Juniper BNG CUPS User Guide

Published

2025-12-22

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper BNG CUPS User Guide

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide

Use this guide to perform initial configuration, monitor and use Juniper BNG CUPS software.

1

CHAPTER

Overview

IN THIS CHAPTER

Juniper BNG CUPS Overview

In an integrated Broadband Network Gateway (BNG), such as a Juniper MX Series router configured with subscriber management services, one control plane is paired with one user plane running on the same hardware platform. The control plane handles functions including, subscriber session state management, AAA, IP address assignment, and policy enforcement. The user plane handles functions including routing and traffic management and collection of subscriber statistics.

A new architecture, called Control and User Plane Separation (CUPS) separates the control plane and user plane functions into different network elements. The control plane and user planes are tethered through a set of defined open interfaces. These interfaces are used for exchanging states and for relaying control packets between the planes. The control plane together with one or more user planes forms a disaggregated BNG.

Juniper BNG CUPS Controller (BNG CUPS Controller) is a cloud-native application that realizes the control plane component of a disaggregated BNG. You install and run BNG CUPS Controller on a Kubernetes cluster such as those created by the Juniper BBE Cloudsetup utility (see BBE Cloudsetup). The BNG CUPS Controller forms a disaggregated BNG with Juniper routing devices that are configured to operate as BNG User Planes.

The BNG CUPS Controller can be deployed in a single geographical cloud environment (see Figure 1 on page 2) or in a multiple geographical cloud environment (see Figure 2 on page 3).

Figure 1: Juniper BNG CUPS Architecture in a Single Geographical Deployment

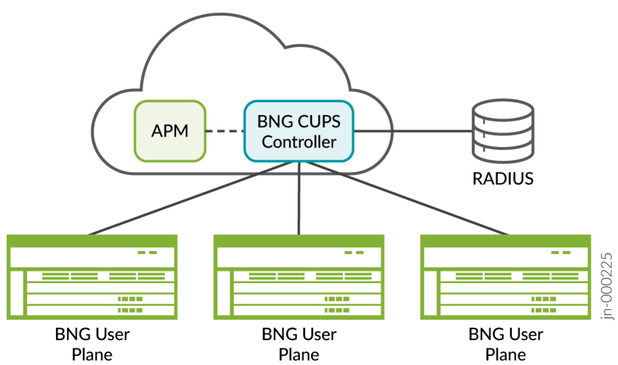

Figure 2: Juniper BNG CUPS Architecture in a Multiple Geographical Deployment

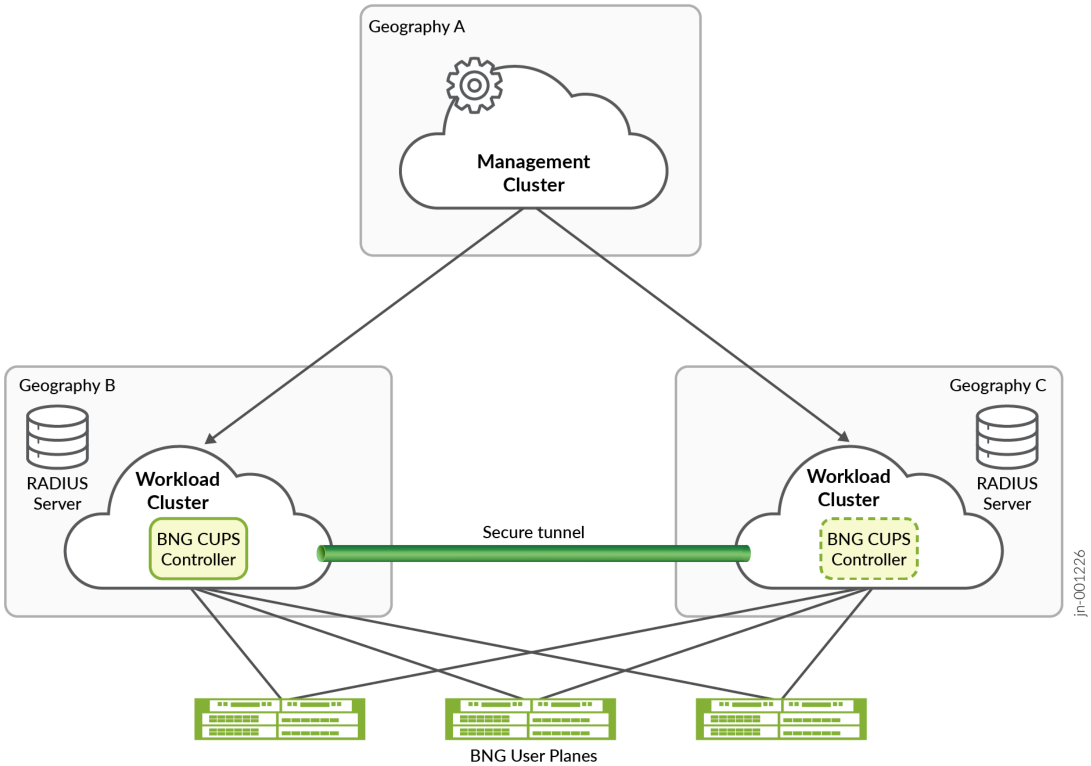

Benefits of Juniper BNG CUPS

A BNG CUPS Controller enables you to use network resources more efficiently through:

- Centralized address allocation.
- User plane load balancing.
- Centralized management and control.
- Increased scale. The cloud environment that Juniper BNG CUPS utilizes enables you to increase the number of subscribers supported.
- Locational independence and separate life-cycle management and maintenance.
- Throughput and latency optimization, because the BNG User Planes are closer to the subscribers.
- Resiliency in responding to network failure events such as a BNG User Plane failure or failure of a transport connection between an access node and the BNG User Plane.
- Live subscriber placement when changes in performance occur or when network congestion occurs.

- Can be deployed in a multiple geographical cloud environment, which ensures the high availability of BNG CUPS Controller.

Why Migrate from an Integrated Broadband Network Gateway to a Disaggregated Broadband Network Gateway

Rising operational costs with declining or flattening revenues have driven telco service providers to rethink the way they plan, design, and operate their networks. Telcos are following the lead of cloud operators looking to apply cloud and data center design principles to their next-generation network architectures as a way to save costs. Further, decoupling the operating system software from the hardware allows you to manage hardware and software life cycles separately.

The following use cases for Juniper BNG CUPS show the benefits of migrating from an integrated BNG to a disaggregated BNG:

- Centralized Address Pool Management

  IP addresses have become a precious resource. If you don't have enough available, subscribers can't access the network. Yet purchasing new addresses has become enormously expensive. Service providers do everything in their power to optimize and efficiently utilize their limited IP address space, but traditional networks with integrated BNGs make it challenging. Operators are required to perform BNG planning and manually distribute (and redistribute) IP address prefixes among the BNGs that are based on expected and changing scale of each BNG.

  Automating IP prefix assignment to adapt to BNG scaling demands and dynamically reclaiming unused IP address prefixes for redeployment to a different BNG as scaling needs decreases, alleviates the need for operators to perform intensive and potentially error-prone IP prefix configurations on each BNG. The need is reinforced by Juniper BNG CUPS resiliency subscriber groups that would otherwise increase operator complexity to manually configure and assign IP address pools on a Subscriber Group basis

  Juniper makes it possible to manage IP address pools as a shared resource, and automatically allocate IP addresses to any user plane across the network. With the cloud-native Address Pool Manager, service providers can do the following:

  - Improve operational efficiency by automatically adding IP addresses when needed—APM proactively monitors IP address pools across all BNG entities in the network. If a user plane crosses a predefined threshold, APM automatically links it to a new address pool. You get the IP address resources you need, where and when you need them, without having to manage address pools manually or build and maintain homegrown tools.

  - Lower costs by maximizing IP address utilization—By monitoring all downstream user planes centrally, APM can identify any BNG nodes with large, underutilized address pools. In a traditional network, those unused addresses would sit idle. APM automatically reclaims and redistributes them across the network where needed, optimizing operational costs for public IPv4 address management.

For more information about APM, see *Address Pool Manager User Guide*

- **Subscriber Stateful Resiliency**

  One of the primary use cases of Juniper BNG CUPS is resiliency to support hitless failover in the event of a an unplanned BNG User Plane failure. You define a resiliency subscriber group where one BNG User Plane operates as the active BNG User Plane and another BNG User Plane serves as a backup. The backup BNG User Plane assume control of the subscriber sessions in the event of a failure. The cloud-hosted BNG CUPS Controller then pre-stages the BNG User Planes and, depending on the redundancy option used, continually programs backup BNG User Planes with the relevant state information. In the event the active BNG User Plane plane fails, the BNG CUPS Controller automatically activates the pre-staged backup and reroutes traffic accordingly.

  You'll be able to choose from three redundancy options, depending on the level of disruption acceptance for a given service or SLA:

  - Hot standby—The controller continually programs all subscriber session state information on the backup BNG User Planes, enabling hitless failover that is undetectable.

  - Warm enhanced oversubscribed standby—A backup BNG User Plane has a limited subscriber forwarding state installed and the full subscriber session state maintained in memory. If an active BNG User Plane fails, the backup assumes forwarding of subscriber sessions and then installs the remaining subscriber state. There is a short time frame until the subscriber session SLA is restored. This approach supports N:1 redundancy.

  - Warm high scale standby—Uses a Routing Engine based backup mode. The backup BNG User Plane holds the full subscriber state on the Routing Engine and doesn't program the Packet Forwarding Engine until the subscriber group becomes active. This approach supports M:N redundancy.

  Also, there are two ways in which the active BNG User Plane is selected for redundancy. They are described in the following:

  - BNG CUPS Controller controlled—The BNG CUPS Controller determines the active BNG User Plane based on the configuration and logical-port and network instance reports from the BNG User Plane.

  - BNG User Plane controlled—Determined by the access network. The state of the connection to the BNG User Plane determines which BNG User Plane is active.

  > **NOTE**: For BNG User Plane controlled subscriber groups, at least one BNG User Plane must be configured as the hot backup (the `backup-mode` set to `hot` on the BNG User Plane).

> The access network must be provisioned to have the preferred active port configured on the BNG User Plane, which has the `backup-mode` command configured as `hot`.

- Hitless BNG User Plane Maintenance

  In traditional vertically integrated networks, most maintenance tasks, such as changing line cards, updating software, and so on, require a scheduled maintenance window. Since you're bringing down the node and all subscribers attached to it, you always risk disrupting services and frustrating subscribers. Additionally, since maintenance windows are typically scheduled late at night, you pay higher overtime costs for that maintenance. A centralized control plane and shared state information make planned maintenance much simpler and less disruptive.

  The process is straightforward:

  1. Operators use the controller to orchestrate the transfer of all subscriber state information from the current user plane to a new one.

  2. They configure the transport network to send traffic to the new user plane instead of the old.

  3. Since the new user plane already has state information for all subscribers, it exists in a warm oversubscribed standby stats and quickly brings up those sessions without service disruption.

  4. Operators perform the maintenance and, once complete, reverse the process and orchestrate traffic back to the original user plane.

  Furthermore, if the subscribers on the user plane to undergo maintenance are all part of one or more resiliency Subscriber Groups, the process is even more straightforward:

  1. Technicians use the controller to initiate subscriber group switchover to the backup User Plane for any subscriber groups in which the User Plane is the Active User Plane.

  2. The User Plane is in Backup mode and can be offlined to perform maintenance. Once complete, the User Plane will resume in a backup role for all subscriber groups and optionally resume the Active role by the technician performing subscriber group switchover.

  The whole procedure can be handled in a streamlined, low risk way during normal business hours, with subscribers never noticing a thing. This means you can continually update your network easily and inexpensively, while improving customer satisfaction and supporting more stringent and profitable SLAs.

- Smart Subscriber Load Sharing

  In traditional broadband networks, BNGs act as siloed entities. If you want to distribute BNG User Planes, you're always at risk of running out of capacity, which means you typically have to over provision. With the centralized control enabled by Juniper BNG CUPS, you can group BNG User Planes together and treat them as a shared pool of resources. In this model, you group together BNG User Planes that are part of the virtual resource pool (called a load-balancing group). The BNG CUPS

Controller proactively monitors their subscriber loads for all BNG User Planes that are part of the same load-balancing group. If a BNG User Plane exceeds a given threshold, the BNG CUPS Controller begins shifting sessions to a less-loaded BNG User Plane. This results in you not having to worry about accurately forecasting or overprovisioning subscriber scale for a given market. Instead, you can share BNG User Planes as needed and continually maximize all available resources in the infrastructure.

Required Configuration Changes

Because the BNG CUPS Controller and the BNG User Planes are separated, you must perform configurations on both the BNG CUPS Controller and the BNG User Planes. You will perform the majority of the configurations on the BNG CUPS Controller.

Configure the following features on the BNG CUPS Controller:

- Subscriber groups
- Load balancing groups
- BNG User Plane profiles
- Dynamic profiles
- Autosense VLANs
- DHCP/DHCPv6 local server and relay
- L2TP
- AAA services
  - RADIUS
  - Access profile
  - Address assignment
  - Domain map
- Subscriber firewall filters
- Subscriber Class of Service (CoS)
  - Routing instances for L3 aware control plane applications (for example, DHCP and DHCPv6)
  - Subscriber groups for resiliency
- Load balancing groups

Configure the following functions on the BNG User Planes:

- Subscriber management mode
  - BNG User Planes
  - BNG CUPS Controller reachability
- Resource monitoring
- Routing instances for forwarding
- Routing protocols for each routing instance

> **NOTE**: Most of the control plane commands from the integrated BNG carry over to Juniper BNG CUPS, with minor extensions for Juniper BNG CUPS.

Operational Changes

Juniper BNG CUPS separates the operational commands into BNG CUPS Controller and BNG User Plane commands. The majority of the BNG commands run on the BNG CUPS Controller. To help with troubleshooting, some of the operational commands run on the BNG User Planes.

The BNG CUPS Controller CLI is where you configure and manage aspects of the following:

- Subscriber management. The following shows some examples:
  - Client protocols
  - AAA
  - Accounting
  - Dynamic profiles
  - Subscriber session state
  - Subscriber groups
  - Load balancing groups
  - And so on
- Associated infrastructure. The following shows some examples:
  - BNG User Planes
  - Control plane instances

- BNG CUPS Controller and BNG User Plane communications
- And so on.

The BNG User Plane CLI is where you configure and manage aspects of the following:

- Subscriber forwarding. The following shows some examples:
  - Routing and routing instances
  - Access-facing interfaces
  - Dynamic autosense VLANs
  - And so on
- Associated infrastructure. The Following shows some examples:
  - BNG User Plane mode operations
  - BNG CUPS Controller associations
  - And so on

Juniper BNG CUPS Feature Support

Juniper BNG CUPS supports most of the same subscriber management features from integrated BNG:

Client Protocol Support

- Dynamic auto-sensed VLANs
- DHCPv4 and DHCPv6 single and dual stack subscribers for local server
- DHCPv4 and DHCPv6 single and dual stack subscribers for relay
- PPP/PPPoE v4/v6 and dual stack subscribers
- L2TP LAC
- Interface combinations—Ethernet, aggregated Ethernet, Pseudowire, and Redundant Pseudowire

AAA Services

- RADIUS based authentication and authorization
- RADIUS change of authorization and disconnect
- Address assignment from:
  - RADIUS—Including framed IP address and framed routes

  - Dynamic address pools created by Address Pool Manager or local reserve
  - Statically configured address pools
- RADIUS-based accounting:
  - Subscriber accounting, including interim accounting
  - Subscriber service accounting
- Subscriber idle timeout and session timeout
- Domain map
- Service profiles

Class of Service (CoS)

- You can use dynamically created scheduler maps, schedulers, and traffic control profiles.
- You can add the following services to dynamic flows:
  - Classifiers
  - Rewrite-rules
  - Output traffic control profiles with scheduler maps
- Hierarchical class of service, including support for interface sets

Firewall Services

- Parameterized filters and policers through a dynamic service profile
- Static filters and policers

Multicast Services Features

- Centralized and distributed multicast services are activated when the subscriber logs in or activated through a RADIUS change of authorization.

Lawful Intercept

- Activation and deactivation of RADIUS-based lawful intercept for a flow-based subscriber during login and logout, on both the BNG CUPS Controller and the BNG User Plane
- Activation and deactivation of RADIUS-based lawful intercept for a flow-based subscriber using RADIUS change of authorization (CoA), on both the BNG CUPS Controller and the BNG User Plane
- Activation and deactivation of Dynamic Tasking Control Protocol (DTCP) based lawful intercept for a flow-based subscriber, on both the BNG CUPS Controller and the BNG User Plane

- Attaching of lawful intercept drop policy for a flow-based subscriber, on both the BNG CUPS Controller and the BNG User Plane

- Reporting of intercept-related events using SNMP traps to a mediation device on the BNG CUPS Controller

Management of Multiple BNG User Planes

- A BNG CUPS Controller can manage up to 32 BNG User Planes. The multiple BNG User Plane architecture defines a BNG User Plane instance per BNG User Plane to encapsulate data and work within a BNG User Plane.

- BNG User Planes are assigned to a control plane instance. A control plane instance initiates an association with a BNG User Plane upon assignment.

Figure 3 on page 11 shows a multiple BNG User Plane topology.

Figure 3: BNG CUPS Controller with Multiple BNG User Planes

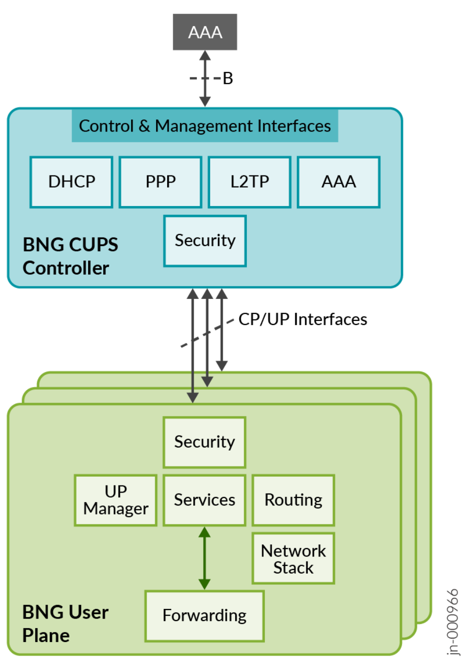

Smart Session Load Balancing

Gives the operator the capability to distribute subscriber loads across the BNG User Planes in the network by moving subscribers from one BNG User Plane to another. Fast failover is a use case for subscriber session load balancing. The fast failover use case occurs when a BNG User Plane's access port goes down and subscribers are rebalanced over to another access port on the same BNG User Plane.

Subscriber Stateful Resiliency

- Ensures resiliency across BNG User Planes where the BNC CUPS Controller holds the primary state for any subscriber session. The BNG User Plane holds the active forwarding state or backup forwarding state for a particular subscriber session.
- Subscriber resiliency is achieved through the use of subscriber groups (`subscriber-groups` configuration).

Subscriber Session Steering

Places subscribers in the desired BNG User Plane based on a RADIUS service group vendor-specific attribute (VSA). This VSA specifies the subscriber services level (SLA) that the BNG CUPS Controller communicates to the user plane selection function. It then uses the SLA in selecting the BNG User Plane that meets the subscriber session service requirements

Additional Information

Forwarding Class Handling

The `forwarding-class` configuration is a special case. You must configure the forwarding class names on the BNG User Planes that you configure on the BNG CUPS Controller.

These matching configurations are required because the number of forwarding classes is limited. Also, other entities in the BNG User Plane use the forwarding class. Thus, the BNG CUPS Controller's forwarding classes must be consistent with the BNG User Plane’s forwarding classes.

> **NOTE**: You can define additional forwarding classes on the BNG User Plane. You do not need to configure these additional forwarding classes on the BNG CUPS Controller.

Juniper BNG CUPS Theory of Operation

SUMMARY

This section describes how Juniper operates and the configurations that you must make to operate Juniper BNG CUPS.

IN THIS SECTION

Operational Overview

The *TR-459 Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, Interface, and Protocol Specifications* (TR-459) document was created by the Broadband Forum to define disaggregated BNG architecture. Figure 4 on page 14 from the TR-459 specification shows the placement of functional blocks on the control plane and the user plane.

Figure 4: TR-459 Functional Separation Between the Control Plane and the User Plane

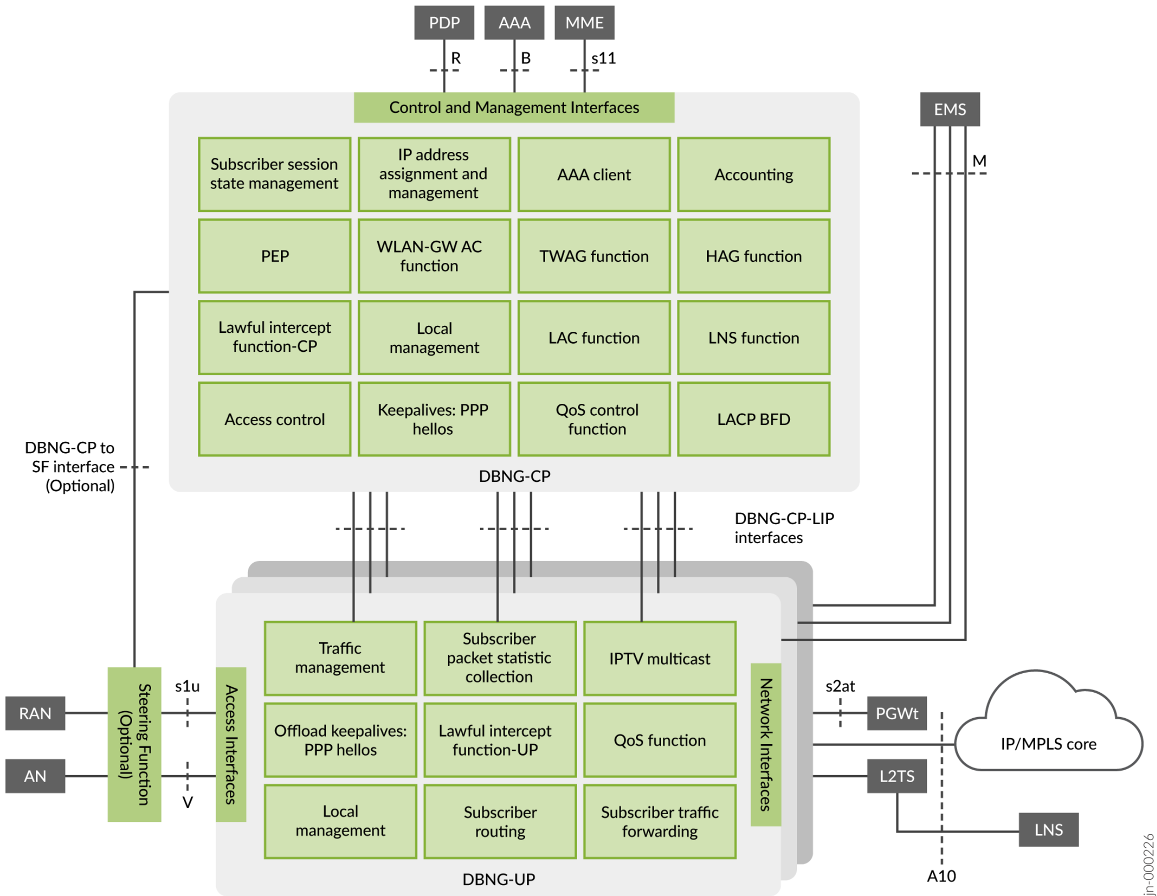

The combination of the control plane functions is referred to as a control plane of the disaggregated BNG. Similarly, a combination of the user plane specific functions is referred to as a user plane of the disaggregated BNG.

Three types of interfaces exist between the control plane and the user plane:

- Management Interface (Mi)—Optionally used for centralized management of the BNG User Planes at the BNG CUPS Controller.
- Control Packet Redirect Interface (CPRi)—Used to direct and exchange control protocol (DHCP, DHCPv6, PPPoE, PPP, L2TP, and so on) traffic between the BNG CUPS Controller and the BNG User Planes to negotiate subscriber sessions.
- State Control Interface (SCi):
  - Used to establish associations between the BNG CUPS Controller and the BNG User Planes.

- Used to program traffic detection and forwarding rules and subscriber state on the BNG User Planes for each subscriber session.
- Used to report session statistics to the BNG CUPS Controller.

The control plane and user plane functions along with the interfaces constitute the disaggregated BNG Architecture as proposed by the TR-459 standard. You can find details in the *TR-459 Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, interface, and Protocol Specifications* document from the Broadband Forum.

Juniper BNG CUPS Controller

The BNG CUPS Controller is a containerized application that runs in a Kubernetes environment. Kubernetes is a container orchestration environment that provides infrastructure to support application and hardware resiliency, automation, application monitoring, application upgrade and rollback, and service discovery.

The BNG CUPS Controller consists of the following micro services:

- Control plane instance—An instance of the subscriber management control plane. The control plane instance manages session states for various access models (for example, DHCP, PPPoE, and L2TP). It also provides AAA services, IP address allocation services, and maintains the SCi and CPRi interfaces to its BNG User Planes. The control plane instance may also interact with a dynamic pool prefix source (Address Pool Manager (remote) or local reserve) to maintain a source of addresses for address allocation. The control plane instance records the session state to the state cache pod. If the control plane instance pod restarts, it recovers its state from the state cache.
- State cache—A persistent in-memory cache that stores subscriber session and other state information generated by the control plane instance. The state cache pod runs on a cluster node other than the node where the control plane instance runs. If the state cache pod restarts, it recovers its state from the control plane instance.

The BNG CUPS Controller components generate log messages through the syslog protocol. You can use the Broadband Edge Event Collection and Visualization (BBE ECAV) application to collect and record the log messages.

Supported Stacking Models

- Juniper BNG CUPS supports the following stacking models:
  - DHCP Server single stack

- DHCPv6 Server single stack
- DHCP Server single session dual stack
- DHCP Relay single stack
- DHCPv6 Relay single stack
- DHCP Relay single session dual stack
- PPPoE single stack (IP or IPv6)
- PPPoE dual stack
- L2TP LAC
- Dynamic VLANs (for DHCP and PPPoE)

Supported Scaling and Topology Requirements

A single BNG CUPS Controller supports the following number of subscribers and BNG User Planes:

- One BNG CUPS Controller can support up to 512K subscribers.
- One BNG CUPS Controller can support up to 32 BNG User Planes.

BNG CUPS Controller runs in a Kubernetes environment.

The Kubernetes environment requires the following devices:

- Control plane node (you must have at least three)
- Worker nodes (you must have at least three)

> **NOTE**: For system requirements, see BNG CUPS Controller Installation Requirements.

Configure BNG CUPS Controller

IN THIS SECTION

The BNG CUPS Controller configuration consists of the following configuration groups:

- `bbe-bng-director`—Contains controller-wide configuration items such as BNG User Plane definitions, control plane instance definitions, BNG User Plane assignments, subscriber and load balancing group definitions.
- `bbe-cpi-0`—Contains the bulk of the subscriber management configurations including the following:
  - Dynamic profiles
  - Class of service classifiers
  - Rewrite rules
  - Traffic control profiles
  - Schedulers and Scheduler maps
  - Firewall filters and policers
  - Authentication, authorization, and accounting (AAA) services at the access and access profile level

Figure 5 on page 18 shows the configuration group hierarchy.

