End-to-End Flow Control

There are various scenarios that might generate sudden spike in the traffic that goes to the cnBNG control plane (CP). To handle these spikes in traffic, it is necessary to flow control and rate limit the CP ingress to ensure that service applications are not overwhelmed with these bursts. The end-to-end flow control feature optimizes flow control and rate limit of the traffic toward the CP ingress. This chapter covers the end-to-end flow control feature on cnBNG user plane (UP).

For details on end-to-end flow control functionality on cnBNG CP, see the *End-to-End Flow Control* chapter in the *Cloud Native BNG Control Plane Configuration Guide*.

End-to-End Flow Control

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
End-to-End Flow Control Between User Plane and Control Plane	Release 7.4.2	In cloud native BNG (cnBNG) networks, this feature allows you to define the rate at which the user plane (UP) sends messages, such as control packets and data packet notifications, to the control plane (CP). The feature regulates the CP ingress traffic flow at the source-of-origin (UP), helping avoid network congestion and packet loss. The feature introduces these commands: • ipoe fsol-flow-control • pppoe fsol-flow-control • up-cp-notification flow-control • up-cp-stats flow-control

The Cloud Native Broadband Network Gateway (cnBNG) manages residential subscribers from different access planes in a centralized way. It accepts and identifies subscriber control plane (CP) traffic coming from multiple user planes (UPs) associated with the CP. When the number of UPs scale, the amount of traffic that eventually reaches the CP from each UP also multiplies.

There are various scenarios where the traffic flow between the CP and UP must be regulated. This is to ensure that the CP attends all the service requests without service interruption. The scenarios that create burst or higher flow rates in the traffic flows are as follows:

- Power outage in a residential area
- Access network outage for a specific period
- Catastrophic events like process crash, route processor reboot, chassis reload and so on, on UP

These scenarios generate sudden spike in traffic going to the CP. To handle these spikes in traffic, it is necessary to use flow control to rate limit the CP ingress traffic to ensure that service applications are not overwhelmed with these bursts. The end-to-end flow control feature optimizes flow control and rate limit of the traffic toward the CP ingress.

How it Works

There are two types of traffic that enter or exit the CP—the control traffic that is responsible for subscriber session creation and the control traffic on already provisioned subscriber session. The application infrastructure (App-Infra) features such as Dispatcher and Overload Control, facilitate the cnBNG CP ingress packet flow control. For details on end-to-end flow control functionality on cnBNG CP, see the *End-to-End Flow Control* chapter in the *Cloud Native BNG Control Plane Configuration Guide*.

On UP side, the end-to-end flow control feature allows you to define the rate at which packet notifications from the UP reach the CP. The protocol packets from CPE, and messages such as statistics and notification events that are locally generated on the UP are subjected to this flow control such that the number of messages reaching the CP are controlled at the source-of-origin itself.

cnBNG UP provides flow control functionality for the following packets:

- Protocol packets from CPE which include:
    - IPoE DHCPv4 DISCOVER rate control from UP to CP
    - IPoE DHCPv6 SOLICIT rate control from UP to CP
    - PPPoE-PTA (PPP Termination and Aggregation) or LAC (L2TP Access Concentrator) PADI (PPPoE Active Discovery Initiation) rate control from UP to CP
    - PPPoE-PTA DHCPV6 SOLICIT rate control from UP to CP
- Locally generated messages on UP (this is common for IPoE, PPPoE-PTA, and PPPoE-LAC):
    - Subscriber delete notification (say, during mark-and-sweep procedure, session deletion by UP administrator, and so on)
    - PPP keep alive timer expiry notification
    - Subscriber session or service periodic statistics notification
- DHCP packets:
    - Network processing unit-level (NPU-level) local packet transport service (lpts) flow control for DHCP broadcast packets
- PPPoE packets:
    - NPU-level lpts flow control for PADI broadcast packets

Based on these packet types, the flow control functionality on UP is broadly classified into:

- UP protocol packet punt flow control
- UP notifications events flow control, which includes:
    - UP delete notification flow control
    - UP statistics report flow control

You can use specific commands on UP to set various flow control limits that define the number of messages sent from UP to CP for each second. The excessive messages are queued up with a limited queue size. The maximum packet holding in the queue is 64K messages. For command details, see the *Configuration* section.

