probing rate during healthy operation.
- Reduced subscriber impact: Minimizes the duration of service disruption during CP-UP link failures by triggering failover mechanisms more rapidly.

How CP-UP liveliness optimization works

Summary

The CP and UP coordinate through an adaptive heartbeat mechanism to monitor the health of their connection. This mechanism enables rapid failure detection and improves SRG convergence by dynamically adjusting heartbeat intervals based on the CP-UP link status

Workflow

These stages describe how CP-UP liveliness optimization works.

1. The UP sends heartbeat messages at a user-configured interval, which ranges from 5 to 60 seconds, during normal operation.
2. The CP receives and responds to these heartbeat messages to confirm link health.
3. The CP detects a missed heartbeat request, indicating a potential link failure.
4. Upon each consecutive missed heartbeat, CP progressively reduces the heartbeat interval to accelerate failure detection.
5. When a successful heartbeat request is received, CP immediately restores the heartbeat interval to the configured value.
6. After the configured number of consecutive heartbeat failures, the CP marks the CP-UP connection as inactive and triggers failover procedures.

Configure the CP-UP heartbeat timer

Use this task to configure the CP-UP heartbeat timer for liveness detection between the CP and UP.

Procedure

Set the heartbeat interval.

You can configure the heartbeat either globally or for a specific user plane. Configuration for a user plane takes precedence over the global configuration.

The heartbeat configuration must be identical on both CP and UP instances to ensure synchronization.

Example:

This example shows how to configure the heart beat at the global level.

```
user-plane
 pfcg-heartbeat interval 10 miss-count 4
exit
```

This example shows how to configure the heart beat for the specific user plane configuration

```
user-plane
instance 1
 user-plane asr9k-0-1
  pfcg-heartbeat interval 10 miss-count 4
  exit
 user-plane asr9k-0-2
  pfcg-heartbeat interval 10 miss-count 4
  exit
exit
exit
```

Example:

Note

If no values are configured, the cnBNG CP applies a default interval of 30 seconds and a miss-count of 3.

When the interval is set to 30 seconds and the miss-count is 3, either by configuration or default, the heartbeat interval is not reduced and remains fixed at 30 seconds. This ensures backward compatibility with existing UPs.