Figure 5: Configuration Group Hierarchy

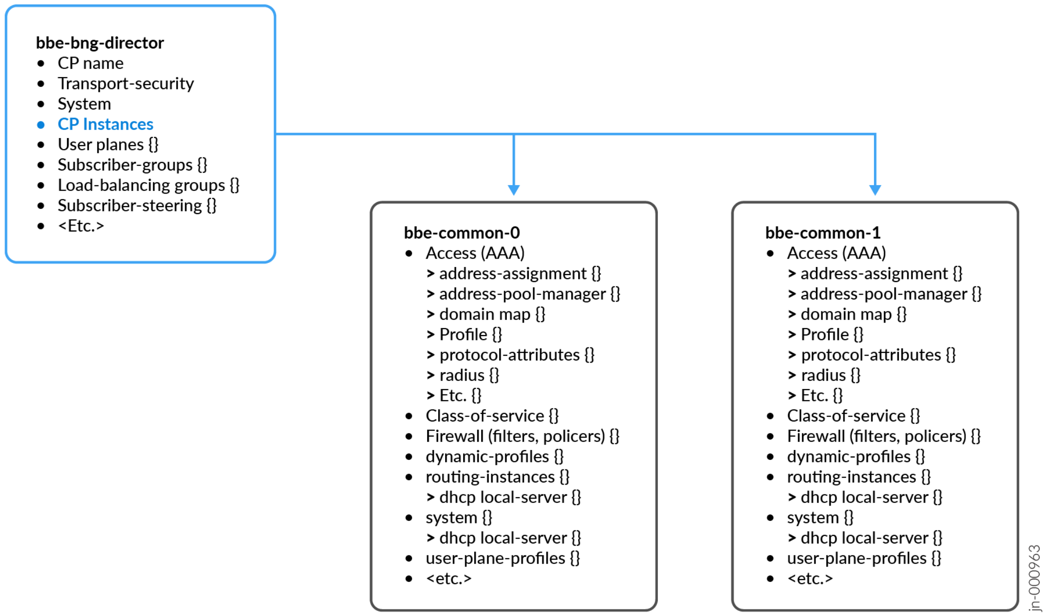

Configure the bbe-bng-director Group

The `bbe-bng-director` configuration group contains the `bng-controller` stanza. You should minimally configure the `bng-controller-name`, `user-planes`, and `control-plane-instances` settings in the `bng-controller` stanza.

See the following `bbe-bng-director` group configuration example:

```
groups {
    bbe-bng-director {
        bng-controller {
            bng-controller-name new-england;
            user-planes {
                billerica {
                    transport {
                        198.20.33.4;
                    }
                    dynamic-address-pools {
                        partition middlesex;
                        v6-na-partition v6-na-partition;
                        v6-dp-partition v6-dp-partition;
                    }
```

```
                    user-plane-profile up-std;
                }
                canton {
                    transport {
                        198.20.48.7;
                    }
                    dynamic-address-pools {
                        partition middlesex;
                        v6-na-partition v6-na-partition;
                        v6-dp-partition v6-dp-partition;
                    }
                    user-plane-profile up-std;

                }
            }
            control-plane-instances {
                cpi-boston {
                    control-plane-config-group bbe-cpi-0;
                    user-plane [billerica canton];
                 }
            }
        }
    }
}
```

In the above example, there are two BNG User Planes defined (billerica and canton). As part of the BNG User Plane configuration, the contact IP address of the BNG User Plane is configured in the `transport` stanza. Dynamic address pool partitions are configured under the `dynamic-address-pools` stanza. Also, the `user-plane-profile`, which defines the BNG User Plane's interfaces and capabilities, is defined and assigned to each BNG User Plane.

The `user-plane-profile` is configured in the common configuration group (for example, `bbe-cpi-0`). So, when the BNG User Plane is configured or assigned to a control plane instance, its user plane profile must be defined in the common configuration group assigned by the `control-plane-config-group` for the control plane instance.

The `bng-controller-name` provides the name of the controller (the collection of control plane instances and the associated microservices).

As part of the control plane instance configuration, you are configuring the following:

- The control plane instance name—The control plane instance name must match the control plane instance name that you assigned to the control plane instance pod created during the `cpi add` configuration in the initial setup of BNG CUPS Controller (see *Juniper BNG CUPS Installation*).

- The name of the configuration group (for example,`bbe-cpi-0`) to use for subscriber management configuration.
- The list of BNG User Planes assigned to the control plane instance.

Configure the bbe-cpi-0 Group

The common configurations for subscriber management are configured in a common configuration group. Up to five common configuration groups can be defined. The name of the common configuration group is fixed. The name must be one of the following: `bbe-cpi-0`, `bbe-cpi-1`, `bbe-cpi-2`, `bbe-cpi-3`, or `bbe-cpi-4`.

See the following `bbe-cpi-0` group configuration example (for simplicity, the example is only partially elaborated) :

```
[edit]
groups bbe-cpi-0 {
    system {
        services {
            dhcp-local-server {
                dhcpv4 {
                    group dhcp-v4-client {
                        dynamic-profile dhcp-client-demux;
                        interface-tag access001;
                    }
                }
            }
        }
    }
    access-profile acc001;
    access {
        address-pool-manager {
            inet 198.19.224.134;
            port 20557;
            local-reserve {
                partition v6-na-partition {
                    family {
                        inet6 {
                            prefix 173:162:1::/96;
                        }
                    }
                }
                partition v6-dp-partition {
```

```
                    family {
                        inet6 {
                            prefix 3000::/8;
                        }
                    }
                }
            }
        }
        radius-server {/* not elaborated */}
        profile acc001 {/* not elaborated */}
        address-assignment {
            domain-profile v4pool {
                family {
                    inet {
                        preferred-prefix-length 24;
                        excluded-address last-octet 255;
                        dhcp-gateway-address-last-octet 1;
                        install-discard-routes {
                            tag 77;
                            backup-tag 88;
                    }
                }
            }
            domain-profile dpPool {
                family {
                    inet6 {
                        partition-type delegated-prefix;
                        preferred-prefix-length 48;
                        allocation-length 56;
                        install-discard-routes {
                            tag 77;
                            backup-tag 88;
                        }
                    }
                }
            }
            domain-profile naPool {
                family {
                    inet6 {
                        partition-type non-temporary-address;
                        preferred-prefix-length 120;
                        allocation-length 128;
                        install-discard-routes {
```

```
                        tag 55;
                        backup-tag 66;
                    }
                }
            }
        }
    }
}
user-plane-profiles {
    up-std {
        interfaces xe-1/1/0 {
            interface-tag access001;
            auto-configure {
                stacked-vlan-ranges {
                    dynamic-profile dhcp-server-demux {
                        accept [ dhcp-v4 dhcp-v6 ];
                        ranges {
                            any,any;
                        }
                    }
                }
                remove-when-no-subscribers;
            }
        }
    }
}
dynamic-profiles {
    dhcp-client-demux {/* not elaborated */}
}
}
```

In this common group configuration, the `dhcp-local-server group` references an interface by its tagged name. An interface tag is defined in the `user-plane-profile` configuration. This allows the same DHCP server group configuration to be used for all BNG User Plane logical ports assigned to the same interface tag.

A user plane profile is a template that is used for a BNG User Plane's interface configuration and other configuration such as lawful intercept, captive portal content delivery, resource monitor, and so on. It is assumed that most of your BNG User Planes will have similar configurations. The user plane profile allows you to avoid constantly having to repeat the BNG User Plane configuration. The DHCP local server can universally represent a BNG User Plane's interface by its tag name (instead of, **up:billerica:xe-1/1/0**). The combination of the tag name and the BNG User Plane context (provided by the BNG CUPS infrastructure) is sufficient enough to identify the interface to the DHCP local server component. This also allows the configuration to avoid specifying the interface for each logical port for each BNG User Plane to be assigned to the DHCP local server group. The same interface tag can be assigned to each logical interface and referenced once in the DHCP local server group.

The common group configuration also includes configurations for Address Pool Manager (APM). In this case, a remote APM instance is used for IPv4 partitions and a local reserve is defined for local IPv6 partitions used to source prefixes for IPv6 non-temporary addresses and delegated prefixes.

Apply the bbe-cpi-0 Group Configuration

After you complete the `bbe-cpi-0` group configuration, you must apply the configuration for it to take affect.

See the following `bbe-cpi-0` group configuration apply example:

```
[edit]
groups {
    bbe-bng-director {}
    bbe-cpi-0 {}
}
apply-groups [ bbe-bng-director bbe-cpi-0 ];
```

Securing BNG CUPS Interfaces

The BNG CUPS Controller interfaces to Address Pool Manager (APM), BNG User Planes, and the observer (multiple geographic deployments only, see "Use Juniper BNG CUPS with Multiple Geographical Redundancy" on page 33) are secured through TLS (TCP-based interfaces) or DTLS (UDP-based interfaces).

Keys and Certificates

To secure interfaces with Transport Level Security (TLS) or Datagram Transport Level Security (DTLS), you must provide the following:

- Private key—The key is used to facilitate exchange of session keys and for encrypting session data.
- Certificate Authority (CA) Certificate—A root certificate (root CA) is used to sign a certificate or public key. The root certificate is also used to authenticate the peer’s certificate.
- Certificate—A signed public key with metadata about the certificated holder and the certificate itself.

Keys and certificates are generated using a variety of available tools (for example, openssl). Certificates may be self-signed or signed by an external Certificate Authority. It is important that the certificates that are used between the BNG CUPS Controller and its external peers (APM, BNG User Planes, and so on) can be authenticated using the root certificate (root CA) that you provide.

Specifying Keys and Certificates to BNG CUPS Controller

To configure keys and certificates in the BNG CUPS Controller, the keys and certificates must be made available to the BNG CUPS Controller's microservices. Keys and certificates are passed by file reference during the setup step of the BNG CUPS Controller installation.

During the setup step you are prompted for the following key and certificate references.

NOTE: If a Kubernetes secret exists in the jnpr-bng-controller namespace, you will not be prompted for the corresponding file references.

```
APM TLS secret name (remove) > K8sSecretName
APM certificate (default: ) > BNGControllerCertificateFile
APM private key (default: ) > BNGControllerKeyFile
APM root certificate (default: ) > BNGControllerRootCertificateFile

BBE observer connection TLS secret name (remove) > K8sSecretName
BBE observer connection certificate (default: ) > BNGControllerCertificateFile
BBE observer connection private key (default: ) > BNGControllerKeyFile
BBE observer connection root certificate (default: ) > BNGControllerRootCertificateFile

DTLS secret name (remove) > K8sSecretName
DTLS certificate (default: ) > BNGControllerCertificateFile
DTLS private key (default: ) > BNGControllerKeyFile
DTLS root certificate (default: ) > BNGControllerRootCertificateFile
```

You can use the same set of keys and certificates for all three interfaces (APM, BNG User Planes, and the observer).

Configuring Keys and Certificates in BNG CUPS Controller

Once keys and certificates are specified during the setup step of the BNG CUPS Controller installation process, they are available for configuration in the BNG CUPS Controller under the **/config** directory path.

APMi Keys and Certificates

APMi keys and certificates are mounted at **/config/apmi** as **ca.crt** (**BNGControllerRootCertificateFile**), **tls.crt** (**BNGControllerCertificateFile**) and **tls.key** (**BNGControllerKeyFile**) files. To secure the BNG

Controller’s interface to APM with TLS, create the following configuration using fully qualified path names:

```
[edit access address-pool-manager]
secrets {
    ca-cert apmi/ca.crt;
    certificate apmi/tls.crt;
    key tls.key;
}
```

SCi and CPRi Keys and Certificates

SCi and CPRi keys and certificates are mounted at the **/config/dtls** location as **ca.crt** (**BNGControllerRootCertificateFile**), **dtls.crt** (**BNGControllerCertificateFile**) and **dtls.key** (**BNGControllerKeyFile**) files. To secure the interfaces to the BNG User Plane (SCi and CPRi) with DTLS, perform the following configuration using file path names relative to **/config**:

```
[edit groups bbe-bng-director]
bng-controller {
    user-planes {
        bng-user-plane-name {
            transport {
                 security-profile security-profile-name;
            }
        }
    }
    security-profiles {
        security-profile-name {
            ca-cert-file-name dtls/ca.crt;
            cert-file-name dtls/dtls.crt;
            key-file-name dtls/dtls.key;
        }
    }
}
```

Configuring Keys and Certificates in the BNG User Plane

Key and certificate files must be copied to the BNG User Plane's file system in the **/config** directory. Once the key and certificate files are transferred to a nonvolatile directory in the BNG User Plane's file system, configure the DTLS key and certificates for the SCi and CPRi interfaces. Perform the following configuration using file path names relative to **/config**:

NOTE: For example of the file path name, if the DTLS key was stored in **/config/up/mx2009-a.key**, the key would be configured as: `key-file-name up/mx2009-a.key`.

```
[edit system services subscriber-management mode user-plane]
security-profiles {
    security-profile-name {
        ca-cert-file-name file-path-to-root-ca;
        cert-file-name file-path-to-cert;
        key-file-name file-path-to-key;
    }
}
transport {
    security-profile security-profile-name;
}
```

Configure BNG User Planes

The BNG User Plane is responsible for applying the subscriber session state originated by the BNG CUPS Controller and acting as the forwarding plane for subscriber traffic. Also, it is responsible for redirecting control protocol packets to the BNG CUPS Controller to negotiate and configure the subscriber session.

The BNG User Plane configuration for subscriber management is a simpler configuration, because most of the configurations for subscriber management are done on the BNG CUPS Controller.

See the following BNG User Plane configuration example:

```
configuration-database {
        max-db-size 419430400;
    }
    subscriber-management {
        enable;
        mode {
            user-plane {
                user-plane-name billerica;
                transport {
```

```
                    inet 198.19.20.33;
                }
                control-plane {
                    bng-controller-name new-england;
                }
            }
        }
    }
```

NOTE: Also, you will need to perform a similar configuration for the BNG User Plane *canton*.

The `user-plane` mode configuration is performed under the `subscriber-management` stanza. The IP address that the BNG User Plane uses to communicate with the BNG CUPS Controller is defined under the `transport` stanza. The BNG CUPS Controller name that the BNG User Plane has been assigned to, and will accept associations from, is defined under the `control-plane` stanza.

The rest of the BNG User Plane's configuration should be focused on other system configurations (for example, interface, telemetry, routing, DDoS protections, resource monitoring, and so on).

Completing Your BNG CUPS Controller Deployment

After you complete the BNG CUPS Controller installation process (see Juniper BNG CUPS Installation Guide), only the state cache service is currently running. You can verify this by running the `dbng status` command.

```
$  dbng status --context context-name
scache        1/1    0
```

To complete the deployment of BNG CUPS Controller, you must create a control plane instance. This is required before you configure control plane instances in the `bbe-bng-director` configuration group. You create a control plane instance using the `cpi-add` command.

```
$  sudo -E dbng cpi add -context context-name --version 23.4R2 cpi-boston
```

This creates the control plane instance pod. You can run the `dbng status` command again to verify that the control plane instance was created. In this example, you can see that *cpi-boston* was created.

```
$  dbng status --context context-name
MICROSERVICE  PODS  RESTARTS
cpi-boston    1/1   0
scache        1/1   0
Storage: Healthy
```

The name you assign to the control plane instance must match the name you use in the `bbe-bng-director` group configuration for the control plane instance. Now that the control plane instance is created, you can proceed to configuring the BNG CUPS Controller by entering the CLI.

```
$  dbng cli -context context-name
root@cpi-boston>
```

BNG CUPS Controller and BNG User Plane Protocol Operations

Now with the BNG User Planes (*billerica* and *canton*) both configured and the assigned control plane instance (*cpi-boston*), the BNG CUPS Controller and BNG User Planes form a disaggregated BNG system by signaling over the state control interface. The signaling consists of PFCP message exchanges to establish an association between the BNG CUPS Controller and each BNG User Plane assigned to it. The signaling also includes additional PFCP message exchanges before subscriber session negotiation and signaling proceeds.

The following diagram shows the initial PFCP message exchanges between the BNG CUPS Controller and each assigned BNG User Plane.

Figure 6: PFCP Message Exchanges

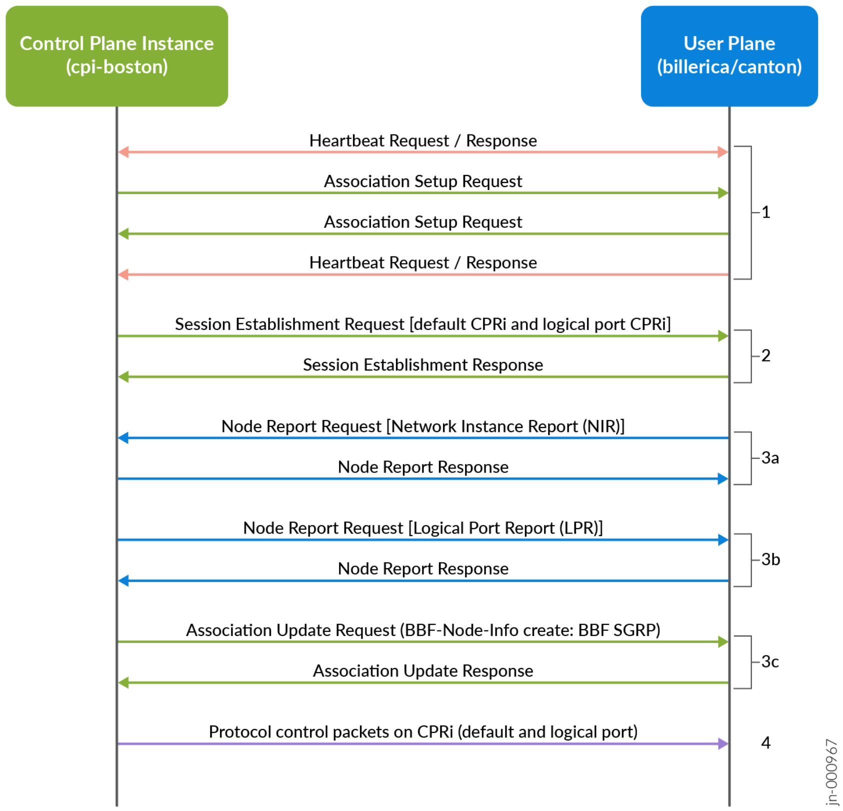

The initial PFCP exchanges occur in three basic steps before the BNG User Plane initiates the forwarding of control protocol packets (for example, PPPoE, PADI, DHCP DISCOVER, DHCPv6, SOLICIT) to the BNG CUPS Controller in step 4 below.

Each BNG User Plane that is assigned to a BNG CUPS Controller, goes through the following steps.

1. - The BNG CUPS Controller initiates the heartbeat request to the BNG User Plane. The BNG User Plane responds to the heartbeat request and initiates its own heartbeat request to the BNG CUPS Controller.

   - The BNG CUPS Controller initiates an association to the BNG User Plane with an association setup request. The BNG User Plane does not initiate a PFCP association and waits to be contacted by the BNG CUPS Controller. If the request is from the configured `bng-controller-name`, the BNG User Plane responds with a PFCP association setup response. A BNG CUPS Controller association is then formed with the BNG User Plane.

   - Heartbeat messages are sent bi-directionally between the BNG CUPS Controller and The BNG User Plane periodically based on the configured interval. It is recommended that the BNG CUPS Controller and the BNG User Planes use the same interval and retry configuration.

The `show health user-plane` command can be performed from the BNG CUPS Controller to confirm a successful association with each assigned BNG User Plane.

```
user@host> show health user-plane
Name        Address       CPi           State        Health     Up-time     Active/Backup-sess
billerica   198.20.33.4   cpi-boston    connected    healthy    00:03:07    0/0
canton      198.20.48.7   cpi-boston    connected    healthy    00:00:18    0/0
```

2. The BNG CUPS Controller initiates session establishment request exchanges to configure the following CPRi tunnels:

   - The default CPRi to allow forwarding of control packets from the BNG User Plane to the BNG CUPS Controller to start subscriber session negotiations.

   - If the user plane profile assigned to the BNG User Plane in the BNG CUPS Controller configuration specifies interfaces configured for auto-sensed VLANs, a logical port CPRi is created for each interface configured for the auto-sensed VLANs. A session establishment request is initiated for each interface and includes both the logical port name and the VLAN ranges from the `auto-configure` stanza for the interface. The logical port CPRi is used to support delayed session creation and thus the exchange of control protocol packets between the BNG User Plane and BNG CUPS Controller to negotiate subscriber sessions.

3. This step consists of three sub-steps. The sub-steps can occur in any order but are expected to occur before subscriber session negotiation is performed:

   a. The BNG User Plane initiates one or more node-level network instance reports. The reports shows each configured network instance and its initial connectivity status (connected or isolated). This action is performed in accordance with TR-459.

   b. The BNG User Plane initiates one or more node-level logical port reports. The reports show each access-facing logical port and its initial forwarding capacity. This action is performed in accordance with TR-459.

   c. The BNG CUPS Controller initiates one or more association update request exchanges to create one or more provisioned subscriber groups. The assigned logical port from the BNG User Plane is included in the subscriber groups creation message.

4. The BNG CUPS Controller receives control protocol packets from the BNG User Plane over the default or logical-port CPRi. Subscriber session negotiation commences based on control packet exchanges between the BNG User Plane and BNG CUPS Controller, resulting in BNG CUPS Controller initiated session establishment requests to create a subscriber session CPRi.

   Note the following:

- The BNG User Plane does not forward received control protocol packets arriving from an access-facing logical port to the BNG CUPS Controller until a node-level subscriber group creation request for the logical port has been received from the BNG CUPS Controller.
- The BNG CUPS Controller discards received control packets arriving on the CPRi until the association update response to create or modify the corresponding subscriber group for the logical port is received from the BNG User Plane.

2

CHAPTER

Use Juniper BNG CUPS

IN THIS CHAPTER

Use Juniper BNG CUPS with Multiple Geographical Redundancy

In a disaggregated BNG model, the control plane provides services to many user planes and their associated subscribers. The control plane's role in the disaggregated model enables new use cases, services, and levels of network redundancy. The use cases require that the control plane has new levels of redundancy. Moving the control plane into a cloud backed by a Kubernetes cluster enables these redundancies.

Kubernetes brings scalability, operational efficiency, and reliability to the solution. The modularity of a Kubernetes cloud enables cluster architectures to have unparalleled redundancy. But, even the most redundant cluster architectures are susceptible to events such as natural disasters or cyberattacks which might target a specific location or geography. A multiple geographic, multiple cluster setup mitigates these susceptibilities.

Figure 7 on page 33 shows an example of a multiple geographic, multiple cluster setup.

Figure 7: Multiple Geographies with Multiple Cluster Setup

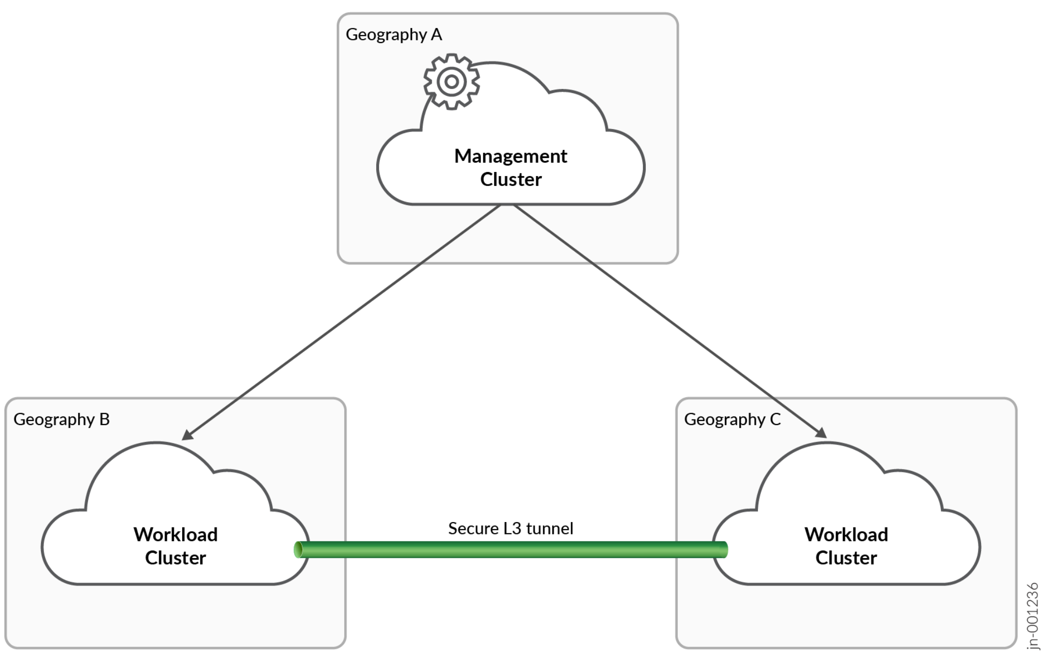

In a multiple geographic, multiple cluster setup, the management cluster maintains a separate context for running multiple cluster scheduling and monitoring functions and is connected to both workload clusters. The multiple cluster context is driven by a policy engine that informs the scheduler how to distribute the application across the workload clusters. Applications that use the multiple cluster setup for multiple geographical redundancy, have policy rules that distribute application microservices involved in state replication to both workload clusters. Other application microservices are distributed to one workload cluster that is chosen as the primary workload cluster.

The workload clusters accept work from the management cluster through the Kubernetes REST API. The workload clusters are standard Kubernetes clusters. A secure L3 tunnel is maintained between the workload clusters. The tunnel facilitates the exchange of application state and general communication between the two workload clusters. As a standard Kubernetes cluster, a workload cluster monitors pods and deployments and performs scheduling tasks for the worker nodes in the cluster, maintaining the deployed application components. The workload cluster does not require the presence of the management cluster to maintain its application workloads. When applications are deployed, it is the workload cluster's responsibility to maintain the application deployment.

If the management cluster detects that a workload cluster has failed or that an application's microservice cannot be satisfactorily scheduled on a workload cluster, the management cluster drives a switchover event. The switchover action is controlled by the policies that are defined for the application. In a switchover event, any application's microservices that exist only on the failed workload cluster are redeployed to the other workload cluster.

The BNG CUPS Controller can be deployed in a multiple geographic, multiple cluster environment. The BNG Controller’s Helm charts include propagation policy rules that instruct the management cluster’s multiple cluster context to deploy an instance of the state cache microservice on both workload clusters. The two state cache instances communicate over the secure tunnel to mirror the subscriber state between the two geographies. The control plane instance is deployed only on one workload cluster. The control plane instance mirrors its state to its local state cache instance, which is then replicated to its state cache peer in the other workload cluster.

Figure 8 on page 35 Shows a BNG CUPS Controller in a multiple geographic, multiple cluster setup.