Restrictions

End-to-end flow control feature is applicable only for L2-connected topology— where the DHCP client is connected to cnBNG UP through a direct link, without a relay agent being present in between them. Whereas, the feature is not applicable for a relay chaining topology— where a lightweight DHCPv6 relay agent (LDRA) is present in between the DHCP client and the cnBNG UP node. This is because, the cnBNG UP does not perform full packet decoding to identify the SOLICT or DISCOVER packets in such scenarios.

Configure End-to-End Flow Control on cnBNG User Plane

End-to-end flow control configuration on cnBNG UP involves these high-level tasks:

- UP protocol packet punt flow control, using **fsol-flow-control** command
- UP notifications events flow control, which includes:
  - UP delete notification flow control, using **up-cp-notification flow-control** command
  - UP statistics report flow control, using **up-cp-stats flow-control** command

Guidelines for Enabling End-to-End Flow Control Feature

Enabling end-to-end flow control feature is subjected to these guidelines:

- With flow control feature applied, you might experience packet loss in the following high availability (HA) scenarios such as *cnbng-nal* process restart, RP fail over, *cnbng-nal* SMU activation, and bring down of CP-UP association.
- The queue size of the flow control packet is limited to 64K messages due to memory constraints. If queue is already full, UP drops the new packets.
- To ensure a robust system performance, you must choose the configuration parameters for the flow control feature based on your network topology and bandwidth requirement.
- Refer the *Cloud Native BNG Control Plane Configuration Guide* for details on various flow control configuration parameters on CP.

Configuration Example

- **UP protocol packet punt packet flow control:**

  The UP protocol packet punt flow control limit ranges from 50 to 400 packets for each second, default being 100.

  ```
  Router#configure
  Router(config)#cnbng-nal location 0/RSP0/CPU0
  Router(config-cnbng-nal-local)#ipoe fsol-flow-control 70
  Router(config-cnbng-nal-local)#commit
  ```

  Similarly, for PPPoE packets, use:

  ```
  Router(config-cnbng-nal-local)#pppoe fsol-flow-control 60
  ```

- **UP notification events flow control:**
  - **UP delete notification flow control**:

The UP delete notification flow control limit ranges from 20 to 400 packets for each second, default being 100.

```
Router(config-cnbng-nal-local)#up-cp-notification flow-control 70
```

- **UP statistics report flow control**

  The UP statistics report flow control limit ranges from 20 to 500 packets for each second, default being 150.

  ```
  Router(config-cnbng-nal-local)#up-cp-stats flow-control 70
  ```

Running Configuration

This is the running configuration on cnBNG UP that includes basic UP configuration as well:

```
cnbng-nal location 0/RSP0/CPU0
 hostidentifier CNBNG-UP2
 up-server ipv4 192.0.2.1 vrf default
 cp-server primary ipv4 198.51.100.1
 auto-loopback vrf default
  interface Loopback0
   primary-address 1.0.0.1
  !
 !
disable-secondary-address-notification
cp-association retry-count 10
ipoe fsol-flow-control 70
pppoe fsol-flow-control 60
up-cp-notification flow-control 70
up-cp-stats flow-control 70
max-create-in-progress 600
secondary-address-update enable
!
```

Verification

Use the following **show** command to see the respective packet counters and to check the packet drops. Check the **V4 DHCP FSOL drop flow ctrl** and **V6 DHCP FSOL drop flow ctrl** parameters in this output.

```
Router#show cnbng-nal counters type all
Thu Feb  3 11:13:49.767 UTC

Location: 0/RSP0/CPU0
.
.
.
Packet Counters
------------------

.
.
.
Counter name                               Value
=============                              =====
V4 DHCP FSOL drop flow ctrl                9490
V6 DHCP FSOL drop flow ctrl                241350
.
```