Figure 8: BNG CUPS Controller in a Multiple Geographic and Multiple Cluster Setup

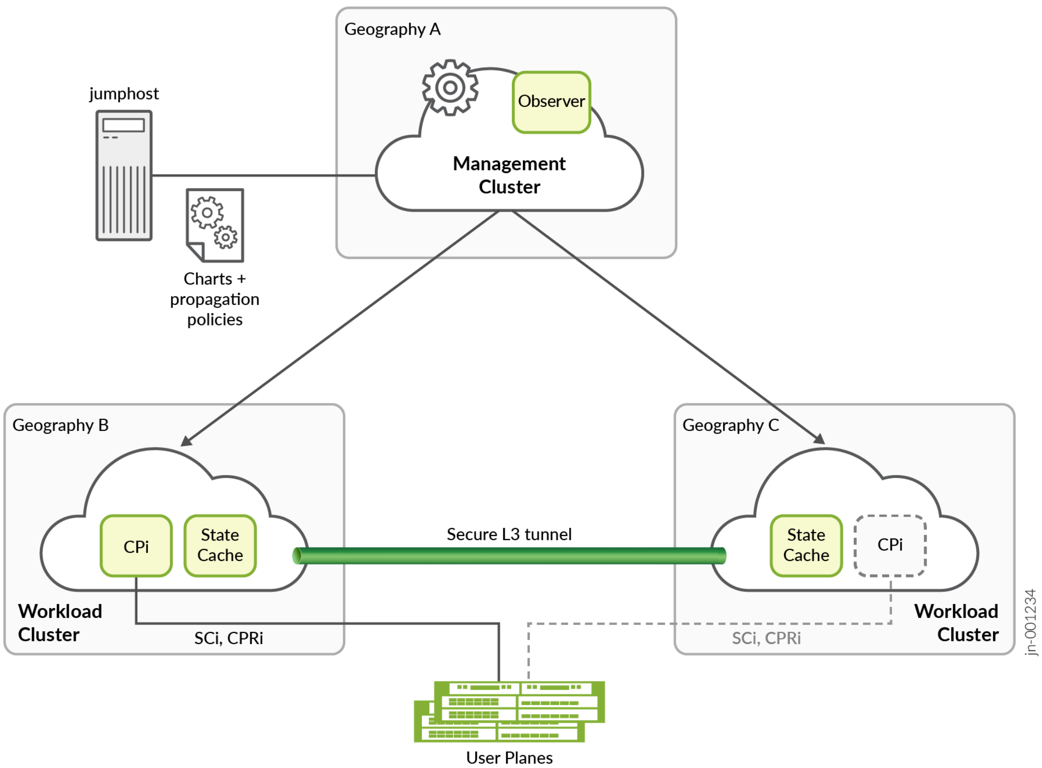

If the workload cluster, where the control plane instance is deployed fails, the management cluster reschedules the control plane instance on the other workload cluster. When the control plane instance initializes on the second workload cluster, it recovers its configuration from a replicated configuration cache (not shown). The control plane instance also recovers its subscriber state from the local state cache instance just as it would on any microservice restart. Since the local state cache received replication state information from the previous workload cluster, all stable states are recovered. Once the state is recovered, the control plane instance establishes its associations with the BNG User Planes. The BNG User Plane association logic detects that the new association is originating from the same BNG CUPS Controller (but in a different geography).

The BNG CUPS Controller also deploys a microservice on the management cluster that is called the observer. The observer runs in the regular context of the management cluster and watches control plane instance scheduling events in the multiple cluster context (associated with the management cluster). In switchover situations where a control plane instance might exist temporarily in both workload clusters, the observer allows the BNG CUPS Controller to resolve any ambiguity over which control plane instance should be running.

For example, in the event that the management cluster can no longer reach or monitor a workload cluster, the management cluster declares that the workload cluster has failed. The management cluster initiates a switchover of workloads from the failed workload cluster to the other workload cluster.

A scenario where the failed workload cluster is operational, but not reachable from the management cluster creates an ambiguous deployment. Multiple instances of application workloads exist on both workload clusters. Since the management cluster is the final arbiter on where workloads run in a multiple cluster deployment, a mechanism is needed to force the duplicate workloads on the workload cluster that has perceived to have failed to enter an inactive or dormant state in deference to their switched over counterparts.

As previously mentioned, the observer microservice runs on the management cluster. The observer watches for scheduling events for application workloads. Each time an application workload is scheduled on a workload cluster, the observer assigns a unique generation number to that workload. When the same workload is switched over to the other workload cluster, the generation number is incremented. As the applications workloads initialize, they request their generation number from the observer. The generation number passes between the workload clusters. Application workloads on the failed workload cluster take note that the same workloads have a higher generation number on the other workload cluster and transition the application into a dormant state (all connections are dropped, no state is generated or consumed).

The generation number helps correct the ambiguous deployment caused by the management cluster's inability to view the true state of the failed workload cluster. When reachability is restored to the failed workload cluster, the management cluster removes the dormant application workloads.

When a cluster switchover occurs, the set of external IP addresses used to contact the BNG CUPS Controller change to the address space used in the new workload cluster. Some back-office applications are more sensitive to IP address changes. For instance, RADIUS wants the value of the NAS-IP-Address to remain constant across cluster switchovers.

The BNG CUPS Controller supports a consistent value for the NAS-IP-Address through a microservice called the Multi-geo IP Route Prioritization Operator (MIRP). The Multi-geo IP Route Prioritization Operator is a microservice that is deployed to each workload cluster.

When a CPi is added, you can provide a persistent external address to use for the CPi’s RADIUS listener address by adding the `ip-aaa` *`external-radius-address`* option to the `cpi add` command.

When the CPi microservice is deployed, a custom resource is created for the MetalLB IP address pool containing the configured external address. A separate load balancer service is created for the RADIUS listener port with MetalLB annotations to reference the IP address pool. This ensures that the service is assigned the listed external address.

The Multi-geo IP Route Prioritization Operator in the workload cluster, where the CPi’s RADIUS load balancer service is created, generates a BGP advertisement custom resource. This resource contains a `LOCAL_PREF` value derived from the CPi’s generation number as obtained by the observer (the generation number is incremented by the observer when a cluster switchover event is observed). The BGP advertisement is exchanged with the network load balancer's IBGP peers. The back-office systems (RADIUS in this case) sit behind the IBGP peers. When the CPi switches over to the other workload cluster, its generation number increments. The associated `LOCAL_PREF` value of the Multi-geo IP Route Prioritization Operator-generated BGP advertisement is also incremented. Traffic bound for the service's external IP address is routed to the cluster with the higher `LOCAL_PREF` value which is also the active workload cluster for the CPi.

Use Dynamic Address Pools in Juniper BNG CUPS

IN THIS SECTION

BNG User Plane high availability within Juniper BNG CUPS is based on subscriber groups. Each subscriber group tracks its own set of subscriber prefixes to successfully switchover all session states, including pool prefix routes, to a backup BNG User Plane. Therefore, subscriber groups are allocated their own set of pool prefixes. Rather than pre-provisioning a set of pools for each subscriber group, a dynamic prefix source is used.

Dynamic prefix sources used in the BNG CUPS Controller include the following:

- Address Pool Manager (APM)—APM is a cloud-native application that maintains a set of prefix partitions from which sub-prefixes may be apportioned for use as pool prefixes. APM communicates with the BNG CUPS Controller's CPi through the APMi, a gRPC-based protocol. Currently, APM serves only IPv4 prefixes.

- Local reserve—Local reserve is a BNG CUPS Controller configured set of prefix partitions from which sub-prefixes may be apportioned for use as pool prefixes. Local reserve serves both IPv4 and IPv6 prefixes. Local reserve can also act as a backup prefix source for APM when the APMi is disconnected. Currently a local reserve must be used for IPv6 prefixes to assign IPv6 non-temporary addresses, delegated prefixes, and router advertisement prefixes.

As part of the BNG CUPS Controller's configuration of the BNG User Planes, the `dynamic-address-pools` stanza defines the source partition names from which pool prefixes are apportioned and from which they will be reclaimed.

Following are the four types of partitions:

- partition—IPv4 Partition name
- v6-dp-partition—IPv6 delegated prefix partition name
- v6-na-partition—IPv6 non-temporary address partition name
- v6-ra-partition—IPv6 route advertisement partition name

Local Reserve

The local reserve is a BNG CUPS Controller configured set of partitions. Partitions can be either IPv4 or IPv6. Local reserve partitions are configured under the `access address-pool-manager` stanza. See the following example:

```
access {
    address-pool-manager {
        inet 198.19.224.134;
        port 20557;
        auto-recovery drain-delay 120;
        apportion-delay 60;
        local-reserve {
            partition middlesex {
                family {
                    inet {
                        prefix 192.168.192.0/20;
                    }
                }
            }
            partition v6-na-partition {
                family {
                    inet6 {
                        prefix 173:162:1::/96;
                    }
                }
            }
            partition v6-dp-partition {
                family {
                    inet6 {
```

```
                            prefix 3000::/8;
                        }
                    }
                }
            }
        }
```

The IPv4 partition (*middlesex* in this example) is a backup partition for a partition of the same name in the APM configuration. In this case the `local-reserve` partition has one prefix of private addresses. If the connection to APM is lost, the subscriber groups associated with the BNG User Planes that have specified *middlesex* as their IPv4 partition apportion private prefixes from the local reserve after the APMi connection has been down for the configured apportion delay time. Once the APM apportioned public pool prefixes are exhausted, incoming subscribers in the subscriber group are allocated addresses from the private pool prefixes sourced from the local reserve.

Subscribers may have limited access with private addresses, but they will be able to login to the network. Once the APMi connection is restored, it is desirable to readdress the subscribers who were allocated private addresses with public addresses from APM-sourced pools. After the configured auto recovery drain delay period, the BNG CUPS Controller enables an active drain on the pools apportioned from the local reserve. As subscribers reconnect, additional public pool prefixes are apportioned from APM and the subscribers are allocated public addresses and regain full service.

There are also two IPv6 partitions configured as part of the local reserve. These partitions apportion IPv6 pool prefixes for non-temporary addresses and prefix delegated addresses for IPv6 subscribers respectively. Since APM does not support IPv6 partitions, local reserve is the only option to source dynamic address pools for subscriber groups serving IPv6 subscribers.

Address Pool Manager

APM is a separate cloud-native application that can be deployed in the same Kubernetes cluster as the BNG CUPS Controller or in a different cluster altogether. APM can source IPv4 partitions for many BNG CUPS Controller control plane instances or integrated BNGs.

See the following APM configuration example:

```
apm {
    inet-pool {
        partition middlesex {
            prefix 192.32.0.0/16 {
                max-prefix-length 24;
            }
```

```
        }
    }
    entity-match cpi-massachusetts {
        pool-domain-profile domainTemplate;
    }
    pool-domain-profile domainTemplate {
        monitoring {
            apportion-threshold 200;
            reclaim-threshold 457;
        }
        auto-reclamation {
            active always;
        }
    }
}
```

In the APM configuration, partition *middlesex* has a public IPv4 prefix from which pool prefixes are apportioned and reclaimed.

The `entity-match` stanza identifies the CPis that APM will accept connections from. In this case, only CPi with the systemID of *cpi-massachusetts* will be allowed to connect. The CPI uses apportion and reclamation settings for created pool domains as defined by the `pool-domain-profile` *domainTemplate*.

The corresponding BNG Controller configuration elements necessary to use APM as a dynamic prefix source are shown in the following example:

```
groups {
    bbe-cpi-0 {
        access {
            address-pool-manager {
                inet 198.19.224.134;
                port 20557;
            }
            address-assignment {
                domain-profile v4FramedPoolName {
                    family {
                        inet {
                             preferred-prefix-length 24;
                             excluded-address-last-octet 255;
                        }
                    }
                }
            }
```

```
        }
    }
}
```

In the `address-pool-manager` stanza, the `inet` statement contains the external IP address used by APM. This can be retrieved by using the `apm ip` utility script command (see APM User Guide). The default port that APM listens on is 20557. The system identifier that the control plane instance uses to identify itself to APM is the `control-plane-instance` name (for example, *cpi-massachusetts*). APM must have a corresponding *entity-match* entry in its configuration.

In the *address-assignment* stanza, *domain-profiles* must match the `FramedPool` names that are supplied during the subscriber authentication phase and include the preferred prefix length to request pool prefixes from the prefix source (either APM or local reserve) and any address exclusions to use for the apportioned dynamic pools.

Domain Creation, Apportionment and Reclamation

The *domain-profile* statement configured under the BNG CUPS Controller's `access address-assignment` stanza in the `bbe-cpi-0` group aligns with the `address-pool` or `FramedPool` attribute returned during the authentication phase of subscriber login. The domain profile defines the size of the prefix to apportion from the partition, any address exclusions, and whether to install a discard route for each pool prefix.

See the following `domain-profile` example configuration:

```
domain-profile v4pool {
    family {
        inet {
            preferred-prefix-length 24;
            excluded-address last-octet 255;
            install-discard-routes {
                tag 77;
                backup-tag 88;
            }
        }
    }
}
domain-profile dpPool {
    family {
        inet6 {
```

```
            partition-type delegated-prefix;
            preferred-prefix-length 48;
            allocation-length 56;
            install-discard-routes {
                tag 77;
                backup-tag 88;
            }
        }
    }
}
```

As a subscriber logs into the network, a `FramedPool` attribute is returned from a successful authentication phase. If the `FramedPool` matches a `domain-profile` in the configuration, the CPi checks to see if a domain has been created for the associated subscriber group. If no domain exists, the CPi coordinates with the partition source (either APM or the local reserve) to create a domain name by connecting the values of the `FramedPool` name, the subscriber group name, and the associated routing instance.

Once the domain is created, the CPi raises an apportion request with the partition source to stock the domain with pool prefixes. As more subscribers associate with the subscriber group during login, the CPi apportions more pool prefixes when the number of available addresses in the domain drops below the domain's apportion threshold. Similarly, when the number of available addresses rises above the domain's reclamation threshold, the CPi raises a reclamation request with the partition source to return pool prefixes to the partition until the available addresses drops below the reclamation threshold. When all prefixes in the domain are reclaimed, the domain itself is cleaned up.

Juniper BNG CUPS High Availability

There are two aspects to Juniper BNG CUPS high availability, high availability for the BNG CUPS controller and high availability for the BNG User Planes.

BNG CUPS Controller High Availability

The BNG CUPS Controller consists of two micro services which run as pods on a Kubernetes cluster. The State Cache pod backs up all sessions and the BNG CUPS Controller state in high-availability mode.

In the case of a BNG CUPS Controller container failure, Kubernetes creates a new BNG CUPS Controller container. The BNG CUPS Controller gets its information from the State Cache container and builds a new state. After creating all the states, the BNG CUPS Controller reconnects to the BNG User Planes and continues from where it left off. BNG User Planes continue to forward traffic during a BNG CUPS Controller failure. No new logins are allowed until the BNG CUPS Controller recovers.

Figure 9 on page 43 shows the BNG CUPS Controller container and the State Cache container.

Figure 9: BNG CUPS Controller High Availability

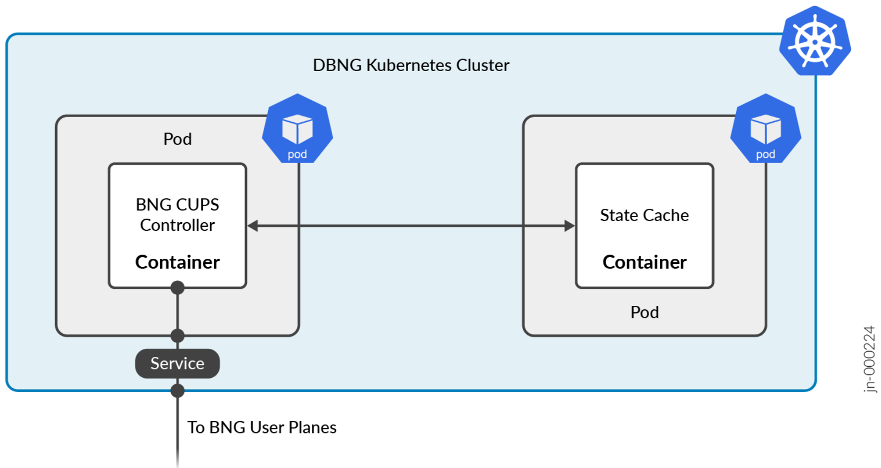

BNG User Plane High Availability

High availability between BNG User Plane's routing Engines, also known as Graceful Routing Engine Switchover (GRES) is used in conjunction with BNG CUPS subscriber resiliency. On GRES, State and other information is replicated in a high availability mode across the routing engines. During GRES, the standby routing engine takes over as the active routing engine immediately.

NOTE: For more information regard BNG User Plane high availability, see "Use Juniper BNG CUPS Subscriber Groups" on page 44.

Figure 10 on page 44 shows the BNG User Plane high availability andwith GRES support between RE0 and RE1.

Figure 10: BNG User Plane High Availability and GRES

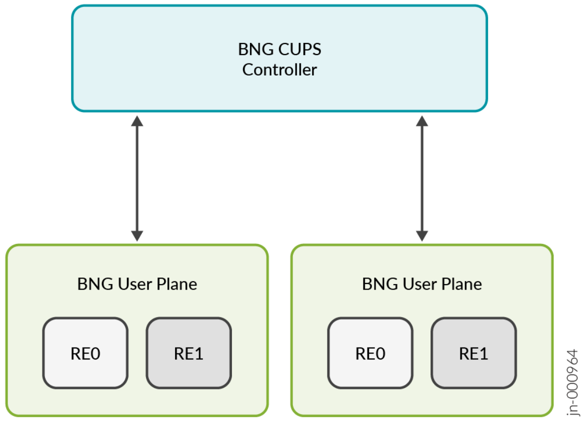

Use Juniper BNG CUPS Subscriber Groups

IN THIS SECTION

Subscriber Groups Overview

The *TR 459 Multi-Service Disaggregated BNG with CUPS, Reference Architecture, Deployment Models, Interface, and Protocol* (TR-459) specification explains how the disaggregation of the BNG improves subscriber resilience. This is due to the fact that the disaggregated control plane (in this case, the BNG CUPS Controller), contains a centralized master state database for all of the disaggregated user planes (BNG User Planes) under its control. Figure 11 on page 45 (from TR-459) shows how the BNG User Planes provide resilience across the BNG User Planes where the BNG Cups Controller holds the master state for any subscriber session. The BNG User plane then holds active forwarding state, or backup forwarding state for a particular subscriber session.

Figure 11: Subscriber Session Resiliency

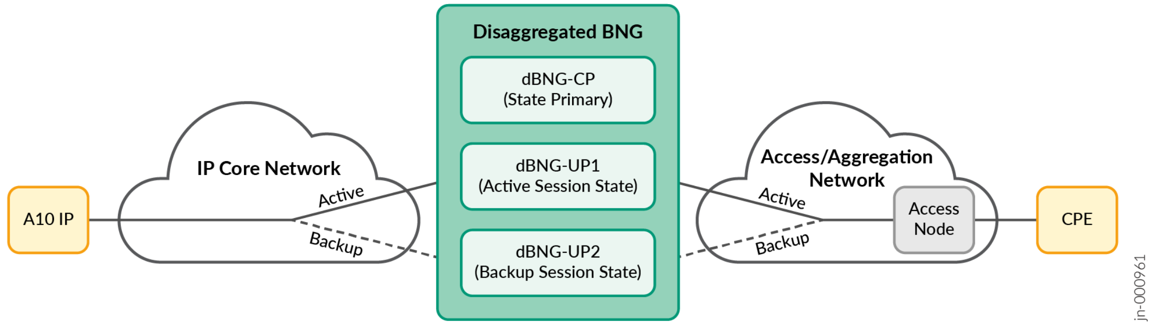

Subscriber sessions that are subject to the same restoration capability are placed into the same subscriber group. Grouping subscribers together helps to increase core routing efficiency.

The use of subscriber groups minimizes the messaging, which reduces the elapsed time between the detection of a failure (or any request to switchover from active to backup) and the restoration of the service.

The *active* or *backup* state is set at the subscriber group level and communicated to the relevant BNG User Plane by the BNG CUPS Controller. Subscriber sessions are tagged with the subscriber group to which they belong when the session is established. All resiliency actions are communicated at the subscriber group level rather than at the session level.

Subscriber groups are created based on the BNG User Plane interfaces.

Resilient subscriber groups have the following characteristics:

- Spans at least two BNG User Planes.
- Contains one or more redundancy interfaces. Redundancy interfaces consist of one interface on each BNG User Plane.
- Is active, only on one BNG User Plane at a time. Subscribers are only serviced by the subscriber's active BNG User Plane. Also, all interfaces in the subscriber group move at the same time.
- BNG User Planes can have more than one subscriber group associated to them.

Subscriber address management operates with subscriber groups in the following ways:

- Subscriber IP addresses must come from its subscriber group's defined address domains (made up of prefixes). Domains are created dynamically based upon, the RADIUS VSA, SGRP name and routing instance.
- Address prefixes are advertised differently on *active* and *backup* BNG User Planes.

Subscriber management switchover consists of the following:

- Switchover can be controlled by either the BNG CUPS Controller or the BNG User Planes.
- Route advertisement metrics are changed during subscriber group switchover.

The address domain prefixes and their associated metrics allow policies to be applied per BNG User Plane. This is so that the routing policy fits within any local variations, and the preferred metric can be applied upon subscriber group switchover.

You configure a subscriber group on the BNG CUPS Controller with the following settings:

- A subscriber group name and a subscriber group identifier (a unique 32bit unsigned integer)
- State—Active, backup, or Track-Logical-Port
- Active BNG User Plane and backup BNG User Plane
- Logical ports, and virtual MAC address
- Prefixes and tags

In the BNG CUPS Controller, the subscriber group is configured with details about the BNG User Planes and the list of logical ports for each of the BNG User Planes. In a BNG CUPS Controller managed subscriber group, the BNG CUPS Controller sends subscriber group notifications to the BNG User Planes with either, an *active* or *backup* state and the respective port list.

In a resilient subscriber group, when the subscriber logs into the subscriber session, services are simultaneously created at both of the BNG User Planes. The services are also tagged with the subscriber group ID. At this time, the subscriber sessions associated with the backup BNG User Plane discards all packets in both directions.

There are two types of subscriber groups, either a BNG CUPS Controller managed subscriber group or a BNG User Plane managed subscriber group.

A BNG User Plane managed subscriber group is a resilient subscriber group with a single logical port pair and its state set to `Track-Logical-Port` (TLP). A TLP or BNG User Plane managed subscriber group (also referred to as subscriber group type TLP) requires that the BNG CUPS Controller set the subscriber group state to `Track-Logical-Port` on both the active and backup BNG User Plane instances. This specifies the logical ports for which the BNG User Planes track the state. The BNG User Plane tracks the operational state of the access network's connectivity on the logical ports. This determines if and when the switchover occurs. The two BNG User Planes that belong to a BNG User Plane managed subscriber group are assumed to be linked by an active to backup connection on the access side. The two BNG User Planes decide, by themselves, which one of the two handles the sessions. This decision is made based on their relevant logical ports.

The following example shows the configuration of a BNG User Plane managed subscriber group on a BNG CUPS Controller:

NOTE: We recommend that you use a BNG User Plane managed subscriber group rather than a BNG CUPS Controller managed subscriber group. We recommend this because, for BNG User Plane managed subscriber groups, the protocols running between the BNG User Planes and the access network establish the active and the backup link in the network connectivity. So, it is the BNG User Plane that tracks the operational state changes or failures of the access network and initiates a subscriber group switchover, directly following the network changes.

```
[edit groups bbe-bng-director bng-controller]
subscriber-group {
   SGRP-TLP {
     virtual-mac aa:bb:01:01:01:01;
     user-plane-managed-mode {
      redundancy-interface GAMMA {
         logical-ports up:boston:ps1,up:nashua:ps3;
      }
     }
     user-plane boston {
       backup-mode hot;
     }
     user-plane nashua {
       backup-mode hot;
     }
  }
}
```

NOTE: The `user-plane-profile` stanza changes and the `bng-controller` stanza changes need to be separated by different configuration **commit** actions. You cannot commit `user-plane-profile` and `bng-controller` stanza changes using a single commit action.

Before configuring a subscriber group, make sure that the `user-plane` logical ports are already configured in the `user-plane-profile`.

A BNG CUPS Controller managed subscriber group is a resilient subscriber group with one or more logical port pairs, where the BNG CUPS Controller only programs active and backup states on the BNG User Planes (also known as subscriber group Type A/B). As a best practice, the BNG CUPS Controller subscriber group should be configured with a single redundancy interface (or a single port pair).

The following example shows the configuration of a BNG CUPS Controller subscriber group:

```
[edit groups bbe-bng-director bng-controller]
subscriber-group {
   SGRP-AB {
     virtual-mac aa:01:01:01:01:01;
     control-plane-managed-mode {
        preferred-user-plane-name jersey;
        redundancy-interface GAMMA {
            logical-ports up:jersey:xe-1/0/0,up:boston:xe-2/0/0;
        }
     }
  }
}
```

Hot backup support ensures that upon switchover, subscriber activity and traffic is unaffected with little or no packet loss.

To check the state of the subscriber group on the BNG CUPS Controller, you can run the `show subscriber-group` command:

```
user@host> show subscriber-group SGRP-AB
Name: SGRP-AB
ID: 5
User-Plane: jersey (active) (hot)
User-Plane: boston (backup) (hot)
Health status: healthy
Mode: Control Plane
VMAC: AA:01:01:01:01:01
Logical port mapping:
  BB device    Name         Logical-port            Sessions  Logical-port             Sessions
  bb0.6        GAMMA        up:jersey:xe-1/0/0      2          up:boston:xe-2/0/0      2
Address domains:
  Name                          Prefixes     User-Plane       Programmed     User-Plane
Programmed
  suburbs:SGP-AB:default        1            jersey            1             boston
1
```

Switchover triggers the use of the subscriber group. Switchover can be split into a BNG CUPS Controller initiated switchover or a BNG User Plane initiated switchover.

You use the `request subscriber-group switchover` command to initiate a BNG CUPS Controller initiated switchover.

```
request subscriber-group switchover SGRP-AB
```

After the BNG CUPS Controller initiated switchover, the BNG User Plane *jersey* is no longer the active BNG User plane, but is now the backup BNG User Plane. See the following `show subscriber-group` command output:

```
user@host> show subscriber-group SGRP-AB
Name: SGRP-AB
ID: 5
User-Plane: boston (active) (hot)
User-Plane: jersey (backup) (hot)
Health status: healthy
Mode: Control Plane
VMAC: AA:01:01:01:01:01
Logical port mapping:
  BB device   Name         Logical-port                Sessions   Logical-port
Sessions
  bb0.6       GAMMA       up:boston:xe-2/0/0   2               up:jersey:xe-1/0/0   2
Address domains:
  Name                                    Prefixes    User-Plane       Programmed    User-
Plane    Programmed
  suburbs:SGP-AB:default            1          jersey                1
boston           1
```

Default Subscriber Groups

A default subscriber group is the subscriber group that is automatically created when a BNG User Plane is associated with the BNG CUPS Controller. If there are no additional BNG User Planes assigned to the default subscriber group, the subscriber group is not resilient.

In the hitless (meaning, subscriber activity and traffic is unaffected with little or no packet loss) maintenance use case, a backup BNG User Plane gets automatically added to the default subscriber group for the BNG User Plane that is being serviced. This action preserves the existing subscriber traffic and state while maintenance is performed on the BNG User Plane. When the maintenance is completed, the backup BNG User Plane is removed from the default subscriber group of the serviced BNG User Plane.

Additional Subscriber Group Information

Each subscriber group has its own address prefixes that do not overlap with other subscriber group's address prefixes.

When a subscriber logs out or a subscriber cleanup is triggered by deleting a subscriber group, the BNG CUPS Controller collects the final statistics from both the backup and active BNG User Planes.

For BNG User Plane managed subscriber group, active and backup pseudowire or EVPN can be used in the Access Network.

BNG CUPS Controller managed switchover and BNG User Plane managed switchover might be mutually exclusive depending on the Access Network technology.

In the hitless maintenance use case, there should be little or no disruption to a subscriber's activity and network traffic should remain uninterrupted while the BNG User Plane is serviced. Hitless maintenance is one of the use cases that uses the BNG CUPS Controller managed subscriber group and the BNG CUPS Controller initiated switchover.

To perform maintenance operations using configured subscriber groups (subscribers and services are already installed in on the subscriber group's backup BNG User Planes) you use BNG CUPS Controller initiated switchover to seamlessly move traffic to the backup BNG User Planes for the subscriber groups that contain the BNG User Plane that is under maintenance.

When maintenance is completed, you then perform a BNG CUPS Controller initiated switchover again and the BNG User Plane that was serviced becomes the active BNG User Plane for the subscriber groups.

There are many switchover triggers that change the active state for a particular subscriber group from one BNG User Plane to a different BNG User Plane:

- Operator driven trigger through the management interface of the BNG CUPS Controller.
- Failure of an entire BNG User Plane.
- Failure of a component of the BNG User Plane that impacts a set of active subscriber sessions.
- The failure of a link or interface directly connected to the BNG User Plane that impacts a logical port and active subscriber sessions.
- A change in the negotiated status of a resilient connection between the BNG User Plane and the Aggregation Network
- A change in the IP core network that isolates a BNG User Plane from the rest of the network.

Subscriber Groups and DHCP Relay

BNG CUPS supports a DHCP relay L3 wholesale model. To support this model, a dedicated routing instance is configured for each retailer and assigned to a subscriber group. In the DHCP relay L3 wholesale model, the retailers DHCP server is responsible for address assignment. The IPv4 and IPv6 prefixes that represent the pool of IPv4 and IPv6 addresses that are assigned to the subscriber by the retailer are configured as part of a subscriber group, using the `prefix-advertisement` command. You can also, configure the `prefix-advertisement` command with optional route tags for the purpose of route advertisement.

The following example shows the configuration of the `prefix-advertisement` command in the subscriber group configuration:

```
[edit groups bbe-bng-director bng-controller]
subscriber-group
subscriber-group-name subscriber-group-name
   routing-instances
        boston {
            prefix-advertisement {
                family inet {
                    tag 300;
                    backup-tag 400;
                    prefix 1.2.0.0/16;
                    prefix 1.3.0.0/16;
                }
                family inet6 {
                    tag 310;
                    backup-tag 410;
                    prefix 2:3::0/64;
                    prefix 2:4::0/64;
                }
            }
        }
    }
```

Supporting a DHCP local server in which there is a DCHP relay in front of the BNG User Plane requires the use of static IP address pools to support the addressing model. The subscriber group allows the assignment of a statically configured address assignment pool, The statically configured address assignment pool is defined outside the SGRP, as follows:

IPv4 pool—`access address-assignment pool` *`pool-name`* `family inet network` *`prefix`*

IPv6 pool—`access address-assignment pool` *`pool-name`* `family inet6 prefix` *`prefix`*

The global name of a static IP address pool is referenced in the subscriber group configuration using the `static-address-pools` command, The network prefix is derived from the address pool and advertised using configured route tags (optional).

The following example shows the configuration of the `static-address-pools` command in the subscriber group configuration:

For a complete description of the `static-address-pools` command, see "static-address-pools" on page 151.

```
[edit groups bbe-bng-director bng-controller]
subscriber-group
subscriber-group-name subscriber-group-name
   routing-instances {
        default {
            static-address-pools {
                pool-1-name {
                    tag 100;
                    backup-tag 200;
                }
                pool-2-name {
                    tag 110;
                }
             }
        }
   }
```

Use BNG User Plane Maintenance

IN THIS SECTION

BNG User Plane Maintenance Overview

Juniper BNG CUPS in accordance with the TR 459 Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, Interface, and Protocol specification introduces a new maintenance (hardware and software maintenance) approach for BNG User Planes. Juniper BNG CUPS enables you to perform maintenance operations on your BNG User Planes without impacting the subscribers' traffic, therefore improving network operations and the subscribers' experience.

BNG User Plane maintenance relies on the BNG User Plane redundancy that is enabled through the BNG CUPS Controller. Instead of triggering a failure in the BNG User Plane, the BNG CUPS Controller assumes an operational procedure is occurring, which can be a maintenance repair.

BNG User Plane maintenance is applied to any session model (DHCP, IPoE, PPPoE and LNS) and assumes an access transport based on pseudowires EVPN-VPWS, with active and standby, or Ethernet with the access node controlled active and standby links to the BNG User Planes.

How to Use BNG User Plane Maintenance

Figure 12 on page 54 illustrates the hitless BNG User Plane maintenance use case. It shows an example of you performing an in service maintenance on a distributed BNG User Plane (*BNG-UP1* in the illustration), without disrupting any live subscriber sessions.

NOTE: meaning, subscriber activity and traffic is unaffected with little or no packet loss

Figure 12: Hitless BNG User Plane Maintenance

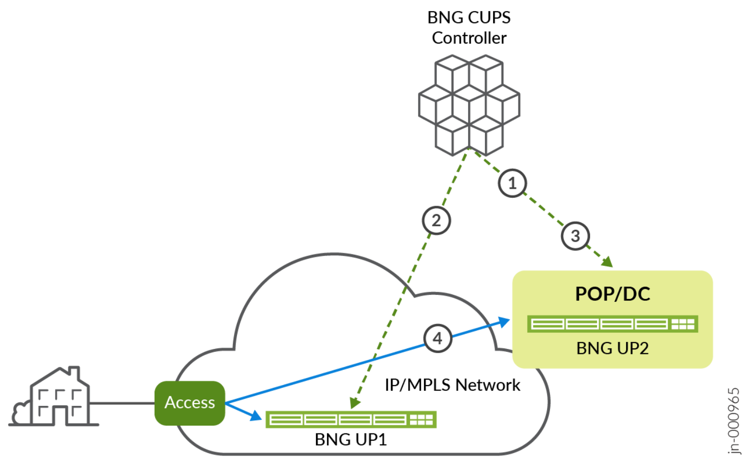

Following are the steps that occur when performing maintenance on a BNG User Plane (see Figure 12 on page 54.

1. Prepare for performing maintenance on your BNG User Plane. You use the BNG CUPS Controller to program an alternate BNG User Plane (*BNG-UP2*) with subscriber state information from BNG User Plane *BNG-UP1*.

2. The BNG CUPS Controller activates BNG User Plane *BNG-UP2*.

3. You enable the access network to start forwarding traffic from the access node to BNG User Plane *BNG-UP2* and the core network.

4. As the hot standby BNG User Plane, BNG User Plane *BNG-UP2* is now preprogrammed with the subscriber state information from BNG User Plane *BNG-UP1*. As subscriber traffic arrives on BNG User Plane *BNG-UP2*, it forwards the subscriber traffic. Once maintenance is complete, you perform the same work flow in reverse to revert traffic back to BNG User Plane *BNG-UP1*.

BNG User Plane Maintenance Process

This section describes the process that is required when you perform a maintenance operation on a BNG User Plane. The procedure refers to Figure 12 on page 54.

1. Create a backup for BNG User Plane *BNG-UP1* on BNG User Plane *BNG-UP2*.

   At this step you must first associate the active BNG User Plane *BNG-UP1* ports with the backup BNG User Plane *BNG-UP2* port. Then synchronize the ports, existing subscribers, and the address domain state from the active BNG User Plane *BNG-UP1* to the backup BNG User Plane *BNG-UP2*. Subscribers that are on BNG User Plane *BNG-UP1* are now also programmed on the backup BNG User Plane *BNG-UP2's* logical ports together with the address domain (prefixes, tags, routing-instances, and so on are also programmed on *BNG-UP2*).

```
user@host# request user-plane maintenance associate serviced-user-plane BNG-UP1 serviced-port port1 backup-user-plane BNG-UP2 backup-port port2
```

2. Setup network to route traffic to the chosen backup BNG User Plane (*BNG-UP2*). This step is provider and operator specific and the actions taken at this step vary greatly with the various access and core network topologies deployed.

   For example, at this step the operator could setup the core network for attracting subscriber traffic to the backup BNG User Plane by setting the routing policy to import these prefixes. Also, it is expected that at this step the operator is done setting up the access network for the backup BNG User Plane.

3. Make the backup BNG User Plane *BNG-UP2* the active BNG User Plane. Now subscribers are programmed on the BNG User Plane *BNG-UP2* and it is ready to take over.

   This step executes the switchover from BNG User Plane *BNG-UP1* to BNG User Plane *BNG-UP2*. After completing this step, BNG User Plane *BNG-UP2* is the active BNG User Plane and BNG User Plane *BNG-UP1* is the backup.

```
user@host# request user-plane maintenance switchover serviced-user-plane BNG-UP2
```

4. Perform the require maintenance on BNG User Plane *BNG-UP1*. The service can be various activities, such as servicing a line card or a software upgrade.

5. Restore the subscribers on BNG User Plane *BNG-UP1* as backup. At this step the ports, subscribers, and domain state are automatically synchronized from the active BNG User Plane *BNG-UP2* to the backup BNG User Plane *BNG-UP1*.

> **NOTE**: At this steps, BNG User Plane *BNG-UP* 1 is expected to be back online. You can verify this, by checking the BNG User Plane *BNG-UP1's* node association state on the BNG CUPS Controller, using the `show health user-plane` command.

6. Clean up the backup BNG User Plane *BNG-UP2*. During this step the ports and subscribers on BNG User Plane *BNG-UP2* are cleaned up. After completing this step, BNG User Plane *BNG-UP2* will be placed back into its original state.

```
user@host# request user-plane maintenance disassociate serviced-user-plane BNG-UP1 backup-user-plane BNG-UP2
user@host# request user-plane maintenance complete serviced-user-plane BNG-UP1
```

Use Juniper BNG CUPS Smart Session Load Balancing

SUMMARY

This section describes how Juniper BNG CUPS uses smart session load balancing. This includes a description of the standards for broadband access network, a description of Juniper's BNG CUPS load balancing, and configuration requirements.

IN THIS SECTION

Juniper BNG CUPS Smart Session Load Balancing Overview

IN THIS SECTION

BNG CUPS smart session load balancing gives the operator the capability to distribute subscriber loads across the BNG User Planes in the network by moving subscribers from one BNG User Plane to another. Fast failover is a use case for subscriber session load balancing. The fast failover use case occurs when a BNG User Plane's access port goes down and subscribers are rebalanced over to another access port on the same BNG User Plane.

BNG CUPS smart session load balancing, operates in accordance with the Broadband Forums *TR 459 Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, Interface, and Protocol* specification. This smart session load balancing model takes into account the session load on a BNG User Plane and the throughput capacity used. It can be applied across different types of BNG User Planes, for any type of session access model (DHCP IPoE and PPPoE, single stack or double stack) and is controlled through the BNG CUPS controller. It assumes that there is Ethernet bridged access to the BNG User Planes, or an alternative like VPLS or EVPN. Smart session load balancing requires that the same residential gateway's first sign of life packet be received by multiple BNG User Planes. The first sign of life packets also, can be either DHCP Discover or PPPoE Active Discovery Initiation (PADI).

NOTE: A BNG User Plane's subscriber limit should be configured for each linecard's PIC and it should be set to the specific linecard PFE maximum limit. This is because the maximum limit varies for each linecard PFE type. The subscribers limit for a PFE is used by resource monitoring to enforce resource consumption and thresholds on the PFE at different calls per second (CPS) rates. For DHCP access models, only 95% of the subscriber limit is supported.

You should not use the **any** option in the `accept` stanza of the `auto-configure` configuration. Instead, you should use the specific client protocol type in the `accept` stanza (for example, DHCP, DHCPv6, PPPoE or a combination of the protocol types).

Figure 13 on page 58 shows how BNG CUPS Controller implements subscriber load balancing.

Figure 13: Juniper BNG CUPS Subscriber Load Balancing

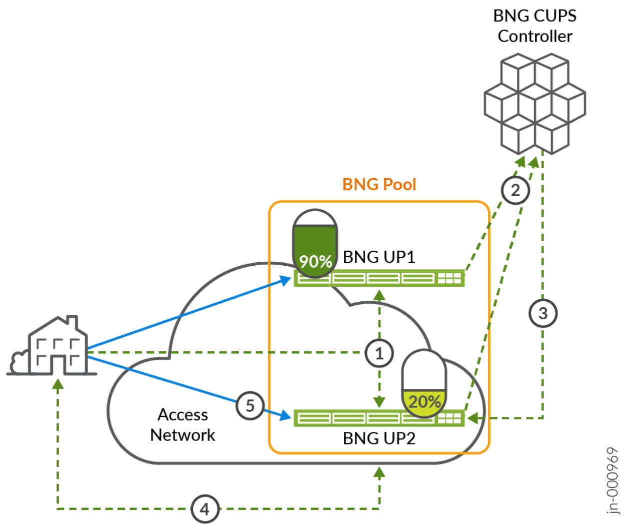

Following is the work flow that Juniper BNG CUPS uses for subscriber load balancing (see Figure 13 on page 58).

1. The subscriber session connects to the broadband access network. Both BNG User Planes (*BNG-UP1* and *BNG-UP2)* in the shared BNG pool receive the broadcasted first sign of life request and forward it to the BNG CUPS Controller.

2. The BNG CUPS Controller receives the first sign of life requests from both BNG User Planes. Because BNG User Plane *UP1* is currently loaded at 80%, the BNG CUPS Controller selects the less loaded BNG User Plane in the pool (*BNG-UP2*).

3. The BNG CUPS Controller replies to BNG User Plane *BNG-UP2*, letting it know that it is the anchor BNG User Plane for the subscriber.

4. BNG User Plane *BNG-UP2* forwards the reply that it received from the BNG CUPS Controller to the subscriber's residential gateway.

5. The subscriber's traffic now flows through BNG User Plane *BNG-UP2*.

The BNG CUPS session load balancing model is based on the following two mutually exclusive criteria:

- Load balancing at the BNG CUPS Controller is based on a live BNG User Plane reported load. The load is report as a percentage.

- Weight is configured in the `dynamic-profile` configuration on the BNG CUPS Controller. Weight can be either IFL-set weight or subscriber weight.

Report-based Subscriber Session Load Balancing

The BNG User Plane reported load balancing model assumes the following:

- It uses a logical-port Packet Forwarding Control Protocol (PFCP) Information element (IE) as described in the TR-459 technical report.
- It is dependent on the BNG User Plane sending the PFCP logical port usage reports to the BNG CUPS Controller.
- It is done in-line in the control packet I/O processing, by allowing or denying the first sign of life packet when comparing the BNG User Plane logical port candidates. It chooses the BNG User Plane with the lowest usage (lowest percentage utilization). The logical port utilization for the logical port candidates is stored in the load balancing database.

The following configuration example shows a BNG User Plane reported load balancing configuration on the BNG CUPS Controller.

```
[edit groups bng-director bng-controller]
load-balancing-groups {
    lb-report-group {
        report-based-mode {
            port up:boston:xe-5/0/5:1;
            port up:nashua-c:xe-0/1/2;
            port up:manchester:xe-1/3/1;
       }
    }
}
```

On each BNG User Plane that is part of a report-based load balancing group, the `subscribers-limit` configuration must be set for the line card or the forwarding engine that the load balancing port is on.

```
[edit configuration system services resource-monitor]
subscribers-limit {
    client-type any {
        fpc 0 {
            limit 8500;
        }
    }
}
```

Weight-based Subscriber Session Load Balancing

Weight can be defined in different ways, based on your needs: Weight can be subscriber bandwidth, logical interface set bandwidth, or an even number of subscribers per logical interface set.

Weight-based load balancing can work with hierarchical class of service (HCoS) or independently.

Weight-based load balancing does not use the BNG User Plane logical port reported load. You can still examine the reported load from the BNG User Plane logical port. Use the **show system subscriber load balancing group** commands to examine the reported load.

When you configure weight-based load balancing, the BNG User Plane reported load is used only for monitoring purposes and troubleshooting.

Weight in the BNG CUPS Controller dynamic profile has the following characteristics:

- It is dependent on the operator needs. It can be subscriber bandwidth, (subscriber or logical interface set) bandwidth, or the number of subscribers.
- It compares the configured logical port maximum weight to the computed weight.
- Computed weight is dynamic. It operates in the following ways:
  - It increases when each weighted item (subscriber or logical interface set) is instantiated.
  - It decreases when each weighted item (subscriber or logical interface set) is de-instantiated.
  - It compares the logical port configured maximum weight to allow or deny a subscriber on the logical port.
- It works with hierarchical class of service (HCoS) and it can work independently.
- It is part of the dynamic profile configuration. Weight based load balancing has a tolerance of one element above the maximum weight configured.
- When load balancing weight is configured the BNG User Plane logical port reported load is ignored.

The following configuration example shows a weight based load balancing configuration on the BNG CUPS Controller.

```
[edit groups bng-director bng-controller]
load-balancing-groups {
   lb-weight-group {
      weight-based-mode {
         port up:boston:xe-5/0/5:1 {
            max-weight 10;
         }
```

```
            port up:nashua:xe-0/1/2 {
               max-weight 20;
            }
            port up:manchester:xe-1/3/1 {
                max-weight 30;
            }
         }
      }
   }
```

Example: Configure Subscriber Session Load Balancing

Consider the use case in Figure 14 on page 61, where a BNG CUPS Controller manages two BNG User Planes (*UP-example-1* and *UP-example-3*). They both can receive the same residential gateway's PADI by being configured each with an active pseudowire that carry the same PADI to both BNG User Planes.

Figure 14: Load Balancing Combined with Subscriber Session Steering

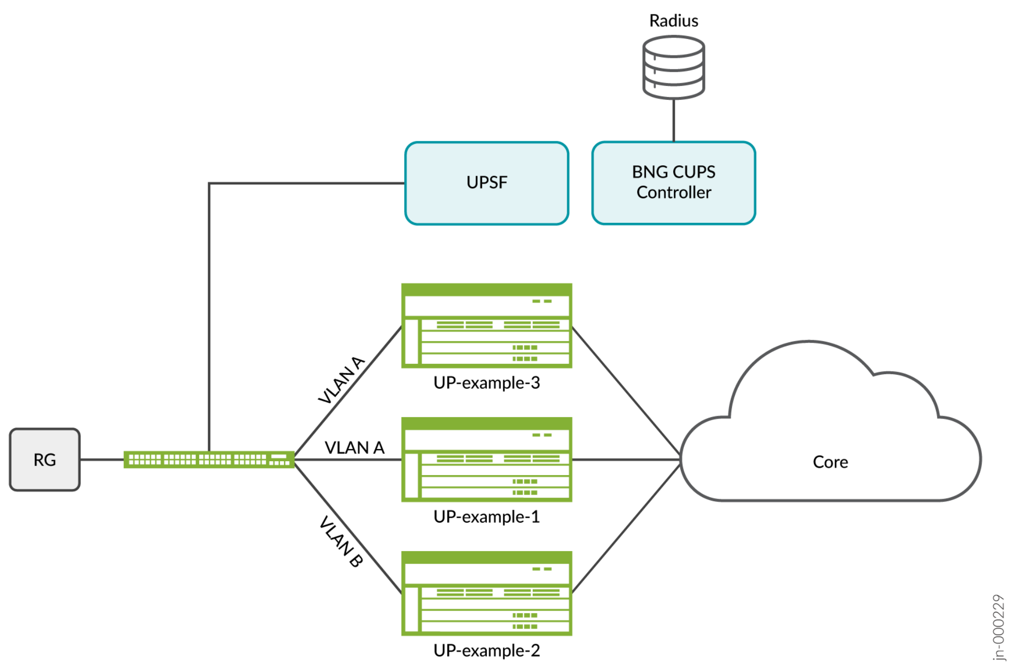

For weight-based load balancing, you use the dynamic profile configuration to specify subscriber weight or logical interface set weight.

To configure subscriber weight, perform the following configuration on the BNG CUPS Controller:

1. On the BNG CUPS Controller, define the load-balancing groups and the BNG User Plane with logical ports.

```
[edit groups bng-director bng-controller]
User@host# set load-balancing-groups group-name user-plane user-plane-name preferred logical-
port port-id
```

2. Configure the logical port maximum weight.

```
[edit groups bng-director bng-controller]
User@host# set load-balancing-groups group-name user-plane user-plane-name preferred logical-
port port-id max-weight max-weight-number
```

3. Configure the dynamic profile to specify either subscriber weight or logical interface set weight.

   - Configure subscriber weight.

   ```
   [edit]
   User@host# set dynamic-profiles dynamic-profiles-name interfaces $junos-interface-ifd-name
   unit $junos-interface-unit load-balance weight weight-number
   ```

   - Configure logical interface set weight.

   ```
   [edit]
   User@host# set dynamic-profiles dynamic-profiles-name interfaces interface-set $junos-phy-
   ifd-interface-set-name load-balance weight weight-number
   ```

Report-based Load Balancing Operational Behavior

Consider the example of a PPPoE subscriber login using BNG User Plane load reports for load balancing. In this example, the same PADI that the residential gateway sends arrives at both BNG User Plane *UP-example-1* and BNG User Plane *UP-example-3*.

Also, you define the load balancing group to contain *UP-example-1* and *UP-example-3* logical-ports to the pseudowires that carry the subscriber PADI.

```
[edit groups bng-director bng-controller]
user@host# set load-balancing-groups group-name user-plane UP-example-1 port UP-example-1:ps0.30
```

```
[edit groups bng-director bng-controller]
user@host# set load-balancing-groups group-name user-plane UP-example-3 port UP-example-3:ps0.25
```

Suppose that *UP-example-1* exceeds an incremental threshold for which an upper limit exists, resulting in *UP-example-1* reporting a load percentage that doesn't allow any more subscribers.

As was mentioned earlier, the same PADI that the residential gateway sends arrives at both *UP-example-1* and *UP-example-3*. Both *UP-example-1* and *UP-example-3* forward the PADI to the BNG CUPS Controller. The BNG CUPS Controller discards the *UP-example-1* PADI and allow the PPPoE subscriber to log in to *UP-example-3*.

On each BNG User Plane that is part of a report-based load balancing group, the `subscribers-limit` configuration must be set for the linecard or the forwarding engine that the load balancing port is on. This limit must be higher than the maximum number of subscribers expected.

For example, if the expected maximum is 8000, you set the `subscribers-limit` to 8500.

```
[edit configuration system services resource-monitor]
subscribers-limit {
    client-type any {
        fpc 0 {
            limit 8500;
        }
    }
}
```

You can use the following load balancing `show` command to examine the percentage load reported by the BNG User Planes for their logical ports.

```
user@host#> show load-balancing-group group lb-report
Logical-Port            % Usage     CPU Exceeded  Computed weight     Max weight
up:mx204-b:ae4            20          no              0                 0
up:mx204-i:xe-0/1/0       45          no              0                 0
up:mx204-b:ae4            10          no              0                 0
```

Weight-based Load Balancing Operational Behavior

Consider the example of a PPPoE subscriber login using BNG User Plane load reports for load balancing. In this example, the same PADI that the residential gateway sends arrives at both BNG User Plane *UP-example-1* and BNG User Plane *UP-example-3*.

Consider the example of a PPPoE subscriber login using weight for load balancing. In this example, the PADI that the residential gateway sends arrives at both BNG User Plane *UP-example-1* and BNG User Plane *UP-example-3*.

In this example, you configure the logical-port maximum weight on the BNG CUPS Controller. Define the load-balancing group to contain *UP-example-1* and *UP-example-3* logical-ports.

```
[edit groups bng-director bng-controller]
user@host# set load-balancing-groups group-name user-plane UP-example-1 port UP-example-1:ps0.30
max-weight 10
```

```
[edit groups bng-director bng-controller]
user@host# set load-balancing-groups group-name user-plane UP-example-3 port UP-example-3:ps0.25
max-weight 10
```

After you configure the weight, you then configure the logical interface set in the dynamic profile.

```
[edit]
user@host# set dynamic-profiles profile-name interfaces interface-set interface-set-name load-
balance weight 2.5
```

The first PPPoE subscriber that logs in creates the logical interface set on BNG User Plane *UP-example-1*. Each logical interface set weight is added up to a computed weight that must be less than 10 (the max logical port weight).

After the subscriber's log in creates the logical interface set and places the logical interface set on a BNG User Plane, it doesn't move. All subscribers belonging to that logical interface set follow the logical interface set (placed on the same BNG User Plane as their corresponding logical interface set).

After that, every new PADI coming in for this logical interface set is placed on BNG User Plane *UP-example-1* and dropped from BNG User Plane *UP-example-3*.

As subscribers for a new logical interface set login, the new logical interface set weight is added to the computed weight and compared to the maximum weight. When the computed weight is greater than the maximum weight, the new logical interface set is no longer placed on BNG User Plane *UP-example-1*. Instead, the logical interface set is placed on BNG User Plane *UP-example-3*.

Use Juniper BNG CUPS for Subscriber Steering

SUMMARY

This section describes how Juniper BNG CUPS uses subscriber steering. This includes a description of the standards for broadband access network, a description of Juniper's subscriber session steering and configuration requirements for subscriber session steering.

IN THIS SECTION

Standards Overview

In a traditional broadband access network, the access nodes connect customers to the network. Service gateways (such as the broadband network gateways) connect customers to network services. Today, the connectivity between the access node and the broadband network gateway (BNG) is generally very static. The subscribers on a particular access node usually connect to the same BNG (also referred to as the service gateway). Typically, subscribers make changes to configurations only when deploying or upgrading the network.

However, the requirements and the architecture of the broadband access network are changing. The world is becoming more dependent upon broadband, with home working placing more demands on the broadband network. Video streaming is no longer just about entertainment; it is an important part of how we learn and work.

Edge compute services and user needs require connectivity to service gateways that are closer to the user. This connectivity reduces the latency between the user and the service.

Service gateway nodes such as the BNG are evolving to become disaggregated. This separation of the control functions from the user plane (or data plane) functions allows for more scalability and flexibility.

With services moving further to the edge, scalability requirements change. Requiring more BNGs or smaller BNGs drives the need for disaggregation and scale-out.

You need to perform maintenance activities and upgrades more often to react to customer needs. Virtualization enables new network functions including service gateway creation, upgrades, and removal on demand.

Broadband Forum WT-474 Subscriber Session Steering (WT-474) requirements standardizes a more flexible and dynamic broadband access network to meet these new requirements.

From the WT-474 requirements, “WT474 is an architecture to enable dynamic real time decisions about the placement of subscribers in the network.”

Figure 15 on page 67 shows the WT-474 subscriber session steering architecture as defined by the WT-474 requirements.

The figure shows a disaggregated broadband network gateway as defined by the Broadband Forum's *TR-459 Control and User Plane Separation for a disaggregated BNG (TR-459)* technical report.

Figure 15: WT-474 Subscriber Session Steering Architecture

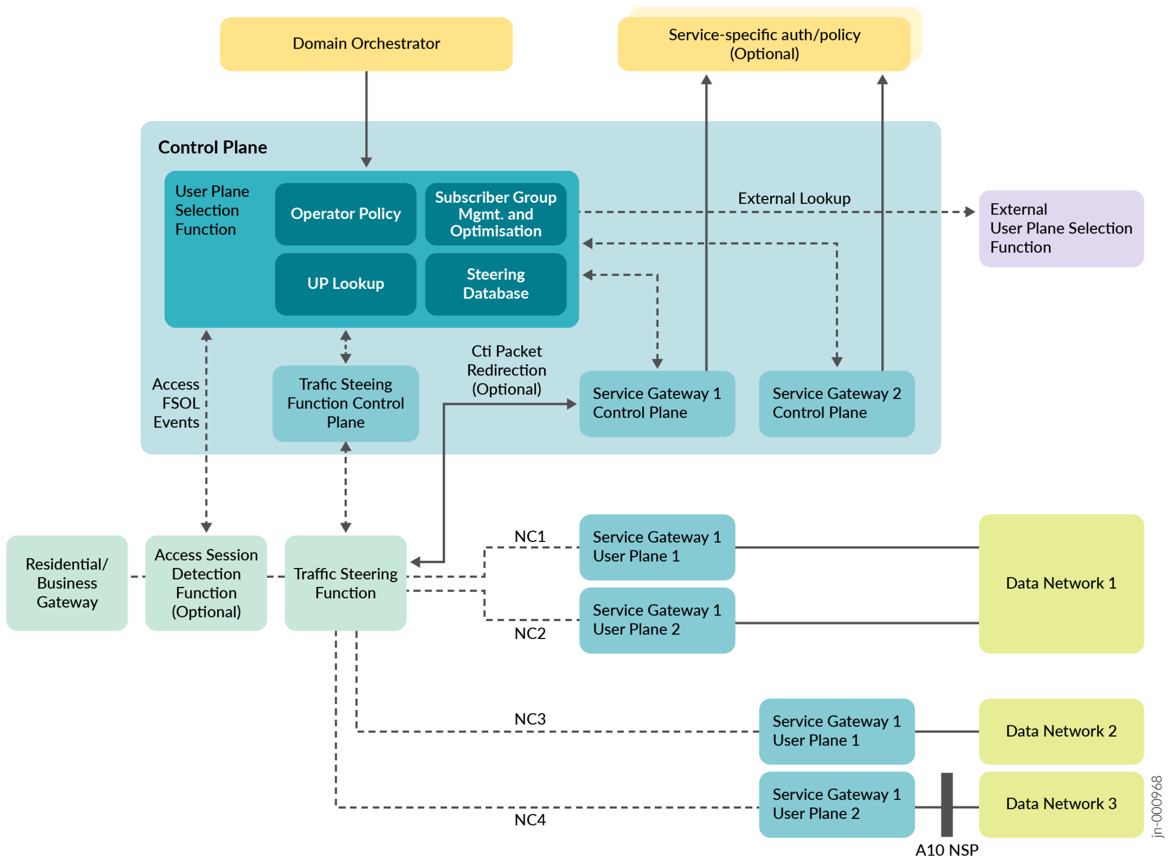

Following are new functions that are listed in the WT474 architecture:

- The access session detection function—Used to identify when a new subscriber is connecting to the network.
- The user plane selection function—Responsible for making the real-time decisions as to which service gateway and to which Juniper BNG User Planes (BNG User Planes) to connect the subscriber to.
- The traffic steering function control lane—Responsible for the configuration of the Traffic Steering Functions.
- The traffic steering function—Forwards the traffic of the subscribers to and from the identified BNG User Plane.

As described in the WT474 architecture, "There is no requirement in the architecture for these new functions to be implemented in dedicated boxes – for example, the Traffic Steering Function is expected simply to be an integral part of the existing Access Node, or aggregation switches, and the traffic steering function control plane and user plane selection function might be implemented as dedicated software, or as part of an SDN controller. The purpose of this architecture is to standardize the approach, interfaces and data models for session steering such that it can become a standard capability of an access network."

Benefits of Subscriber Steering and Load Balancing

Juniper BNG CUPS provides key operational and service-differentiating benefits.

Following are the operational benefits:

- Active load balancing of subscribers on BNG User Planes across the network
- Seamlessly moving subscribers away from BNG User Planes that require maintenance
- Enabling a Continuous Deployment approach to software upgrades
- Optimizing power consumption by moving subscribers onto a smaller number of BNG User Planes

Following are the service-differentiating benefits:

- Customer on-demand connecting to edge-service locations that can then deliver the required end user experience (for example, low latency)
- Mapping of specific service types to dedicated slices of the network
- Flexibility of trying new capabilities without requiring entire network upgrades

Juniper BNG CUPS Subscriber Session Steering Overview

IN THIS SECTION

As described in the WT-474 architecture, the user plane selection function together with the traffic steering function on the Juniper BNG CUPS Controller (BNG CUPS Controller) place subscriber sessions based on specific operator defined characteristics.

Subscriber session steering aggregates the user plane selection function and the traffic steering function control plane into the user plane selection function module. The user plane selection function module triggers the subscriber BNG User Plane placements based on the specific operator-defined characteristics.

Juniper BNG CUPS subscriber steering provides a one-touch mechanism for steering a subscriber's traffic through the access network to the selected BNG User Plane (service application point).

The steering works per subscriber and service using a RADIUS policy.

How Subscriber Session Steering Works

The user plane selection function module starts when a subscriber logs in. This module validates that the Juniper BNG User Plane supports the subscriber's services. If it cannot support the subscriber's services, the subscriber's login ends. The steering function then directs the subscriber to an appropriate BNG User Plane.

The user plane selection function selection uses the subscriber's service group vendor-specific attribute (VSA).

NOTE: Service group VSA is a new RADIUS VSA added to the subscriber for use with the user plane selection function.

The user plane selection function module chooses the BNG User Plane that hosts the subscriber based on the RADIUS service group VSA. (See Figure 16 on page 70.)

Figure 16: Subscriber Session Steering

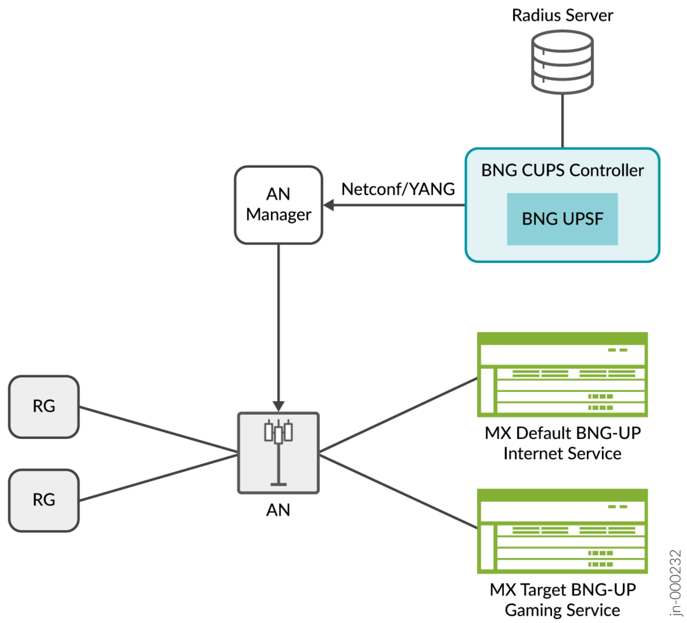

The default BNG User Plane is the ingress BNG User Plane for the subscriber login control packets. The target BNG User Plane is where user plane selection function places the subscriber. Depending on the network architecture, the default BNG User Plane and the target BNG User Plane might be the same physical BNG User Plane.

For example, the residential gateways can connect to the access network using known C-TAG and S-TAG VLANs. The user plane selection function module implements traffic steering through the access network to the proper BNG User Plane. It does this by mapping the residential gateway's VLAN tags to the correct access node's connected link (for example, pseudowire) that ends at the desired BNG User Plane.

Subscriber session steering assumes that the access node manager can communicate with the access node. Also, that it can change the mapping between the residential gateway's VLANs and the access node to the BNG User Plane connected link.

A cluster is a set of BNG User Planes that can service an access node. A subscriber that an access node services ends at the cluster. Each BNG User Plane sends the user plane capabilities to the Juniper BNG CUPS Controller. The capabilities include the name of the cluster to which the BNG User Plane to belongs and the name of the service group that the BNG User Plane supports.

The BNG CUPS Controller stores the BNG User Plane capabilities and sends an event to the user plane selection function module. Upon receiving this event the user plane selection function module writes this BNG User Plane data into the user plane selection function placement database.

The subscriber login sequence proceeds through the following steps:

1. You configure a BNG User Plane with a list of the service groups that it can support and the name of the cluster to which it belongs.

2. When connecting to the BNG CUPS Controller, the BNG User Plane provides a list of service groups and the cluster to the BNG CUPS Controller as capabilities.

3. The placement application takes the BNG User Plane service group capabilities and cluster from the BNG CUPS Controller. It then enters the BNG User Planes into its local database.

4. RADIUS creates a new service group VSA for the subscriber that contains the service group name.

5. The AAA Service Framework provides the subscriber’s service group name and the BNG User Plane identifier as part of the subscriber's login.

6. The user plane selection function module looks up whether the default BNG User Plane that the subscriber arrived on can support the service group.

   - A. Yes—The user plane selection function module sends an ACK login request to AAA.

   - B. No—The decision goes out to the user plane selection function module.

     a. The user plane selection function module looks for a BNG User Plane in the cluster that supports the required service group.

     b. The user plane selection function module tells the access node manager to connect the access node to the correct BNG User Plane to route the subscriber to.

     c. The user plane selection function module sends a NACK login request to AAA.

After the above sequence is completed, the following occurs: If a subscriber requires a service that is not supported on the default BNG User Plane, the subscriber reconnects and is placed on a BNG User Plane that does support the required service group.

Configuring Subscriber Session Steering

Consider the following use case: A BNG CUPS Controller manages two BNG User Planes (*UP-example-1*and *UP-example-2*). They both are part of the same cluster. The *UP-example-1* BNG User Plane can provide only Internet service. The *UP-example-2* BNG User Plane can provide premium services with low latency, such as gaming. Therefore, subscribers connecting to *UP-example-1* can get only Internet services, whereas subscribers connecting to *UP-example-2* can get gaming services.

You perform the subscriber session steering configuration on the BNG User Plane. RADIUS users must have the new service group VSA set to the desired service group.

On the BNG User Planes, define the clusters and service groups supported on the BNG User Planes. The service group names that you configured on the BNG User Planes must match the RADIUS service group VSA for the users.

To configure subscriber session steering, perform the following procedure on the BNG User Planes:

1. For *UP-example-1*, define the cluster named *example-cluster*.

   ```
   [edit system services subscriber-management]
   user@host# set mode user-plane selection-function cluster example-cluster
   ```

2. For *UP-example-2*, define the same cluster.

   ```
   [edit system services subscriber-management]
   user@host# set mode user-plane selection-function cluster example-cluster
   ```

3. Configure *UP-example-1* to support the *service-internet* service group.

   ```
   [edit system services subscriber-management]
   user@host# set mode user-plane selection-function service-group service-internet
   ```

4. Configure *UP-example-2* to support the *service-gaming* service group.

   ```
   [edit system services subscriber-management]
   user@host# set mode user-plane selection-function service-group service-gaming
   ```

Subscriber Session Steering Operational Behavior

Using the previous example, assume that a user subscribes to a gaming service. Also, use a subscriber VLAN-Tag of 100. The links between the access node and the BNG User Planes are pseudowires.

For example, the link from the access node to *UP-example-1* is ps0.25. The link from the access node to *UP-example-2* is ps0.35.

When a subscriber logs into *UP-example-1* over ps0.25, the BNG CUPS Controller receives the subscriber packet and notifies the user plane selection function module. The user plane selection function module looks up whether *UP-example-1* can support the *service-gaming* service group. Because *UP-example-1* can support only the *service-internet* service group, the user plane selection function module looks up which other BNG User Planes in the cluster can support the *service-gaming* service group.

The user plane selection function module finds the BNG User Plane *UP-example-2*, which supports the *service-gaming* service group. The user plane selection function module then tells the access node manager to cross-connect the subscriber's VLAN-Tag 100 to this link (pseudowire ps0.35). The access node manager communicates the steering information to the access node. So, during the subscriber's next login attempt, the subscriber is redirected to the correct BNG User Plane (*UP-example-2*).

Last, the user plane selection function module sends a NACK to the AAA Service Framework on the BNG CUPS Controller, which causes the subscriber to log in again. The second login attempt is redirected to the desired BNG User Plane.

Operational Behavior of Subscriber Session Steering and Load Balancing Combined

The most flexible and powerful use case is when you get all the benefits of network load balancing and service differentiation together in one topology.

In Figure 14 on page 61, you learn about this use case: If the BNG User Planes belong to the same cluster, you can steer subscribers based on different service requirements between VLAN A and VLAN B (for example, internet on BNG User Plane *UP-example-1* or gaming on BNG User Plane *UP-example-2*). You can also load balance subscribers on VLAN A between BNG User Plane *UP-example-1* and User Plane *UP-example-3*.

Use Juniper BNG CUPS Multicast

SUMMARY

This document presents conceptual information and sample configurations that you can use to help you set up multicast in your Juniper BNG CUPS environment.

IN THIS SECTION

Multicast Overview

Juniper BNG CUPS supports multicast functionality as it is defined in the The TR-459.3: Multi-Service Disaggregated BNG with CUPS: IPTV Multicast function — Reference Architecture, Deployment Models, Interface, and Protocol Specifications (TR-459) document.

Multicast provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many. Routing devices in IP multicast networks use multicast routing protocol to build a distribution tree that connects multicast receivers to sources.

The root of the distribution tree lies at the multicast source. Each subnetwork with hosts on the routing device that has at least one interested receiver is a leaf on the distribution tree.

When a new leaf subnetwork is added to the tree, a new branch is built, a leaf is added to the tree, and replicated packets are sent out on the outgoing interface. When a branch contains no leaves, because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is removed from the distribution tree, and no multicast packets are sent out to that interface.

Internet Group Management Protocol (IGMP) is a host to router signaling protocol used to support IPv4 multicasting. Similarly, Multicast Listener Discovery (MLD) Protocol manages membership of hosts and routers in IPv6 multicast groups.

Figure 17 on page 75 displays what BNG CUPS multicast in subscriber management scenario looks like.

Multicast goes through the following process:

- During subscriber login or change of authorization (CoA), the BNG CUPS Controller sends a multicast service activation or deactivation request to the BNG User Plane.
- Subscribers (multicast receivers) send either IGMP or MLD group membership reports to the BNG User Plane. This report indicates their interest in joining a multicast group. Similarly, when a subscriber wants to leave a group, it sends an IGMP or MLD leave group message to the BNG User Plane.
- Using the information contained in the IGMP or MLD join (leave) protocol messages, the BNG User Plane programs multicast on the subscriber flows in the forwarding plane and sources the multicast traffic from the core.

Figure 17: BNG CUPS Multicast Subscriber Management

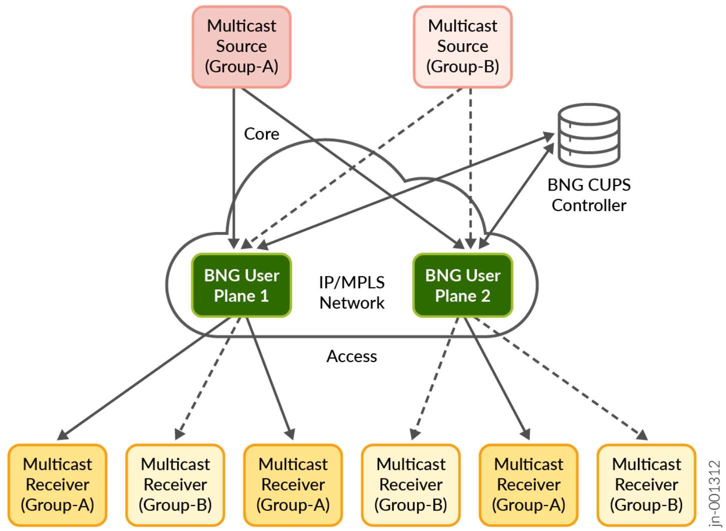

BNG CUPS Controller Multicast Configuration

You enable IGMP and MLD through dynamic profiles.

Configure IGMP for multicasting on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles profile-name protocols]
    IGMP {
        interface interface-name {
            (accounting | no-accounting);
            disable;
            distributed;
            group-limit limit;
            group-policy policy-name;
            group-threshold value;
            immediate-leave;
            log-interval seconds;
            oif-map map-name;
            passive;
            promiscuous-mode
            ssm-map-policy map-policy-name;
            static {
```

```
                group multicast-group-address {
                     exclude;
                     group-count number;
                     group-increment increment;
                     source ip-address {
                          source-count number;
                          source-increment increment;
                     }
                 }
            }
            version version;
      }
}
```

Configure MLD for multicasting on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles profile-name protocols]
    mld {
        interface interface-name {
            (accounting | no-accounting);
            disable;
            distributed;
            group-limit limit;
            group-policy policy-name;
            group-threshold value;
            immediate-leave;
            log-interval seconds;
            oif-map map-name;
            passive;
            ssm-map-policy map-policy-name;
            static {
                group multicast-group-address {
                     exclude;
                     group-count number;
                     group-increment increment;
                     source ip-address {
                          source-count number;
                          source-increment increment;
                     }
                 }
            }
            version version;
```

```
        }
}
```

BNG User Plane Multicast Configuration

IN THIS SECTION

Optional features are supported for multicast through policies. The policy is configured on the BNG User Plane and the name of the policy is specified in the BNG Controller's dynamic profile for a subscriber. Separate policies are required for IGMP and MLD.

You can use the group policy configuration to limit subscribers that can join a certain multicast group. The `igmp-group-policy` statement enables you to filter unwanted IGMP reports at the subscriber interface level.

Configure the IGMP and MLD group policy for multicasting on a BNG User Plane, as follows:

```
[edit policy-options]
policy-statement igmp-group-policy {
     term A1 {
           from {
                route-filter ip-address orlonger;
           }
           then accept;
     }
     then reject;
}
policy-statement mld-group-policy {
     term A1 {
           from {
                route-filter ip-address orlonger;
           }
           then accept;
```

```
      }
      then reject;
}
```

In outgoing interface (OIF) mapping, the subscriber interface is used for unicast data and for IGMP and MLD to join and leave requests. A separate multicast VLAN is used only for multicast streams and is shared across subscribers. The policy is configured in the BNG User Plane and the name of the policy is specified in the BNG Controller's dynamic profile.

Configure the OIF mapping policy for multicasting on the BNG User Plane, as follows:

```
[edit policy-options]
    policy-statement igmp-oif-map-policy {
     term A {
           from {
                route-filter ip-address orlonger;
                route-filter ip-address exact;
           }
           then {
                map-to-interfce interface-name;
                accept;
           }
     }
     then reject;
}
policy-statement mld-oif-map-policy {
     term A {
           from {
                route-filter ip-address orlonger;
                route-filter ip-address exact;
           }
           then {
                map-to-interfce interface-name;
                accept;
           }
      }
      then reject;
}
```

Source-specific multicast (SSM) mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

Configure the SSM mapping policy for multicasting on the BNG User Plane, as follows:

```
[edit policy-options]
    policy-statement igmp-ssm-map-policy {
     term A1 {
           from {
                route-filter ip-address orlonger;
           }
           then {
                ssm-source ip-address;
                accept;
           }
     }
}
policy-statement mld-ssm-map-policy {
     term A1 {
           from {
                route-filter ip-address orlonger;
           }
           then {
                ssm-source ip-address;
                accept;
           }
      }
}
```

BNG CUPS Controller Operational Commands for Multicast

You can view the subscribers that are enabled for IGMP and MLD by using the following commands on the BNG CUPS Controller:

- `show user-plane igmp interface up-name` *`user-plane-name`*
- `show user-plane mld interface up-name` *`user-plane-name`*

BNG User Plane Operational Commands for Multicast

You can view details about current membership and membership statistics by using the following commands on BNG User Planes:

- *show igmp interface*

- *show igmp group*
- *show igmp statistics*
- *show igmp statistics continuous*
- *show mld interface*
- *show mld group*
- *show mld statistics*
- *show mld statistics continuous*

Use Juniper BNG CUPS Lawful Intercept

SUMMARY

This document presents conceptual information and sample configurations that you can use to help you set up lawful intercept in your Juniper BNG CUPS environment.

IN THIS SECTION

Lawful Intercept Overview

Lawful intercept is a process for obtaining communications network data related to a target individual or organization, as authorized by a judicial or administrative order. Subscriber secure policy (lawful intercept) provides capability to mirror traffic on a per-subscriber basis. Subscriber traffic can be mirrored as well as the events related to the subscriber session, that is being mirrored, can be monitored. (For more information about lawful intercept, see *Lawful Intercept on Junos Multi-Access User Plane*.)

Subscriber secure policy mirroring can be based on information provided by either RADIUS or Dynamic Tasking Control Protocol (DTCP) and can mirror both IPv4 and IPv6 traffic. After subscriber secure policy is triggered, the subscriber's incoming and outgoing traffic are both mirrored. The original traffic is sent to its intended destination, and the mirrored traffic is sent to a mediation device for analysis. The actual mirroring operation is transparent to subscribers whose traffic is being mirrored. A special UDP/IP plus lawful intercept header is prepended to each mirrored packet sent to the mediation device. The mediation device uses the LI header (8 Bytes) to differentiate different users

BNG CUPS Subscriber Secure Policy functionality in Juniper BNG CUPS is split between the BNG CUPS Controller and the BNG User Plane in accordance with the "Support of Lawful Intercept" section of the *3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Interface between the Control Plane and the User Plane Nodes; Stage 3* document.

Figure 18 on page 81 shows RADIUS server based subscriber lawful Interception:

1. The mediation device sends intercept provisioning information to the RADIUS server.
2. The RADIUS server sends lawful intercept VSAs to the BNG CUPS Controller, either in an ACCESS-ACCEPT or change of authorization (CoA) message.
3. BNG CUPS Controller sends lawful intercept traps to the mediation device (if there are any) and programs the BNG User Plane.
4. Subscriber traffic is mirrored to the mediation device from the BNG User Plane.

Figure 18: RADIUS Based Subscriber Lawful Intercept

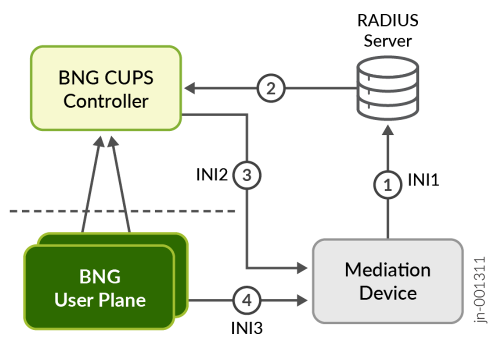

Figure 2 shows DTCP Based Subscriber Lawful Interception:

1. The mediation device sends a DTCP message to the BNG CUPS Controller.
2. The BNG CUPS Controller sends lawful intercept traps to the mediation device (if there are any) and programs the BNG User Plane.
3. Subscriber traffic is mirrored to the mediation device from the BNG User Plane.

Figure 19: DTCP Based Subscriber Lawful Intercept

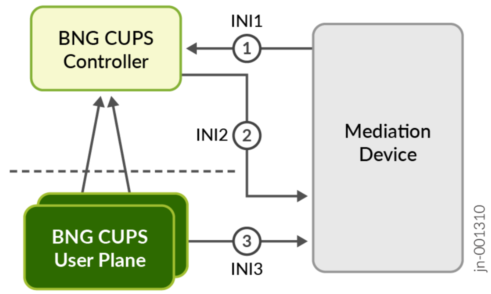

Configure Lawful Intercept

IN THIS SECTION

The `radius-flow-tap` configuration commands are split between the BNG CUPS Controller and the BNG User Planes.

BNG CUPS Controller Lawful Intercept Configuration

Configure the `radius-flow-tap` service for subscriber secure policy mirroring on the BNG CUPS Controller, as follows:

```
[edit services]
radius-flow-tap {
   policy policy-name {
…
   }
}
```

Configure Dynamic Tasking Control Protocol

You run all the Dynamic Tasking Control Protocol configurations on the BNG CUPS Controller.

Perform the following configuration on the BNG CUPS Controller:

```
[edit]
System {
       login {
         class class-name {
             permissions flow-tap-operation;
         }
         user user-name {
             uid uid;
             class class-name;
             authentication {
                 encrypted-password <string>
             }
         }
     }
     services {
         flow-tap-dtcp {
             ssh {
                 connection-limit connection-limit;
                 rate-limit rate-limit;
             }
         }
     }
}
```

BNG User Plane Lawful Intercept Configuration

Configure the `radius-flow-tap` service for subscriber secure policy mirroring parameters for the mediation device on the BNG User Plane, as follows:

```
[edit services]
radius-flow-tap {
  forwarding-class class-name;
  logical-system logical-system-name routing-instance routing-instance-name;
  routing-instance routing-instance-name;
```

```
  source-ipv4-address ipv4-address;
}
```

Use Juniper BNG CUPS Captive Portal and Content Delivery Services

SUMMARY

This document presents conceptual information and sample configurations that you can use to help you set up Captive Portal and Content Delivery services in your Juniper BNG CUPS environment.

IN THIS SECTION

Captive Portal and Content Delivery Overview

Captive Portal and Content Delivery (CPCD) also known as HTTP redirect services is supported in Juniper BNG CUPS. When CPCD is enabled, HTTP request traffic from subscribers is intercepted and redirected to an external captive portal device. The captive portal is often the initial page a subscriber sees after logging into the network. The captive portal typically provides authentication and authorization services for redirected subscribers before granting access to protected servers outside of a walled garden. A walled garden, also known as an allow list, defines a group of servers where access is provided to subscribers without reauthorization through a captive portal. These walled gardens enable you to increase revenue by marketing various services to your customers.

Figure 20 on page 85 shows the network elements that are involved in the BNG CUPS CPCD function.

Figure 20: BNG CUPS Captive Portal and Content Delivery

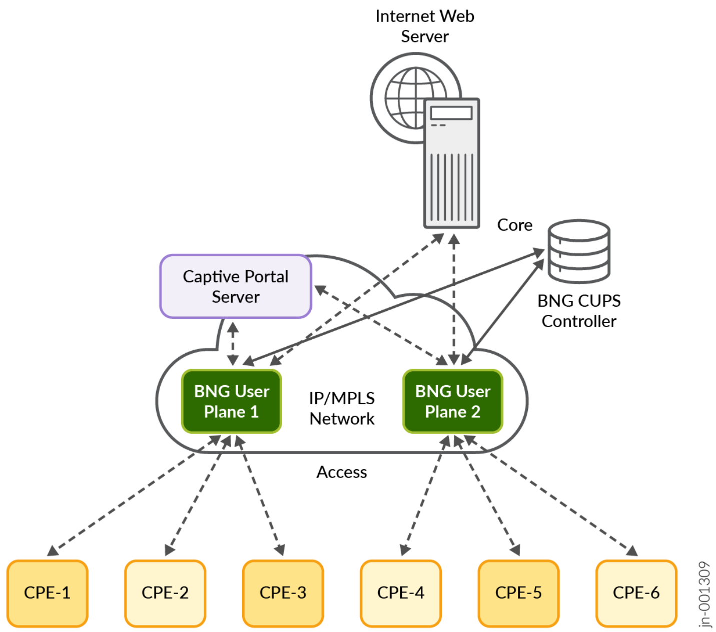

The BNG CUPS CPCD functionality is disaggregated to run on both the BNG CUPS Controller and the BNG User Planes. This functionality is described in the "Traffic Redirection" section of the *3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Interface between the Control Plane and the User Plane Nodes; Stage 3* specification.

Enabling subscribers for CPCD and provisioning service filter is handled in the BNG CUPS Controller while processing of packets for redirection is done in the BNG User Planes. The Routing Engine of the BNG User Plane performs the redirect function and additional service card is not required in the BNG User Plane.

There are two modes which are supported, redirect and rewrite. For redirect, the portal server URL is returned for the HTTP requests. For rewrite, the IP address destination is modified and traffic is sent to the portal server.

Configure CPCD

IN THIS SECTION

BNG CUPS Controller CPCD Configuration

Configure the service filter on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| firewall]
family inet {
    service-filter http-reverse-v4 {
        term http {
            from {
                source-address {
                    192.100.1.100/32;
                }
            }
            then {
                count httpv4_reverse_count;
                service;
            }
        }
        term skip {
            then {
                count httpv4_reverse_skip_count;
                skip;
            }
        }
    }
    service-filter redirect-in {
        term http {
            from {
                destination-port [ 80 8080 https ];
            }
            then {
                count httpv4_redirect_count;
                service;
            }
```

```
        }
        term skip {
            then {
                count httpv4_redirect_skip_count;
                skip;
            }
        }
    }
}
family inet6 {
    service-filter http-reverse-v6 {
        term http6 {
            from {
                source-address {
                    1000:1:2::1/128;
                }
            }
            then {
                count httpv6_reverse_count;
                service;
            }
        }
        term skip6 {
            then {
                count httpv6_reverse_skip_count;
                skip;
            }
        }
    }
    service-filter redirectv6-in {
        term http6 {
            from {
                destination-port [ 80 8080 https ];
            }
            then {
                count httpv6_redirect_count;
                service;
            }
        }
        term skip6 {
            then {
                count httpv6_redirect_skip_count;
                skip;
```

```
            }
        }
    }
}
```

Configure the service set on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| services]
captive-portal-content-delivery {
    profile cpcd_dyn_cpcd {
        dynamic;
    }
}
service-set sset-redirect-dyn {
    service-set-options {
        routing-engine-services;
    }
    captive-portal-content-delivery-profile cpcd_dyn_cpcd;
}
```

Configure the HTTP redirect profile on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles http-redirect-re-dyn]
variables {
    redirect-url default-value https://www.juniper.net;
}
interfaces {
    pp0 {
        unit "$junos-interface-unit" {
            family inet {
                service {
                    input {
                        service-set sset-redirect-dyn service-filter redirectin;
                    }
                    output {
                        service-set sset-redirect-dyn service-filter http-reverse-v4;
                    }
                }
            }
            family inet6 {
                service {
```

```
                    input {
                        service-set sset-redirect-dyn service-filter redirectv6-in;
                    }
                    output {
                        service-set sset-redirect-dyn service-filter httpreverse-v6;
                    }
                }
            }
        }
    }
}
services {
    captive-portal-content-delivery {
        rule r1 {
            match-direction input;
            term 1 {
                then {
                    redirect "$redirect-url";
                }
            }
        }
    }
}
```

To enable the subscriber to use the HTTP redirect service, the following must be provisioned on the AAA server:

- cpcduser Cleartext-Password—"cpcd123"
- Auth-Type—Local,
- Service-Type—Framed-User
- Framed-Pool—"v4pool"
- Framed-IPv6-Pool—"v6-na-pool-0"
- ERX-Service-Activate:1 +—'http-redirect-re-dyn(https://www.juniper.net)'

Configure the HTTP rewrite profile on the BNG CUPS Controller, as follows:

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles http-rewrite-re-dyn]
variables {
    rewrite-ip default-value 7.1.1.2;
```

```
}
interfaces {
    pp0 {
        unit "$junos-interface-unit" {
            family inet {
                service {
                    input {
                        service-set sset-redirect-dyn service-filter redirectin;
                    }
                    output {
                        service-set sset-redirect-dyn;
                    }
                }
            }
            family inet6 {
                service {
                    input {
                        service-set sset-redirect-dyn service-filter redirectv6-in;
                    }
                    output {
                        service-set sset-redirect-dyn;
                    }
                }
            }
        }
    }
}
services {
    captive-portal-content-delivery {
        rule r1 {
            match-direction input;
            term t1 {
                then {
                    rewrite destination-address $rewrite-ip;
                }
            }
        }
```

```
    }
}
```

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles http-rewrite-v6-re-dyn]
variables {
    rewrite-ip default-value 1000:1:2::2;
}
interfaces {
    pp0 {
        unit "$junos-interface-unit" {
            family inet {
                service {
                    input {
                        service-set sset-redirect-dyn service-filter redirectin;
                    }
                    output {
                        service-set sset-redirect-dyn;
                    }
                }
            }
            family inet6 {
                service {
                    input {
                        service-set sset-redirect-dyn service-filter redirectv6-in;
                    }
                    output {
                        service-set sset-redirect-dyn;
                    }
                }
            }
        }
    }
}
services {
    captive-portal-content-delivery {
        rule r1 {
            match-direction input;
            term t1 {
                then {
                    rewrite destination-address $rewrite-ip;
                }
```

```
            }
        }
    }
}
```

To enable the subscriber to use the rewrite service, the following must be provisioned on the AAA server.

- cpcdrewrite Cleartext-Password—"joshua"
- Auth-Type—Local
- Service-Type—Framed-User
- Framed-Pool—"v4pool"
- Framed-IPv6-Pool—"v6-na-pool-0"
- ERX-Service-Activate:1 +—'http-rewrite-dyn(7.1.1.2)'

BNG User Plane CPCD Configuration

Configure the service set on the BNG User Plane, as follows:

> **NOTE**: The configured `service-set` name must match the BNG CUPS Controller provisioned `service-set` name.

```
[edit configuration services]
captive-portal-content-delivery {
    profile cpcd_dyn_cpcd {
        dynamic;
    }
}
service-set sset-redirect-dyn {
    service-set-options {
        routing-engine-services;
    }
    captive-portal-content-delivery-profile cpcd_dyn_cpcd;
    interface-service {
        service-interface si-0/0/0;
```

```
    }
}
```

```
[edit configuration interfaces]
si-0/0/0 {
    unit 0 {
        family inet;
        family inet6;
    }
}
```

BNG CUPS Controller Operational Commands

To view the subscribers that are enabled for CPCD, you can use the "`show subscribers user-plane up-name` *`bng-user-plane-name`*" on page 254 command.

BNG User Plane Operational Commands

To view the subscribers that are enabled for CPCD, you can use the *show services captive-portal-content-delivery* command.

Configure Agent Circuit Identifier and IFL Sets in Juniper BNG CUPS

SUMMARY

This document presents sample configurations that you can use to help you set up Agent Circuit Identifier and IFL Sets in your Juniper BNG CUPS environment.

IN THIS SECTION

Configure Agent Circuit Identifier and IFL Sets

NOTE: ACI virtual VLAN sets (ACI sets) are IFL sets of client IFLs, they are not VLAN IFLs.

All service related configurations of ACI sets must be defined in the client's IFL dynamic profile. Service configurations in the ACI dynamic profile are not supported.

For examples of service configurations, see the following:

- "Example: Configure Class of Service" on page 94
- "Example: Configure Dynamic Service VLAN Profile" on page 95
- "Example: Configure Dynamic ACI Virtual VLAN Profile" on page 96
- "Example: Targeted Distribution Configuration for IFL for DHCP" on page 97
- "Example: Targeted Distribution Configuration for IFL for PPPoE" on page 98
- "Example: Targeted Distribution Configuration for IFL Sets for DHCP" on page 98
- "Example: Targeted Distribution Configuration for IFL Sets for PPPoE" on page 98

Example: Configure Class of Service

```
{edit]
class-of-service {
    forwarding-classes {
        queue 0 BestEffort;
        queue 1 Gold;
        queue 2 RealTime;
    }
    scheduler-maps {
        map-name {
            forwarding-class class-name scheduler scheduler-name;
        }
    }
    schedulers {
        scheduler-name {
            transmit-rate percent 25;
            shaping-rate percent 80;
        }
```

```
    }
}
```

Example: Configure Dynamic Service VLAN Profile

```
[edit}
dynamic-profiles {
    profile-name {
        interfaces {
            demux0 {
                unit interface-unit {
                    demux-source [ inet inet6 ];
                    vlan-tags outer stacked-vlan-id inner vlan-id;
                    auto-configure {
                        agent-circuit-identifier {
                            dynamic-profile example-aci-profile-with-no-services;
                        }
                    }
                    demux-options {
                        underlying-interface underlying-interface-name;
                    }
                    family inet {
                        unnumbered-address lo0.0;
                    }
                    family inet6 {
                        unnumbered-address lo0.0;
                    }
                    family pppoe {
                        duplicate-protection;
                        dynamic-profile ppp-dp-pp0;
                    }
                }
            }
        }
        protocols {
            router-advertisement {
                interface interface-name;
            }
        }
```

```
    }
}
```

Example: Configure Dynamic ACI Virtual VLAN Profile

```
[edit}
dynamic-profiles {
    example-aci-profile-with-no-services {
        interfaces {
            interface-set interface-set-name {
                interface interface-ifd-name;
            }
        }
    }
}
Dynamic Client IFL profile:
dynamic-profiles {
    example-client-dhcp-demux-profile {
        interfaces {
            interface-set interface-set-name {
                interface demux0 {
                    unit interface-unit;
                }
            }
            demux0 {
                unit interface-unit {
                    demux-options {
                        underlying-interface underlying-interface-name;
                    }
                    family inet {
                        demux-source {
                            subscriber-ip-address;
                        }
                        unnumbered-address lo0.0;
                    }
                    family inet6 {
                        demux-source {
                            subscriber-ipv6-address;
                        }
                        unnumbered-address lo0.0;
                    }
```

```
                }
            }
        }
        protocols {
            router-advertisement {
                interface jnterface-name;
            }
        }
        class-of-service {
            traffic-control-profiles {
                tcpSet {
                    scheduler-map schedular-map-name;
                    shaping-rate 1g;
                    guaranteed-rate 100m;
                }
            }
            interfaces {
                interface-set interface-set-name {
                    output-traffic-control-profile profile-name;
                }
            }
        }
    }
}
```

Configure Aggregated Ethernet Targeted Distribution

Targeted distribution configuration for IFL and IFL set is supported only when configured in a client dynamic profile. Targeted distribution is not supported for IFL and IFL set in a VLAN dynamic profile.

Example: Targeted Distribution Configuration for IFL for DHCP

```
set dynamic-profiles client-dhcp-demux interfaces demux0 unit interface-unit targeted-
distribution
```

Example: Targeted Distribution Configuration for IFL for PPPoE

```
set dynamic-profiles ppp-dp-pp0 interfaces pp0 unit interface-unit targeted-distribution
```

Example: Targeted Distribution Configuration for IFL Sets for DHCP

```
set dynamic-profiles client-dhcp-demux interfaces interface-set svlan-interface-set-name
targeted-distribution
set dynamic-profiles client-dhcp-demux interfaces demux0 unit interface-unit targeted-
distribution
```

Example: Targeted Distribution Configuration for IFL Sets for PPPoE

```
set dynamic-profiles ppp-dp-pp0 interfaces interface-set svlan-interface-set-name targeted-
distribution
set dynamic-profiles ppp-dp-pp0 interfaces pp0 unit interface-unit targeted-distribution
```

Understanding DVLAN (Single/Dual tag) for Subscriber Services Scaling (Junos Evolved for ACX7100-48L, ACX7332, and ACX7348 Devices)

IN THIS SECTION

DDOS Policers Configuration for Broadband Edge

Starting Junos Evolved release 24.4R1, the default DDOS policer support feature provides configuration options for BBE DDOS protocol groups namely – DHCPv4, DHCPv6, PPPoE, PPP, L2TP. These configuration values are supported on ACX7100-48L with the desired scale and performance required to bring up BBE subscribers.

This configuration of DVLAN (single and dual tag) for subscriber services performance scaling also supports:

- DHCP (IP-DEMUX and IP-DEMUX lite) & PPPoE subscribers (IPv4, IPV6, and dual stack) with CoS lawful intercept and filter support.
- DVLAN (Single and dual tag) with L2TP LAC DDOS policers configuration for BBE protocols. Individual BBE protocol and DDOS protocol group configuration is enabled. Only the aggregate DDOS protocol group `dhcpv4v6` is active.
- Subscriber scale qualification with Class of Support (CoS) for L2TP access concentrator (LAC) subscriber interfaces for IPV4, IPV6 and dual stack.

The combined DDOS policer protocol group `dhcpv4v6` is supported on ACX series devices. This policer group supports the aggregated traffic of BBE protocols (DHCPv4, DHCPv6, PPPoE, PPP and L2TP).

Following are the different BBE-related DDOS protocol groups CLI configurations supported on ACX7100-48L, ACX7332, and ACX7348 devices using the `set system ddos-protection protocols` command:

- `user@root# set system ddos-protection protocols dhcpv6 ?`
- ```
  user@root# set system ddos-protection protocols l2tp ?
  ```
- For individual policers you can configure the DDOS parameters like bandwidth, burst as follows:

  ```
  user@root# set system ddos-protection protocols dhcpv4 aggregate
  ```

Support for BNG PPP and L2TP-LAC on Non-default Routing Interfaces

BBE is disabled by default. The BBE protocols can be enabled using the following configurations:

- User Plane BBE Mode
  ```

User Plane BBE Mode

The subscriber-management mode can be configured as user-plane with following configuration.

```
system
{
  services
  {
    subscriber-management
    {
      enable;
      mode
      {
        user-plane
        {
        user-plane-name bng-user-plane-name;
        transport
        {
        inet <user-plane-ip>;
        }
        control-plane {
        bng-controller-name bng-cups-controller-name;
        transport {
          inet <control-plane-ip>;
        }
        }
      }
   }
}
}
```

An additional PPPoE underlying configuration option is available for the BBE support in the control-plane mode as follows:

```
[edit system services subscriber-management] interfaces <interface name>
        {
        unit <unit #>

        {
             pppoe-underlying-options
               {
```

```
                    ….
                }


        }
          auto-configure
        {
         …
        }
    }
```

In the above DDOS protocol group, there is one aggregate policer and one or more per-packet type policers. Aggregate policer runs for all packet traffic belonging to that protocol, while per-packet type policers run for traffic belonging to an individual packet type of that protocol, for example, PPPoE or L2TP.

Configurations related to different subscriber-management features are enabled. All BBE protocol feature specific configurations are enabled, starting Junos Evolved release 23.4R1.

Subscriber Stacking over PWHT Interface (ACX7100-48L, ACX7332, ACX7348)

Starting in Junos OS Evolved release 24.4R1, subscriber access models using DHCP Server and Relay dual-stack and PPPoE dual-stack configurations are supported over Layer-2 pseudowire headend termination (PWHT) interfaces. This feature includes support for:

- Firewall Filters
- Hierarchical CoS
- Lawful Intercept
- Multicast.

Use the `show user-plane subscribers` command in user-plane BBE mode to view the subscriber stacking configuration details. [See Understanding DVLAN (Single/Dual tag) for Subscriber Services Scaling (Junos Evolved for ACX Series Devices), and "show user-plane subscribers" on page 514.]

How to Use the Juniper BNG CUPS Controller Utility Commands

SUMMARY

After you have installed Juniper BNG CUPS Controller (BNG CUPS Controller), you can perform numerous administrative functions.

IN THIS SECTION

Access Juniper BNG CUPS Controller Utility Commands

You can use the BNG CUPS Controller utility script (`dbng`) to administer the application and to access the CLI that you use for configuring operations. The BNG CUPS Controller installation places the utility script in **/usr/local/bin**.

The `dbng` utility script performs the tasks you need to do to manage BNG CUPS but masks the complexity of the `kubectl` command. This masking of the `kubectl` commands simplifies your administrative duties.

The `dbng` utility script uses the Kubernetes `kubectl` utility commands to do the following:

- Create and delete objects.
- Conduct interactive sessions with pod containers.
- Display the status of the BNG CUPS Controller objects.

Table 1 on page 103 lists the commands that you can invoke with the `dbng` utility script and describes the action that each command initiates.

Table 1: BNG CUPS Controller Utility Script Commands

Command Name	Action
`sudo -E dbng clean [--docker] [--release` *`software-release`*`] [--dry-run] [--uninstall]`	Clean up unneeded releases and Docker cache. To run this command, you need sudo root privileges. This command offers the following options: • `log [error, warning, info, debug]`—Adjusts the log level. • `no-color`—Prints messages without colors. • `docker`—Only cleans the local Docker cache, all other files remain. • `release` *`software-release`*—Specify a release to clean or clean all possible releases. • `dry-run`—Identifies releases and docker images for removal and prints them to console. This command does not actually clean any releases or the Docker cache. • `uninstall`—Uninstalls all BNG CUPS Controller materials from the disk. The command does not affect the running application. • `clustr-repos`—Clean the cluster repository of the clusters that have been removed.
`dbng cli --context` *`context-name`* `[-p|--pipe]`	Gives you access to the CLI that you can use to configure BNG CUPS Controller features. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. • `pipe`—Allows you to pipe input into the command.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`dbng contexts [-o|--output json]`	Displays the available contexts for control with BNG CUPS Controller. This command offers the following options: • `contexts`—Lists the available contexts. • `output json`—Allows you to request the output in JSON format.
`sudo -E dbng cpi add --context` *`context-name`* `-h --log --version` *`software-release`* `--ip-aaa` *`ip-address`* `cpi-`*`label`*	Deploys a new control plane instance (CPi) pod. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `h, help`—Shows the help message and exit. • `l, log [error, warning, info, debug]`—Adjusts the log level. • `version` *`software-release`*—The software release for the new CPi pod. Enter a release. • `name`—Name of the new CPi. • `--ip-aaa` *`ip-address`*—You use the `ip-aaa` option to specify a single external IP address to use for the RADIUS listener port on the CPi. The address remains the same across multiple geography switchovers. The `ip-aaa` option requires a L3-enabled MetalLB instance. MetalLB BGP peering is used to direct traffic to the CPi on the active workload cluster. • `cpi-`*`label`*—Specify a label that is used to identify the CPi .

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`sudo -E dbng cpi rm --context` *`context-name`* `cpi-`*`label`*	Removes a control plane instance (CPi) pod. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `cpi-`*`label`*—Specify the CPi's label.
`dbng ip --context` *`context-name`* `[-o| --output json] [--detail]`	Displays the IP addresses of every service with an external IP address. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `output json`—Allows you to request the output in JSON format. • `detail`—Displays detailed IP information.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
Single geography setup use: `sudo -E dbng link --version` *`software-release`* `--context` *`context-name`* `--from-running` Multiple geography setup use: `sudo -E dbng link --context` *`karmada-context-name`* `--from-running --workload-contexts` *`workload-1-context-name`* *`workload-2-context-name`* `--observer-context` *`observer-context-name`*`--version` *`software-release`*	Links a cluster to a specific software version. To run this command, you need sudo root privileges. This command offers the following options: • `version` *`software-release`*—Specify the software release to link to the cluster specific repository. • `context` *`context-name`*—The Kubernetes context name to link to the software release. Enter the name of the context. • `context` *`karmada-context-name`*—The context name of the Karmada context that is created when Karmada is installed on the management cluster. Used with a multiple geography setup only. Enter the name of the context. • `from-running`—Attempts to match the software releases of the running BNG CUPS Controller. The `from-running` option is used to recover the version and user settings from a running application deployment. This can help you recover a failed jump host or for synchronizing a network of jump hosts. • `workload-contexts` *`workload-1-context-name workload-2-context-name`*—The two workload context names. Used with a multiple geography setup only. Enter the name of the context. • `observer-context` *`management-context-name`*—The context name for the management cluster. Used with a multiple geography setup only. Enter the name of the context.
`sudo -E dbng multi-cluster add-cluster` *`workload-context-name`*`--context` *`karmada-context-name`*	Adds a workload context to an existing management context. To run this command, you need sudo root privileges. This command offers the following options: • *`workload-context-name`*—The context name for the workload cluster that is being added. • `context` *`karmada-context-name`*—The context name of the Karmada context that is associated with the management cluster. Enter the name of the context. **NOTE**: This command is only supported in a multiple geography setup.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`sudo -E dbng multi-cluster rm-cluster` *`workload-context-name`* `--context` *`karmada-context-name`*	Removes a workload context from an existing management context. To run this command, you need sudo root privileges. This command offers the following options: • *`workload-context-name`*—The context name for the workload cluster that is being removed. • `context` *`karmada-context-name`*—The context name of the Karmada context that is associated with the management cluster. Enter the name of the context. `context` *`karmada-context-name`*—The context name of the Karmada context that is associated with the management cluster. Enter the name of the context. **NOTE**: This command is only supported in a multiple geography located setup.
`dbng multi-cluster status --context` *`karmada-context-name`*	Displays the status of the workload clusters from the application's perspective. This command offers the following option: `context` *`karmada-context-name`*—The context name of the Karmada context that is associated with the management cluster. Enter the name of the context. **NOTE**: This command is only supported in a multiple geography setup.
`sudo -E dbng multi-cluster switchover --context` *`karmada-context-name`*	Initiates a switchover to the other workload cluster. The CPi moves to the other workload cluster. To run this command, you need sudo root privileges. This command offers the following option: `context` *`karmada-context-name`*—The context name of the Karmada context that is associated with the management cluster. Enter the name of the context.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`sudo -E dbng rename-context --context` *`context-name`* `--new-name` *`new-name`*	Renames a context. The command does not affect the BNG CUPS Controller that is currently running on the cluster. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The old Kubernetes context name to rename. Enter the name of the context. • `new-name` *`new-name`*—The new name of the Kubernetes context (cluster name). Enter a new name. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster.
`sudo -E dbng restart --context` *`context-name`* `[--force] [--wait]` *`microservice-name`*	Restarts a specific BNG CUPS Controller service. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name on which to restart the service. Enter the name of the context. • `force`—Forcibly restart the micro-service without validating that it can be safely restarted. • `wait`—Wait for the new pod to fully come up. • *microservice-name*—Enter the microservice name to restart. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`sudo -E dbng rollout --context` *`context-name`* `[--service` *`service name`* `--version` *`software-release`*`]`	Upgrade a BNG CUPS Controller service. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name on which to roll out the new software version. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `sevice` *`service name`*—The microservice name to roll out. Enter the microservice's name. • `version` *`software-release`*—The software release to roll out. Enter the software release number.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`sudo -E dbng setup --context` *`context-name`* `[--help] [--log info] [--no color] [--bbecloudsetup] [--update] [--ssh` *`ip-address:port-number`*`] [--secrets] [--verbose] [--config` *`file-name`*`] [--template` *`file-name`*`] [--mandatory] [--optional]`	Sets up the BNG CUPS Controller application as part of the installation process. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name on which to run startup. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `h, help`—Shows the help message and exit. • `l, log [error, warning, info, debug]`—Adjusts the log level. • `no-color`—Prints messages without colors. • `bbecloudsetup`—Fills in operational parameters that align with a bbecloudsetup created cluster so that you do not have to interact with BNG CUPS Controller during the setup process (see the BBE Cloudsetup Installation Guide for cluster installation instructions). **NOTE**: Only use either the `bbecloudsetup` option or the `template` *`file-name`* option. Do not use both options. • `update`—You will only be prompted for missing values during setup. • `ssh` *`ip-address:port-number`*—Enables SSH towards the control plane instance. Enter the SSH IP address and port number on which the control plane instance is listening for SSH (when enabled in the configuration). The IP address can also be a DNS name. • `secrets`—Updates the keys, certificates, and secrets used by the BNG CUPS Controller. • `verbose`—Provides a detailed description before each prompted question. • `config` *`file-name`*—The initial configuration file that you want BNG CUPS Controller to use at startup. • `template` *`file-name`*—A YAML formatted file that contains a subset of the configuration file that is created during setup. The values that are entered in

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
	the template file are used automatically by the setup process. When you use the `template` option, you are not required to manually enter the information contained in the template file during the setup process. You should only use the `template` option when using Red Hat OpenShift Container Platform to create the cluster or when creating a multiple geography cluster. Setuop File Field Descriptions describes the information that you need to enter into the `template` configuration file. **NOTE**: Only use either the `bbecloudsetup` option or the `template` *`file-name`* option. Do not use both options. • `mandatory`—Only asks required questions during setup. • `optional`—Only asks questions that are not required during setup.
`dbng shell --context` *`context-name`* `[-p|--pipe]` *`microservice-name`*	Connects you to a running microservice. This command offers the following options: • *microservice-name*—The name of the microservice that you want to connect to. • `context`—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `pipe`—Allows you to pipe input into the command.
`sudo -E dbng start --context` *`context-name`*	Starts a specific BNG CUPS Controller service. To run this command, you need sudo root privileges. This command offers the following option: • `context` *`context-name`*—The Kubernetes context name on which to start a BNG CUPS Controller. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`dbng status --context` *`context-name`* `[-o|--output json] [--terse] [--detail]`	Displays the current status of the BNG CUPS Controller services. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `output`—Allows you to request the output in JSON format. • `terse`—Displays a summarized output of the health of the system. • `detail`—Displays information for each pod.
`sudo -E dbng stop --context` *`context-name`* `[--now]`	Stops all BNG CUPS Controller services. To run this command, you need sudo root privileges. This command offers the following option: • `context` *`context-name`*—The Kubernetes context name on which to stop a BNG CUPS Controller. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `now`—Stops the BNG CUPS Controller immediately, instead of waiting for the two minute delay.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`dbng storage --context` *`context-name`* `[-o|--output json] [--terse]`	Provides the status of the storage drivers for BNG CUPS Controller. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `output`—Allows you to request the output in JSON format. • `terse`—Displays a summarized output of the storage health.
`sudo -E dbng unlink --context` *`context-name`*	Unlink components associated with the context. To run this command, you need sudo root privileges. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name to uninstall. Enter the name of the context. • **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster.

Table 1: BNG CUPS Controller Utility Script Commands *(Continued)*

Command Name	Action
`dbng version [--context` *`context-name`*`] [-o|--output json] [--detail] [--compare] [--release` *`release-number`*`]`	Displays the version of the following: • Every running microservice in the BNG CUPS Controller instance. • The BNG CUPS Controller utility. • All available BNG CUPS Controller software releases on the system. This command offers the following options: • `context` *`context-name`*—The Kubernetes context name. Enter the name of the context. **NOTE**: If you are using this command in a multiple geography setup, you must enter the context name of the Karmada context that is associated with the management cluster. • `output`—Allows you to request the output in JSON format. • `detail`—Displays all available software versions. • `compare`—Compares the designated release to the currently running BNG CUPS Controller release. • `release` *`release-number`*—Displays microservice information for the requested release.

Use the following general syntax to issue a command:

- For a short option:

```
$ dbng command-name -option
```

- For a long option:

```
$ dbng command-name --option
```

To display a list of available commands with a brief description, use either the `h` or `help` option:

```
$ dbng -h
```

```
$ dbng --help
```

To display the options for a specific command:

```
$ dbng command-name -h
```

Start or Stop BNG CUPS Controller Services

Use the `dbng` utility script to start or stop all BNG CUPS Controller services.

- To start all BNG CUPS Controller services:

  ```
  $ sudo -E dbng rollout --context context-name
  ```

- To stop all BNG CUPS Controller services:

  ```
  $ sudo -E dbng stop --context context-name
  ```

Check the Status of BNG CUPS Controller Services

Use the `dbng status` utility script to check the status of each BNG CUPS Controller service (functional component) listed in Table 2 on page 116. The status shows whether a service is running, has exited, or has not started. It also displays the service name on the Kubernetes pod. You can compare uptime for the services to quickly see whether any service has been restarted.

Table 2: Services Displayed with the Status Command

Service	Description
`cpi-`*`label`*	The BNG CUPS Controller instance service—Implements the subscriber management control plane, which includes control plane protocols; authentication, authorization, and accounting (AAA) and supporting infrastructure. The *label* is defined by you, when you run the `dbng cpi add` command.
`scache`	The state cache service—Provides an on-cluster backing storage for subscriber service states generated by the `cp` service. Use this service for state recovery in the event of a restart of the `cp` service. In multiple geography setups, you will see two `scache` services.
`configserver`	Configuration file replication. **NOTE**: This service is only used in a multiple geography setup.
`bbe-mirp`	The Multi-geo IP Route Prioritization (MIRP) operator is used to facilitate the advertisement of the external AAA address (NAS-IP-Address) to iBGP routers in order to steer inbound AAA traffic to the active workload cluster (see the "`ip-aaa`" on page 104 option for the `cpi add` command). **NOTE**: This service is only used in a multiple geography setup.
`bbe-observer`	Observes multiple cluster scheduling events for generation number calculation. **NOTE**: This service is only used in a multiple geography setup.

To check the status of controller services, display the service status:

```
$ dbng status
```

For example:

```
user@host $ dbng status --detail --context context-name
MICROSERVICE  POD                          STATE    RESTARTS  UPTIME
NODE
scache        scache-pod-7f646d56dc-w88sg  Running  0         0:00:38.959603
example-1.juniper.net
```

Juniper BNG CUPS Logging

Juniper BNG CUPS uses the Broadband Edge (BBE) Event Collection and Visualization application for logging purposes.

BBE Event Collection and Visualization collects syslog events and records them in a time-series database. You can view the recorded events through the BBE Event Collection and Visualization Dashboard. The BBE Event Collection and Visualization Dashboard is a GUI-based visualization tool that enables you to view recorded events according to a defined filter, which can be within a specific time range. The Dashboard also provides powerful search and visualization tools through which you can correlate recoded events from multiple sources. To install BBE Event Collection and Visualization, see Broadband Edge Event Collection and Visualization Installation Guide.

Uninstall and Remove BNG CUPS Controller

Use the `dbng` utility script to uninstall the BNG CUPS Controller configuration. The `unlink` command reverts the actions you performed when setting up BNG CUPS Controller. This script returns BNG CUPS Controller to the state it was in immediately after you installed the application but before you did any setup configuration.

To uninstall BNG CUPS Controller:

1. On the jump host where you installed BNG CUPS Controller, run the `stop` command.

```
$ sudo -E dbng stop --context context-name
```

2. Run the `unlink` command.

```
$ sudo -E dbng unlink --context context-name
```

3. Run the `clean` command.

```
$ sudo -E dbng clean --uninstall
```

How to Access BNG CUPS Controller Configuration and Operational Commands

IN THIS SECTION

Access the BNG CUPS Controller CLI

You use the BNG CUPS Controller command-line interface (CLI) to configure BNG CUPS Controller and to monitor its operations. This section describes how to access the CLI.

To access the BNG CUPS Controller CLI prompt:

1. Enter the following `dbng` utility script command.

```
$ dbng cli
root@host>
```

2. Enter a question mark to see the available top-level CLI commands. This command yields a subset of the Junos OS top-level commands.

```
root@host ?
Possible completions:
  clear                Clear information in the system
  configure            Manipulate software configuration information
  help                 Provide help information
  monitor              Show real-time debugging information
  op                   Invoke an operation script
  quit                 Exit the management session
  request              Make system-level requests
  set                  Set CLI properties, date/time, craft interface message
  show                 Show system information
  start                Start shell
```

The CLI available for BNG CUPS Controller is a subset of the Junos OS CLI. For an overview of Junos OS CLI basics, see Day One: Exploring the Junos CLI. For more detailed information, see the CLI User Guide.

Access and Use CLI Configuration Statements

You use configuration statements to configure, set, manage, and monitor BNG CUPS Controller properties.

To configure BNG CUPS Controller components:

1. Use the BNG CUPS Controller utility command `dbng cli` to access the top-level CLI prompt.

```
$ dbng cli

root@host>
```

2. Access configuration mode to configure BNG CUPS Controller and the information that BNG CUPS Controller uses to configure a managed router.

```
root@user> configure
root@user#
```

3. Enter CLI statements to configure the Juniper BNG CUPS components (BNG CUPS Controller and BNG User Planes).
4. Save and activate the configuration. This command succeeds only when no configuration syntax errors exist.

```
root@user# commit
commit complete
```

5. (Optional) Exit configuration mode and return to the top-level CLI prompt.

```
root@user# exit
root@user>
```

For a list of supported configuration statements, see "Juniper BNG CUPS Controller CLI Configuration Statements" on page 124.

Access and Use CLI Operational Commands

You use operational commands to display the current status of Juniper BNG CUPS. You enter operational commands to monitor and to troubleshoot the BNG CUPS Controller and the BNG User Planes.

To monitor BNG CUPS Controller, view BNG CUPS Controller configuration and statistics, or run certain operations manually:

1. Use the BNG CUPS Controller utility command `dbng cli` to access the top-level CLI prompt.

   ```
   $ dbng cli

   root@host
   ```

2. Enter specific commands.

   - Use `show` commands to display statistical information.

   - Use `request` commands to manually initiate certain BNG CUPS Controller operations.

   For a list of supported operational commands, see "Juniper BNG CUPS Controller CLI Operational Commands" on page 166.

Troubleshooting and Monitoring Juniper BNG CUPS

IN THIS SECTION

Many of the existing mechanisms for troubleshooting an MX Series BNG are available for troubleshooting Juniper BNG CUPS. Most of the BNG functionality is on the BNG CUPS Controller; therefore, you perform the majority of the troubleshooting on the BNG CUPS Controller.

The following troubleshooting mechanisms are available on the BNG CUPS Controller:

- Tracelogs

- Shared memory logs

- Operational and troubleshooting commands for the following components:
  - Node management
  - DHCP and DHCPv6
  - PPP
  - L2TP
  - DVLAN
  - AAA
  - Subscriber management
  - Subscriber groups
  - Load balancing groups

The following troubleshooting mechanisms are available on the BNG User Plane:

- Tracelogs
- Shared memory logs
- Operational and troubleshooting commands for the following components:
  - Node management
  - Subscriber management

Centralized Logging Using Broadband Edge Event Collection and Visualization

Broadband Edge (BBE) Event Collection and Visualization is an event collection application that is meant to operate with Juniper's Broadband Edge cloud applications, such as Juniper BNG CUPS Controller and Address Pool Manager (APM).

BBE Event Collection and Visualization collects syslog events and records them in a time-series database. You can view the recorded events through the BBE Event Collection and Visualization Dashboard. The BBE Event Collection and Visualization Dashboard is a GUI-based visualization tool that enables you to view recorded events according to a defined filter, which can be within a specific time range. The Dashboard also provides powerful search and visualization tools through which you can correlate recoded events from multiple sources (for example, from APM or from the Kubernetes cluster).

BBE Event Collection and Visualization can be installed in the same Kubernetes cluster as the Juniper BBE applications (BNG CUPS Controller and APM). The installation follows the same model (installed from the jump host). For BBE Event Collection and Visualization installation instructions, see Broadband Edge Event Collection and Visualization Installation Guide.

3

CHAPTER

Juniper BNG CUPS Controller CLI Configuration Statements

IN THIS CHAPTER

Juniper BNG CUPS Controller CLI Configuration Statements

IN THIS SECTION

This topic provides an overview of configuration commands, including syntax and option descriptions, that you use with Juniper BNG CUPS Controller.

address-pool-manager

IN THIS SECTION

Syntax

```
address-pool-manager {
     inet ip-address;
     port port-number;
     secondary-address{
       inet ip-address;
       port port-number;
     local-reserve{
         partition partition-name{
               prefix ipv4-prefix;
         }
         auto-reclamation {
             drain-delay number;
         }
         apportion-delay number;
         }
     secrets {
         certificate certificate-file;
         key         private-key-file;
         ca-cert     cacertificate-file;
      }
```

```
    }
]
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| access]
```

Description

Configures Juniper Address Pool Manager's (APM) connection to Juniper BNG CUPS. This configuration is done on the Juniper BNG Controller.

Options

`inet` *`ip-address`*	APM's IPv4 address. For a multiple geography, redundant APM setup, this is the IP address for the primary workload cluster.
`port` *`port-number`*	The port that APM is listening on for incoming APM connections. For a multiple geography, redundant APM setup, this is the port number for the primary workload cluster.
`secondary-address`	The secondary address used for a multiple geography, redundant APM setup. Enter the following for the `secondary-address`: • APM's IPv4 address. This is the IP address for the backup workload cluster. • The port that APM is listening on for incoming APM connections. This is the port number for the backup workload cluster.
`local-reserve`	A BNG CUPS Controller configured set of partitions. Partitions can be either IPv4 or IPv6 addresses.
`Partition` *`partition-name`*	The configured partion.
`prefix` *`ipv4-prefix`*	Specify an IPv4 prefix to include in the partition.
`drain-delay` *`number`*	Specifies a hold down time to wait after reconnecting with APM, to start draining local pools.
`apportion-delay` *`number`*	Specifies a hold down time to wait before entering the local apportionment mode, following the loss of connectivity with APM.

`secrets` If the gRPC Network Management Interface (gMI) connection is secured, configure any Transport Layer Security (TLS) keys, as follows:

- `certificate` *`certificate-file`*
- `key` *`private-key-file`*
- `ca-cert` *`ca-certificate-file`*

bng-controller

IN THIS SECTION

Syntax

```
bng-controller {
        bng-controller-name bng-cups-controller-name;
        security-profiles security-profile-name {
               ca-cert-file-name ca-certificate-name;
               cert-file-name    certificate-name;
               key-file-name     key-name;
         }
         user-plane {
            bng-user-plane-name {
                transport {
                    inet ip-address;
                    inet6 ip-address;
                    security-profile security-profile-name {
                 }
                 dynamic-address-pools {
                     partion partition-name;
```

```
                }
                user-plane-profile bng-user-plane-profile-name;
            }
        }
        control-plane-instances {
            control-plane-instance-name {
               control-plane-config-group control-plane-config-group-name;
               user-plane bng-user-plane-name;
             }
         }
        subscriber-group  {
        }
        load-balancing-groups {
        }
      }
   }
 }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director]
```

Description

Configures system wide attributes for the BNG CUPS Controller, such as the `user-plane` configuration and the `control-plane-instances` configuration.

Configures the BNG CUPS Controller system.

Options

bng-controller-name *bng-cups-controller-name*	The *bng-cups-controller-name* is a mandatory reference to the local system and can be 1 to 12 characters long. You can combine uppercase letters and lowercase letters, numbers, hyphens, and periods in this reference but cannot start or end it with a hyphen.
`security-profile`	Secify a security profile. See "security-profiles" on page 146.
`user-plane`	Specify the BNG User Planes to be associated with the BNG CUPS Controller. See "user-plane (bng-controller)" on page 161.

`control-plane-instances`	See "control-plane-instances" on page 131.
`subscriber-group`	See "subscriber-group" on page 152.
`load-balancing-groups`	See "load-balancing-groups" on page 136.

captive-portal-content-delivery

IN THIS SECTION

Syntax

```
captive-portal-content-delivery {
    profile name
        dynamic;
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| services]
```

Description

Configure the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber. Use the statement at the `[edit groups bbe-cpi-|0 - 4| services]` hierarchy level for converged services CPCD. This command runs on the BNG CUPS Controller.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

captive-portal-content-delivery-profile

IN THIS SECTION

Syntax

```
captive-portal-content-delivery-profile profile-name;
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| services service-set service-set-name]
```

Description

Configure converged HTTP redirect services on the Routing Engine. This command runs on the BNG CUPS Controller.

Options

`captive-portal-content-delivery-profile` *`profile-name`*—Name of the CPCD profile.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced before Juniper BNG CUPS Release 23.1.

control-plane-instances

IN THIS SECTION

Syntax

```
control-plane-instances{
   control-plane-instance-name{
       control-plane-config-group control-plane-config-group-name;
       user-plane user-plane-name;
```

```
    }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

Control plane instances to which BNG User Planes are mapped. This mapping enables you to easily move BNG User Planes from one control plane instance to another to adapt to changing scaling or use case demands. A control plane instance is assigned to a control plane configuration group.

Options

control-plane-instance-name	Name of the control plane instance.
control-plane-config-group *control-plane-config-group-name*	Specify the name of an existing control plane configuration group from which the control plane instance obtains its configuration.
user-plane *user-plane-name*	Specify the name of a BNG User Plane assigned to the control plane instance. You can assign more than one BNG User Plane to a control plane instance.

domain-profile

IN THIS SECTION

Syntax

```
domain-profile domain-profile-name{
family{
     inet ip-address | inet6 ip-address{
        partion-type [delegated-prefix | non-temporary-address | router-advertisement];
        preferred-prefix-length number;
        allocation-length number;
        install-discard-routes{
           tag number;
           backup-tag number;
     }
     source-partition-qualifier string;
     excluded-address last-octet number;
     dhcp-gateway-address-last-octet number;
     protocol-attributes dhcp-attribute;
    }
 }
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| access address-assignment]
```

Description

Configures the domain profile. The domain profile defines the BNG attributes for creating domains. The domain is created based on the framed pool received from RADIUS.

Options

`domain-profile` *`domain-profile-name`*	Set the name of the domain profile.
`family`	Specify an address family protocol. Specify `inet` for IPv4 addresses. Specify `inet6` for IPv6 addresses.
`partion-type [delegated-prefix | non-temporary-address | router-advertisement]`	Only applies to `inet6` configurations. The setting corresponds to what is entered in the `groups bbe-bng-director bng-controller user-planes` *`user-plane-name`* `dynamic-address-pools` setting.

preferred-prefix-length *number*	Define the preferred prefix length. • **Range:** 8 through 30
allocation-length *number*	Define the allocation length of the IPv6 address or prefix that are assigned from the dynamic address pool prefix to the subscriber. This option is supported only for the inet6 address family.
source-partition-qualifier *string*	(Optional) A string that is applied as a suffix to the domain's location, to create a partition name that is passed to Juniper Address Pool Manager.
excluded-address last-octet *number*	(Optional) When you configure the preceding code phrase, the domain profile excludes all addresses with a domain pool prefix that matches the specified last-octet value. This option is supported only for the inet address family. • **Range:** 0 through 255
dhcp-gateway-address-last-octet *number*	specifies the value of the last byte to reserve in each dynamic pool prefix to be used as the DHCP gateway address for the DHCP Local Server. For example, if the dynamically allocated pool prefix is 192.32.6.0/24 and `dhcp-gateway-address-last-octet` is set to 1, the system would reserve and program 192.32.6.1 as the DHCP gateway address.
protocol-attributes *dhcp-attribute*	Specifies the name of the protocol attributes profile that defines the DHCP attributes to use for dynamic pools created in the domain.
install-discard-routes tag *number* **backup-tag** *number*	(Optional) Indicates that you must configure a discard route (with the associated route tag supplied with the pool prefix) separately on the BNG User Planes to import these routes into the exported route set. Valid route tags are 0..2^(32-1)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 23.4R1.

igmp

IN THIS SECTION

Syntax

```
igmp {
    interface interface-name {
        (accounting | no-accounting);
        disable;
        distributed;
        group-limit limit;
        group-policy [ policy-names ];
        immediate-leave;
        oif-map map-name;
        passive;
        promiscuous-mode;
        ssm-map ssm-map-name;
        ssm-map-policy ssm-map-policy-name;
        static {
            group multicast-group-address {
                exclude;
                group-count number;
                group-increment increment;
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
```

```
        version version;
    }
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles profile-name protocols]
```

Description

Enable IGMP on an interface and configure interface-specific properties. This command runs on the BNG CUPS Controller.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

load-balancing-groups

IN THIS SECTION

Syntax

```
load-balancing-groups {
   group-name;
     user-plane bng-user-plane-name{
     weight-based-mode{
      port port-number{
         max-weight max-weight-number;
         preferred;
     }
      report-based-mode
       port port-number{
         preferred;
     ]
 }
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

Enables load balancing on Juniper BNG CUPS. This command runs on the BNG CUPS Controller.

Options

group-name	Specify the load balancing group name.
user-plane ***bng-user-plane-name***	Specify the BNG User Plane that is associated with the BNG CUPS Controller for load balancing.

`weight-based-mode`	Used to configure weight-based load balancing.
`report-based-mode`	Used to configure report-based load balancing.
`port` *`port-identifier`*	Specify the logical port that is associated with the BNG CUPS Controller load balancing. You use the format **up:*user-plane-name*:*physical-port-name***.
`max-weight` *`max-weight-number`*	Specify the maximum weight value (1 through 255) for the logical port.

mld

IN THIS SECTION

Syntax

```
mld {
   interface interface-name {
       (accounting | no-accounting);
       disable;
       distributed;
       group-limit  limit;
       group-policy [ policy-names ];
       group-threshold value;
       immediate-leave;
       log-interval seconds;
       oif-map [ map-names ];
       passive;
       ssm-map ssm-map-name;
```

```
        ssm-map-policy ssm-map-policy-name;
        static {
            group multicast-group-address {
                exclude;
                group-count number
                group-increment increment
                source ip-address {
                    source-count number;
                    source-increment increment;
                }
            }
        }
        version version;
    }
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| dynamic-profiles profile-name protocols]
```

Description

Enable MLD on an interface and configure interface-specific properties. This command runs on the BNG CUPS Controller.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

overrides

IN THIS SECTION

Syntax

```
overrides {
    no-unsolicited-ra;
    statistics-reporting-interval seconds;
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| user-plane-profiles user-plane-profile-name]
```

Description

Override the default configuration settings for the enhanced subscriber management software for subscriber management.

Options

`statistics-reporting-interval` *seconds*	The interval at which statistics are reported from a BNG User Plane to the BNG CUPS Controller. The statistics reporting interval is reported in seconds. • **Default:** 60 seconds • **Range:** 60 through 1440 seconds
`no-unsolicited-ra`	Disable the default transmission and periodic refresh of unsolicited Router Advertisement messages by the router when the subscriber interface is created, and at configured periodic intervals thereafter. When you include the `no-unsolicited-ra` statement, the router sends Router Advertisement messages and associated periodic refresh messages only when it receives a Router Solicitation message from the subscriber.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

partition

IN THIS SECTION

Syntax

```
partition partition-name;
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller user-planes bng-user-plane-name dynamic-address-pools-plane]
```

Description

Defines the BNG User Plane partition attribute. The partition attribute defines the source partition defined by a dynamic address pool source. The dynamic pool source can be either APM or "local-reserve" on page 126.

NOTE: For Juniper BNG CUPS to operate with Juniper Address Pool Manager, you must configure the partition attribute.

Options

`partition` *`partition-name`*	Name of the partition.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

pfcp

IN THIS SECTION

Syntax

```
pfcp {
   retransmission-timer seconds;
   retries number;
   heartbeat-interval seconds;
   enable-tracing
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| user-plane-profiles user-plane-profile-name]
```

Description

Sets the Packet Forwarding Control Protocol (PFCP) protocol attributes for the control plane manager and any other daemons using Packet Forwarding Control Protocol to communicate with their peers.

Options

`pfcp` Specify the Packet Forwarding Control Protocol protocol attributes.

> **NOTE**: We recommend that you configure the BNG CUPS Controller and the BNG User Planes with the same Packet Forwarding Control Protocol attributes.

- `retransmission-timer`—Defines the retransmission interval in seconds.
  - **Default:** 5 seconds
  - **Range:** 3 through 30

    seconds
- `retries`—Defines the number of retransmission attempts.
  - **Default:** 5
  - **Range:** 5 through 10

- `heartbeat-interval`—Defines the interval in seconds between keep-alive messages.
  - **Default:** 60

    seconds
  - **Range:** 60 through 600

    seconds

prefix-advertisement

IN THIS SECTION

Syntax

```
prefix-advertisement {
    family inet {
        tag number;
        backup-tag number;
        prefix ip-address;
    {
    family inet6 {
        tag number;
        backup-tag number;
        prefix ip-address;
    }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller subscriber-group subscriber-group-name subscriber-group-name routing-instances routing-instance-name]
```

Description

Supports the DHCP relay L3 wholesale model. You configure a dedicated routing instance for a retailer and assign it to a subscriber group. You use the `prefix-advertisement` command to configure the IPv4 or IPv6 prefixes, that represent the pool of IPv4 or IPv6 addresses and prefixes assigned to the subscriber by the retailer. You can also configure optional route tags for route advertisement.

Options

`family inet`—You can configure the following:

- `tag` *number*—(Optional) Set the active route tag.
- `backup-tag` *number*—(Optional) Set the backup route tag.
- `prefix` *ip-address*—Set the prefix IPv4 Address.

`family inet6`—You can configure the following:

- `tag` *number*—(Optional) Set the active route tag.
- `backup-tag` *number*—(Optional) Set the backup route tag.
- `prefix` *ip-address*—Set the prefix IPv6 Address.

routing-engine-services

IN THIS SECTION

Syntax

```
routing-engine-services;
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| services service-set service-set-name service-set-options]
```

Description

When configuring a Routing Engine-based captive portal service, specify the service set options to apply to a service set. This command runs on the BNG CUPS Controller.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 23.1.

security-profiles

IN THIS SECTION

Syntax

```
security-profiles profile-name{
   ca-cert-file-name ca-certificate-name;
   cert-file-name certificate-name;
   key-file-name key-name;
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

Defines one or more profiles that specify security requirements to secure the BNG CUPS Controller channels to the BNG User Planes using Data Transport Layer Security and Transport Layer Security. If the **security-profiles** is not configured, the related BNG CUPS Controller or BNG CUPS User Plane assumes that the transport interfaces are not secure.

Options

security-profiles ***profile-name***	Give the security profile a name.
ca-cert-file-name ***ca-certificate-name***	Name of the CA profile.
cert-file-name ***certificate-name***	Name of the public certificate.
key-file-name ***key-name***	Name of the private key pair.

service-set

IN THIS SECTION

Syntax

```
service-set service-set-name {
    service-set-options {
        routing-engine-services;
    }
    captive-portal-content-delivery-profile profile-name;
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| service-set service-set-name]
```

Description

Define one or more service sets. Service sets are applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. You can use the predefined dynamic interface variables `$junos-input-service-set`, `$junos-output-service-set`, `$junos-input-ipv6-service-set`, and `$junos-output-ipv6-service-set`. This command runs on the BNG CUPS Controller.

Options

service-set-name	Name of the service set.
`service-set-options`	Specify service set options (see "service-set-options" on page 149).
`routing-engine-services`	Specify routing engine services configuration (see "routing-engine-services" on page 145).
`captive-portal-content-delivery-profile` *profile-name*	Name of the CPCD profile (see "captive-portal-content-delivery-profile" on page 130).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 23.1R1.

service-set-options

IN THIS SECTION

Syntax

```
service-set-options {
        routing-engine-services;
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| service-set service-set-name]
```

Description

Specify the service set options to apply to a service set. This command runs on the BNG CUPS Controller.

Options

routing-engine-services	Specify routing engine services configuration (see "routing-engine-services" on page 145).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 23.1R1.

static-address-pools

IN THIS SECTION

Syntax

```
static-address-pools {
    pool-name {
        tag number;
        backup-tag number
    }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller subscriber-group subscriber-group-name subscriber-
group-name routing-instances default]
```

Description

Sets the assignment of a statically configured address pool used to assign subscriber IPv4 or IPv6 addresses for a subscriber group. The address pool can also be configured with optional route tags on a network instance basis. For each specified static pool for a specified routing instance, a discard route can be derived from the pool network or prefix configuration for route advertisement.

Options

pool-name	Specify the name of the address pool.
tag *number*	(Optional) Set the active route tag.
backup-tag *number*	(Optional) Set the backup route tag.

subscriber-group (control-plane-managed-mode)

IN THIS SECTION

Syntax

```
subscriber-group
 subscriber-group-name subscriber-group-name{
   virtual-mac mac-address;
   control-plane-managed-mode{
    preferred-user-plane-name user-plane-name;
    redundancy-interface alpha{
      logical-port up:user-plane-name:logical-port-name,up:user-plane-name:logical-port-name;
    }
    redundancy-interface beta {
      logical-port up:user-plane-name:logical-port-name,up:user-plane-name:logical-port-name;
    }
   }
   user-plane user-plane-name
```

```
        backup-mode [hot | warm-enhanced | warm-scale]
    }
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

A group of subscribers. Subscriber sessions that are subject to the same restoration capability are placed into the same subscriber group. Grouping subscribers together helps to increase core routing efficiency. The use of subscriber groups minimizes the messaging, which reduces the elapsed time between the detection of a failure (or any request to switchover from active to backup) and the restoration of the service.

Options

`subscriber-group-name` *subscriber-group-name*	Specify the subscriber group name.
virtual-mac *mac-address*	A logical MAC address assigned to the subscriber group that is used for all communication between the BNG CUPS Controller and the subscriber sessions assigned to the subscriber group. This ensures that the same MAC address is used by the BNG CUPS Controller for communication with subscriber sessions, irrespective of which BNG User Plane is currently active for the subscriber group. A virtual MAC address is required for a resiliency subscriber group.
`control-plane-managed-mode`	Establishes that the BNG CUPS Controller determines which BNG User Plane is the active one for a resiliency subscriber group
`preferred-user-plane-name` *user-plane-name*	When operating in control plane managed mode for a resiliency subscriber group, it establishes which of the member BNG User Planes is the preferred active BNG User Plane.
`redundancy-interface alpha`	• `logical-port up:`*user-plane-name*`:`*logical-port-name* Configures a named set of logical ports on a BNG User Plane that is assigned to the subscriber group.
`redundancy-interface beta`	• `logical-port up:`*user-plane-name*`:`*logical-port-name*

For a resiliency subscriber group, you configure a named set of logical ports on the BNG User Planes that are assigned to the subscriber group. The two redundancy interfaces form a resiliency subscriber group, that for `control-plane-managed-mode`, the `preferred-user-plane-name` establishes which of the two BNG User Planes is the preferred active BNG User Plane.

`user-plane` *`user-plane-name`*

- `backup-mode`

  Set the BNG User Plane backup level. Choose from the following options:

  - `hot`—The BNG CUPS Controller continually programs all subscriber session state information on the backup BNG User Planes, enabling undetectable hitless failover.

  - `warm-enhanced`—The backup BNG User Plane has a limited subscriber forwarding state installed and the full subscriber session state maintained in memory. If an active BNG User Plane fails, the backup assumes forwarding of subscriber sessions and then installs the remaining subscriber state. There is a short time frame until the subscriber session SLA is restored. After convergence, the `warm-enhanced` backup mode has the same packet flow recovery as the `hot` backup mode. Services are programmed in the forwarding engine after packet flow recovery occurs. This approach supports N:1 redundancy.

    NOTE: The `warm-enhanced` backup mode is only available for backup devices that are based on Juniper Trio 5 and 6 chipsets.

  - `warm-scale`—Uses a Routing Engine based backup mode. The backup BNG User Plane holds the full subscriber state on the Routing Engine and doesn't program the Packet Forwarding Engine until the subscriber group becomes active. This approach supports M:N redundancy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

subscriber-group (user-plane-managed-mode)

IN THIS SECTION

Syntax

```
subscriber-group
 subscriber-group-name subscriber-group-name{
   virtual-mac mac-address;
   user-plane-managed-mode{
    redundancy-interface alpha{
      logical-port up:user-plane-name:logical-port-name,up:user-plane-name:logical-port-name;
    }
   }
   user-plane user-plane-name
    backup-mode [hot | warm-enhanced | warm-scale]
  }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

A group of subscribers. Subscriber sessions that are subject to the same restoration capability are placed into the same subscriber group. Grouping subscribers together helps to increase core routing efficiency. The use of subscriber groups minimizes the messaging, which reduces the elapsed time between the detection of a failure (or any request to switchover from active to backup) and the restoration of the service.

Options

`subscriber-group-name` *subscriber-group-name*	Specify the subscriber group name.
virtual-mac *mac-address*	A logical MAC address assigned to the subscriber group that is used for all communication between the BNG CUPS Controller and the subscriber sessions assigned to the subscriber group. This ensures that the same MAC address is used by the BNG CUPS Controller for communication with subscriber sessions, irrespective of which BNG User Plane is currently active for the subscriber group. A virtual MAC address is required for a resiliency subscriber group.
`user-plane-managed-mode`	Establishes that the BNG User Plane determines which BNG User Plane is the active one for a resiliency subscriber group
`redundancy-interface alpha`	• `logical-port up:`*user-plane-name*`:`*logical-port-name* Configures a named set of logical ports on a BNG User Plane that is assigned to the subscriber group.
`user-plane` *user-plane-name*	• `backup-mode` Set the BNG User Plane backup level. Choose from the following options: • `hot`—The BNG CUPS Controller continually programs all subscriber session state information on the backup BNG User Planes, enabling undetectable hitless failover. • `warm-enhanced`—The backup BNG User Plane has a limited subscriber forwarding state installed and the full subscriber session state maintained in memory. If an active BNG User Plane fails, the backup assumes forwarding of subscriber sessions and then installs the remaining subscriber state. There is a short time frame until the subscriber session SLA is restored. After convergence, the `warm-enhanced` backup mode has the same packet flow recovery as the `hot` backup mode. Services are programmed in the forwarding engine after packet flow recovery occurs. This approach is typically used to support N:1 redundancy.

> **NOTE**: The `warm-enhanced` backup mode is only available for backup devices that are based on Juniper Trio 5 and 6 chipsets.

- `warm-scale`—Uses a Routing Engine based backup mode. The backup BNG User Plane holds the full subscriber state on the Routing Engine and doesn't program the Packet Forwarding Engine until the subscriber group becomes active. This approach supports M:N redundancy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

subscriber-group-default-tag

IN THIS SECTION

Syntax

```
subscriber-group-default-tag {
 tag number;
 backup-tag number;
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4| access address-assignment]
```

Description

Configures active and backup global tags for subscriber groups.

Options

`tag` *number*	Set the global active tag.
`backup-tag` *number*	Set the global backup tag.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 23.4R1.

transport

IN THIS SECTION

Syntax

```
transport {
   inet ip-address;
   security-profile security-profile-name;
}
```

Hierarchy Level

```
[edit groups bbe-bng-dirctor bng-controller user-planes bng-user-plane-name]
```

Description

Defines the transport security for all BNG CUPS Controller and BNG User Plane inter-communication. You use the transport command to configure either the BNG CUPS Controller or the BNG User Planes, depending on which option you choose at the `mode` level of the hierarchy.

Options

`inet` *`ip-address`*	The IP address of either the BNG CUPS Controller or the BNG User Plane that you are configuring.
`security-profile` *`security-profile-name`*	Specify the configured security profile that lists the CA profile, public certificate, and private key pair.

user-plane-profiles

IN THIS SECTION

Syntax

```
user-plane-profiles {
    user-plane-profile-name {
                capabilities {
                   hardware-family (juniper-mx | juniper-acx)
                }
                pfcp {
                    retransmission-timer seconds;
                    retries number;
                    heartbeat-interval seconds;
                }
                interfaces interface-name {
                     auto-configure {
                        stacked-vlan-ranges {
                            dynamic-profile <dynamic-profile-name> {
                                accept any;
                                ranges {
                                    any,any;
                                }
                            }
                        }
                        remove-when-no-subscribers;
                    }
                 }
               }

    }
}
```

Hierarchy Level

```
[edit groups bbe-cpi-|0 - 4|]
```

Description

The `user-plane-profiles` statement uses a template to configuring a BNG User Plane in terms of interfaces, pfcp behavior, and subscriber management override behavior. The `user-plane-profiles` statement is specified as part of the BNG User Plane configuration in the `bbe-bng-director` configuration group. It is defined in the common group configuration that is part of the `control-plane-instance` configuration. When a BNG User Plane is assigned to a `control-plane-instance` during its configuration, the BNG User Plane's `user-plane-profile` must be defined in the control plane instance's `control-plane-config-group` `(bbe-cpi-|0 - 4|)`.

Options

`user-plane-profile-name`	Name of the user plane profile.

`hardware-family`	Specify configuration pertaining to the capabilities of the BNG User Plane type. Currently, only the `juniper-mx` BNG User Plane type is supported. You can assign this profile to all BNG User Planes with the same characteristics and use case.
`pfcp`	Specify the PFCP configuration to be used for the BNG User Plane (see "pfcp" on page 142).
`interfaces` *`interface-name`*	Specify interfaces configuration to be used for the BNG User Plane (see .

user-plane (bng-controller)

IN THIS SECTION

Syntax

```
user-plane {
    bng-user-plane-name {
            transport {
                inet ip-address;
                security-profile security-profile-name;
             }
             dynamic-address-pools {
                 partion partition-name;
             }
             user-plane-profile bng-user-plane-profile-name;

    }
}
```

Hierarchy Level

```
[edit groups bbe-bng-director bng-controller]
```

Description

Define the BNG User Planes that are authorized to associate with the BNG CUPS Controller. You must list each BNG User Plane.

Options

user-plane-name	Name of the BNG User Plane.
transport	Specify transport information. See "transport" on page 158.
dynamic-address-pools	Specify the dynamic address pool related configuration. You should at least configure the partition name.
partition *partition-name*	The partition from which IPv4 and IPv6 addresses and prefixes are assigned.
user-plane-profile *user-plane-profile-name*	Specify one or more user plane profiles. See "user-plane-profiles" on page 159.

weight

IN THIS SECTION

Syntax

```
weight weight-number;
```

Hierarchy Level

```
[edit groups bbe-cpi-|1 - 4| dynamic-profiles dynamic-profiles-name interfaces interface-name
unit interface-unit load-balance]
[edit dynamic-profiles dynamic-profiles-name interfaces interface-set interface-set-name load-
balance]
```

Description

Sets the load balancing weight for either subscribers or the logical interface set.

You can define weight based on your needs. You can define it by using subscriber bandwidth, logical interface set bandwidth, or an even number of subscribers per logical interface set. This command runs on the BNG CUPS Controller.

Options

`weight` *`weight-number`*	Defines the load balancing weight value (1 through 255).

Required Privilege Level

root—To view this statement in the configuration.

root—To add this statement to the configuration.

4

CHAPTER

Juniper BNG CUPS Controller CLI Operational Commands

IN THIS CHAPTER

Juniper BNG CUPS Controller CLI Operational Commands

IN THIS SECTION

This topic provides an overview of `clear`, `request`, `restart`, and `show` commands, including syntax, option descriptions, and sample output. You use these commands with Juniper BNG CUPS Controller.

clear user-plane ipv6 router-advertisement

IN THIS SECTION

Syntax

```
clear ipv6 router-advertisement up-name
<up-name user-plane-name>
```

Description

Clear IPv6 router advertisement counters.

Options

up-name ***user-plane-name*** Clear IPv6 router advertisement counters for the specified BNG User Plane.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear user-plane ipv6 router-advertisement up-name

```
user@host> clear user-plane ipv6 router-advertisement up-name up1-example
```

clear user-plane pppoe lockout

IN THIS SECTION

Syntax

```
clear user-plane pppoe lockout
<up-name user-plane-name>
```

Description

Clear the lockout condition for the PPPoE client associated with the specified BNG User Plane.

Options

up-name *user-plane-name*	Clear the lockout condition for the PPPoE clients associated with the specified BNG User Plane.

Required Privilege Level

clear

Sample Output

clear use-plane pppoe lockout up-name

```
user@host> clear user-plane pppoe lockout up-name up-test-1
```

clear user-plane pppoe statistics

IN THIS SECTION

Syntax

```
clear user-plane pppoe statistics
<up-name user-plane-name>
```

Description

Reset PPPoE session statistics information.

Options

up-name ***user-plane-name*** Reset PPPoE statistics for the specified BNG User Plane.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear user-plane pppoe statistics up-name

```
user@host> clear user-plane pppoe statistics up-name up-test1
```

clear user-plane statistics

IN THIS SECTION

Syntax

```
clear user-plane statistics
<up-name user-plane-name>
```

Description

Clear subscriber-management statistics.

Options

up-name *user-plane-name* Clear subscriber-management statistics for the specified BNG User Plane.

Required Privilege Level

view and system

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear user-plane statistics up-name

```
user@host> clear user-plane statistics up-name up1-example
```

request network-access aaa address-assignment domain-profile

IN THIS SECTION

Syntax

```
request network-access aaa address-assignment domain-profile profile-name profile-name ri-name
routing-instance-name [enable-logins | disable-logins]
```

Description

Enable or disable logins for existing domains created from the domain profile and to control the creation of new domains from the domain profile.

Options

`ri-name` *routing-instance-name*	Specify the routing instance name.
`profile-name` *profile-name*	Specify the name of the profile.
`[enable-logins | disable-logins]`	Specify the desired action for enabling logins.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request network-access aaa address-assignment subscriber-group

IN THIS SECTION

Syntax

```
request network-access aaa address-assignment subscriber-group subscriber-group-name [enable-
logins | disable-logins]
```

Description

Enable or disable logins for a particular subscriber group.

Options

subscriber-group-name	Specify the name of the subscriber group.
`[enable-logins	disable-logins]`	Specify the desired action for enabling logins.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request network-access aaa address-assignment user-plane

IN THIS SECTION

Syntax

```
request network-access aaa address-assignment user-plane user-plane-name [enable-login | disable-
login]
```

Description

Enable or disable logins for subscribers originating from the specified BNG User Plane. When you use this command, you effectively enable or disable logins for existing domains associated with the BNG User Plane. You also control the creation of new domains for the BNG User Plane.

Options

`user-plane` *`user-plane-name`*	Specify the BNG User Plane name.
`[enable-login	disable-login]`	Specify the desired action.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request subscriber-group *subscriber-group-name* user-plane *user-plane-name* [enable | disable]

IN THIS SECTION

Syntax

```
request subscriber-group
<subscriber-group-name user-plane user-plane-name [enable | disable]>
```

Description

Disable or enable the BNG User Plane port for a BNG User Plane managed subscriber group.

Options

subscriber-group-name	The subscriber group name for which you want to enable or disable the subscriber port.
`user-plane` *user-plane-name*	The subscriber group's BNG User Plane for which you want to either disable or enable the subscriber port.
`[enable	disable]`	Enable or disable the port for the subscriber group’s BNG User Plane.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

Sample Output

request subscriber-group *subscriber-group-name*user-plane *user-plane-name*

```
user@host> request subscriber-group SGRP1 user-plane boston disable
```

request subscriber-group switchover

IN THIS SECTION

Syntax

```
request subscriber-group switchover subscriber-group-name
```

Description

Activate or deactivate a subscriber group on a BNG User Plane. You use this command to switch between active and backup BNG User Planes. This command runs on the BNG CUPS Controller.

Options

subscriber-group-name Specify the subscriber group name that you want to make the active BNG User Plane.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request user-plane maintenance associate serviced-user-plane

IN THIS SECTION

Syntax

```
request user-plane maintenance associate serviced-user-plane user-plane-name serviced-port port-number backup-user-plane user-plane-name backup-port port-number
```

Description

Creates a backup of a BNG User Plane. You can run this command multiple times for each logical port active and backup pair.

Options

`serviced-user-plane` *`user-plane-name`*	Specify the serviced BNG User Plane name.
`serviced-port` *`port-number`*	Specify the serviced port number.
`backup-user-plane` *`user-plane-name`*	Specify the backup BNG User Plane name.
`backup-port` *`port-number`*	Specify the backup port number.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request user-plane maintenance complete serviced-user-plane

IN THIS SECTION

Syntax

```
request user-plane maintenance complete serviced-user-plane user-plane-name
```

Description

Completes the maintenance operation for a BNG User Plane. The command ensures that all resources that were used for the maintenance operation are restored.

Options

`serviced-user-plane` *`user-plane-name`*	Specify the BNG User Plane name that was serviced as part of the maintenance operation.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request user-plane maintenance disassociate serviced-user-plane

IN THIS SECTION

Syntax

```
request user-plane maintenance disassociate serviced-user-plane user-plane-name
request user-plane maintenance disassociate serviced-user-plane user-plane-name serviced-port
port-number backup-user-plane user-plane-name backup-port port-number
```

Description

Remove the active and backup BNG User Plane association and remove the database synchronization.

Options

`serviced-user-plane` *`user-plane-name`*	Specify the serviced BNG User Plane name.
`serviced-port` *`port-number`*	Specify the serviced port number.
`backup-user-plane` *`user-plane-name`*	Specify the backup BNG User Plane name.
`backup-port` *`port-number`*	Specify the backup port number.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

request user-plane maintenance switchover serviced-user-plane

IN THIS SECTION

Syntax

```
request user-plane maintenance switchover serviced-user-plane user-plane-name
```

Description

Switch the role of the active and the backup BNG User Planes for the logical port pairing.

Options

`serviced-user-plane` *`user-plane-name`*	Specify the serviced BNG User Plane name.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback only if an error occurs.

restart authentication-service

IN THIS SECTION

Syntax

```
restart authentication-service
```

Description

Restarts the authentication service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart authentication-service

```
user@host> restart authentication-service
General authentication process started, pid 1687717
```

restart cpm-service

IN THIS SECTION

Syntax

```
restart cpm-service
```

Description

Restarts the Control Plane Manager daemon.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart cpm-service

```
user@host> restart cpm-service
Control Plane Manager for dBNG started, pid <process-id>
```

restart dhcp-service

IN THIS SECTION

Syntax

```
restart dhcp-service
```

Description

Restarts the DHCP service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart dhcp-service

```
user@host> restart dhcp-service
Junos Dynamic Host Configuration Protocol process started, pid 1687787
```

restart dynamic-flow-capture

IN THIS SECTION

Syntax

```
restart dhcp-service
```

Description

Restarts the dynamic flow capture service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart dynamic-flow-capture

```
user@host> restart dynamic-flow-capture
Dynamic flow capture service started, pid 1688041
```

restart gtp-proxy-service

IN THIS SECTION

Syntax

```
restart gtp-proxy-service
```

Description

Restarts the GPRS tunneling protocol proxy service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart l2tp-service

```
user@host> restart gtp-proxy-service
Control Plane GTP Protocol Proxy started, pid 1688164
```

restart l2tp-service

IN THIS SECTION

Syntax

```
restart l2tp-service
```

Description

Restarts the L2TP service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart l2tp-service

```
user@host> restart l2tp-service
Universal edge Layer 2 Tunneling Protocol daemon started, pid 1687917
```

restart pfcp-proxy-service

IN THIS SECTION

Syntax

```
restart pfcp-proxy-service
```

Description

Restarts the PFCP proxy service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart pfcp-proxy-service

```
user@host> restart pfcp-proxy-service
Control Plane PFCP Protocol Proxy started, pid 1688195
```

restart ppp-service

IN THIS SECTION

Syntax

```
restart ppp-service
```

Description

Restarts the Point-to-Point Protocol (PPP) service.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart ppp-service

```
user@host> restart ppp-service
Universal edge PPP process started, pid 1687840
```

restart smg-service

IN THIS SECTION

Syntax

```
restart smg-service
```

Description

Restarts the Enhanced Session Management BNG CUPS Controller process. This command runs on the BNG CUPS Controller.

Options

`control-plane-instance` *`control-plane-instance-name`*	Display information for the designated control plane instance. This option is required when your BNG CUPS Controller system has multiple control plane instances.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart smg-service

```
user@host> restart smg-service
Control Plane Manager for dBNG started, pid <process-id>
```

restart subscriber-statistics-service

IN THIS SECTION

Syntax

```
restart subscriber-statistics-service
```

Description

Restarts the Enhanced Session Management Statistics daemon.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request. This command runs on the BNG CUPS Controller.

Sample Output

restart subscriber-statistics-service

```
user@host> restart subscriber-statistics-service
Control Plane Manager for dBNG started, pid <process-id>
```

restart upsf-service

IN THIS SECTION

Syntax

```
restart upsf-service
```

Description

Restarts the User Plane Selection Function daemon. This command runs on the BNG CUPS Controller.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart upsf-service

```
user@host> restart upsf-service
Control Plane Manager for dBNG started, pid <process-id>
```

show broadband-device

IN THIS SECTION

Syntax

```
show broadband-device [interface | detail | state | summary]
```

Description

Displays information for a broadband edge (BBE) device.

Options

`interface`	Displays information for a the specified BBE device interface.
`detail`	Displays detailed information for all BBE devices.
`state`	Displays the state for each port on the BBE device.
`summary`	Displays summary information for all BBE devices.

Required Privilege Level

root

Output Fields

Table 3 on page 195 lists the output fields for the `show broadband-device` command.

Table 3: show broadband-device Output Fields

Field Name	Field Description
Interface	BBE interface.
Name	BBE device name.
Port	The fully qualified BNG User Plane port name of the format **up:*user-plane-name*:*physical-port-name***.
Sessions	The number of sessions.
SGRP	The subscriber group that the BBE device belongs to.
User Plane	The BNG User Planes that belong to the subscriber group.
Redundancy interface name	The redundancy-interface name containing the member BNG User Plane logical ports.
Subscriber Group	The subscriber group that the BBE device belongs to.
Session ID	Session identifier.

Sample Output

show broadband-device summary

```
user@host> show broadband-device summary
Interface  Name  Port            Sessions Port               Sessions SGRP
bb0.1      alpha up:NYC:xe-2/0/0 125      up:Jersey:xe-1/0/0 124      NYC-1
bb0.2      beta  up:NYC:xe-3/0/0 120      up:Jersey:xe-2/0/0 120      NYC-2
```

show broadband-device interface

```
user@host> show broadband-device interface bb0.1
Interface Name  Port            Sessions Port                Sessions SGRP
bb0.1     alpha up:NYC:xe-2/0/0 125      up:Jersey:xe-1/0/0 124      NYC-1
```

```
user@host> show broadband-device interface bb0.1 detail
UP1 port : NYC:xe-2/0/0 (active) (124)
UP2 port : Jersey:xe-1/0/0 (backup) (123)
Redundancy Interface Name: alpha
Subscriber Group: NYC-1
Subscriber Group ID: 26
Session ID UP1 state UP 2 state
200 Installed Installed-Warm
353 Installed Out of resources
```

show firewall

IN THIS SECTION

Syntax

```
show firewall
<filter filter-name>
```

```
<filter regex regular-expression>
<terse>
```

Description

When running the `show firewall` command on the BNG CUPS Controller, the output displays only the filter names and the associated BNG User Plane. No counters or other information appears. To see the counters or the log or syslog output, you must run the `show firewall` command on the appropriate BNG User Plane.

Options

filter *filter-name*	(Optional) Name of a configured filter.
filter regex *regular-expression*	(Optional) Regular expression that matches the names of a subset of filters.
terse	(Optional) Display firewall filter names and BNG User Plane names.

Required Privilege Level

view

Output Fields

Table 4 on page 197 lists the output fields for the `show firewall` command. Output fields are listed in the approximate order in which they appear.

Table 4: show firewall Output Fields

Field Name	Field Description
`Filter`	Name of a filter that has been configured with the `filter` statement at the `[edit firewall]` hierarchy level
User Plane	BNG User Plane name

Sample Output

show firewall terse

```
user@host> show firewall terse

Filter: finid_UID4003-demux0.3221229982-in     User Plane: up-example-2
Filter: dfwda-demux0.3221229982-in             User Plane: up-example-2
Filter: finid_UID4003-demux0.3221225473-in     User Plane: up-example-1
Filter: dfwda-demux0.3221225473-in             User Plane: up-example-1
```

show health

IN THIS SECTION

Syntax

```
show health
<subsystem microservice-name (endpoint endpoint-name)>
<subsystem microservice-name (service service-name)>

<user-plane user-plane-name (endpoint endpoint-name)>
```

Description

Displays the health information about the BNG CUPS Controller subsystems or the overall health of the BNG User Planes.

Options

none	Displays health information for all BNG CUPS Controller subsystems.
`subsystem` ***`microservice-name`***	You can display health information for the following specified subsystems: • `cpi-`*`name`*—Control plane instance. • `bngd`—BNG Director. This option is only available for multiple geography deployments. • `scache`—State cache. For multiple geography deployments, there are multiple `scache` (`scache-`*`n`*). • `mgmt`—Management.
`service` ***`service-name`***	Displays health information for the associated endpoints for the specified service that is part of the microservice. Not all services have associated endpoints.
`user-plane` ***`user-plane-name`***	Displays BNG User Plane health information.
`endpoint` ***`endpoint-name`***	Displays health information for the specified endpoint of the BNG User Plane.

Required Privilege Level

root

Output Fields

Table 5 on page 200 lists the output fields for the `show health` command.

Table 5: show health

Field Name	Field Description
`Name`	Depending on which show health command output you are viewing, the name field can be one of the following: • Subsystem name • BNG User Plane name
`Cluster`	The name of the workload cluster on which a subsystem is running. This command applies only to multiple geographical redundancy deployments.
`Subsystem-State`	The state of the subsystem. Either `active` or `inactive`.
`Health`BNG CUPS Controller subsystems	Health of the BNG CUPS Controller subsystem. Following are the health levels: • Healthy—All of the following must exist: All services are up, shared memory is healthy, initial state recovery succeeded, and all key endpoints are up. • Unhealthy-major—If any of the following exist: Any of the services are permanently down, shared memory is unhealthy, or the initial state recovery failed. • Unhealthy-minor—If any of the following exist: Any of the services are down, initial state recovery is in-progress, any one of the key endpoints are down.
`Unhealty-services`	The number of unhealthy services.
`Uptime`	The amount of time the service has been up.
`Generation`	The generation number of the subsystem. A subsystem's generation number increments each time it switches the workload clusters. This command applies only to multiple geographical redundancy deployments.
`Subsystem`	The subsystem for which the information is being displayed.
`Shared-memory`	The health of the shared memory.

Table 5: show health *(Continued)*

Field Name	Field Description
`Initial State Recovery`	Displays whether the initial state recovery succeeded.
`Services`	List of services for the subsystem.
`Status`	Current status of the service. Either `up` or `down`.
`Restarts`	The number of times the service restarted.
`Endpoint-Health`	The health of the endpoint for the service.
`Unhealthy-Endpoints`	The number of unhealthy endpoints for the service.
`Key-Endpoints`	List of key endpoints.
`State (Key-Endpoints)`	State of the key endpoint.
`Flapped`	The number of times the key endpoint flapped.
`Memory Usage`	Memory usage of the service.
`CPU%`	The percentage of CPU being used by the service.

Table 5: show health *(Continued)*

<table>
<tr><th>Field Name</th><th>Field Description</th></tr>
<tr><td><code>Health</code>BNG User Plane</td><td>Health of the BNG User Plane. Following are the health levels:
Healthy—All of the following must exist: The state is connected or connecting, the corresponding smd-N service is up, and all of its associated endpoints are connected.
Unhealthy-major—If any of the following exist: The state is not connected or connecting, or security-updating and the corresponding smd-N service is down.
Unhealthy-minor—All of the following must exist: The state is either not connected, connecting, or security-updating, and the corresponding smd-N service is down, and any of its associated endpoints are disconnected.</td></tr>
<tr><td><code>Address</code></td><td>BNG User Plane IP address.</td></tr>
<tr><td><code>Active/Backup-sess</code></td><td>The number of active and backup subscriber sessions served by the BNG User Plane.</td></tr>
<tr><td><code>State (user-plane)</code></td><td>The state of the BNG User Plane. The state can be one of the following:
<code>initializing</code>
<code>ready</code>
<code>connecting</code>
<code>connected</code>
<code>disconnecting</code>
<code>disconnected</code>
<code>security-updating</code>
<code>warm-init</code>
<code>deconfiguring</code>
<code>misconfigured</code></td></tr>
</table>

Table 5: show health *(Continued)*

Field Name	Field Description
`User-plane`	The BNG User Plane for which the information is being reported.
`CPi`	The control plane instance that handling the BNG User Plane.
`Active-sessions`	The number of active subscriber sessions served by the BNG User Plane.
`Backup-sesions`	The number of backup subscriber sessions served by the BNG User Plane.

Sample Output

show health

```
user@host show health
Name            Cluster    Subsystem-State  Health    Unhealthy-services   Uptime
Generation
cpi-groton       wlc001    active         healthy    0             14d 02:22:19    1
cpi-westford     wlc001    active         healthy    0             18d 07:00:03    1
cpi-chelmsford   wlc001    active         healthy    0             18d 07:16:13    1
cpi-littleton    wlc001    active         healthy    0             18d 05:41:14    1
cpi-acton        wlc001    active         healthy    0             16d 11:31:58    1
bngd             wlc001    active         healthy    0             21d 02:24:19    1
mgmt             wlc001    active         healthy    0             21d 02:22:19    1
scache-1         wlc001    active         healthy    0             21d 02:24:01    1
scache-2         wlc002    active         healthy    0             21d 02:24:01    1
```

show health subsystem

```
user@host> show health subsystem cpi-hardening
Subsystem: cpi-hardening
Cluster: -
Health: healthy
```

```
Shared-memory: healthy
Initial State Recovery: succeeded
Services                       Status   UpTime    Restarts  Endpoint-Health  Unhealthy-Endpoints
ppp-service                     up      05:59:04     0       healthy              0
pfcp-proxy-service              up      05:59:08     0       healthy              0
smg-service                     up      05:59:08     0       healthy              0
replication-server-service      up      06:08:20     0       healthy              0
replication-client-service      up      06:08:20     0       healthy              0
authentication-service          up      05:59:09     0       healthy              0
smd-4-service                   up      05:59:10     0       healthy              0
smd-3-service                   up      05:59:11     0       healthy              0
smd-2-service                   up      05:59:11     0       healthy              0
smd-1-service                   up      05:59:12     0       healthy              0
cpm-service                     up      06:08:20     0       healthy              0
l2tp-service                    up      05:59:05     0       healthy              0
dhcp-service                    up      05:59:06     0       healthy              0
upsf-service                    up      05:59:07     0       healthy              0
subscriber-statistics-service   up      05:59:07     0       healthy              0
gtp-proxy-service               up      05:59:09     0       healthy              0

Key-Endpoints                  State         UpTime          Flapped    Sync
cpi-test-Scache-Client         reconciled    -               0          completed
cpi-test-Scache-Server         connected     01:45:05        0          continuous
Apm                            connected     01:44:32        1          -
BngDirector-cpi-test           connected     01:44:38        1          -
cpi-test-BngDirector           connected     01:44:44        1          -
BngDirector-scache             connected     01:44:38        1          -
```

show health subsystem bngd

```
user@host> show health subsystem bngd
Subsystem: bngd
Cluster: -
Health: healthy
Subsystem-State: active
Generation: 1
Shared-memory: healthy
  Services                         Status           UpTime          Restarts       Endpoint-
Health    Unhealthy-Endpoints
  bng-director-service             up               11d 04:44:29    0
healthy           0
```

```
  replication-client-service       up              11d 04:44:29    0
healthy           0
  replication-server-service       up              11d 04:44:29    0
healthy           0
  bngd-podmgr-service              up              11d 04:44:31    0
healthy           0

  Key-Endpoints                     State           UpTime          Flapped      Sync
  BngDirector-scache                connected       11d 04:44:30    0            -
  BngDirector-mgmt                  connected       11d 04:44:30    0            -
  BngDirector-cpi-xxx               connected       09:27:22        0            -
  BngDirector-cpi-yyy               connected       09:27:22        0            -
  bngd-Scache-Client                reconciled      -               0            completed
  bngd-Scache-Server                connected       09:27:22        0            continuous
```

show health subsystem scache

```
user@host> show health subsystem scache
Subsystem: scache
Cluster: -
Health: healthy
Subsystem-State: active
Generation: 1
Shared-memory: healthy

  Services                     Status UpTime          Restarts       Endpoint-Health    Unhealthy-
Endpoints
  scache-service               up     10:51:27        0              healthy            0
  replication-server-service   up     10:51:27        0              healthy            0
  replication-client-service   up     10:51:27        0              healthy            0


  Key-Endpoints                     State           UpTime          Flapped      Sync
  Scache-cpi-xxx-Client           connected       09:30:04          1          continuous
  Scache-cpi-xxx-Server           reconciled      -                 0          completed
  Scache-cpi-yyy-Client           connected       09:30:04          1          continuous
  Scache-cpi-yyy-Server           reconciled      -                 0          completed
  Scache-bngd-Client              connected       09:30:04          1          continuous
  Scache-bngd-Server              reconciled      -                 0          completed
  Scache-mgmt-Client              connected       09:30:04          1          continuous
  Scache-mgmt-Server              reconciled      -                 0          completed
```

```
  scache-BngDirector              connected        10:46:20          0
-
```

show health subsystem mgmt

```
user@host> show health subsystem mgmt
Subsystem: mgmt
Cluster: -
Health: healthy
Shared-memory: healthy

  Subsystem: mgmt
  Cluster: -
  Health: healthy
  Subsystem-State: active
  Generation: 1
  Shared-memory: healthy

  Services                      Status  UpTime       Restarts      Endpoint-Health    Unhealthy-
Endpoints
  mgmt-podmgr-service           up      10:50:49     0             healthy            0
  replication-server-service    up      10:51:27     0             healthy            0
  replication-client-service    up      10:51:27     0             healthy            0

  Key-Endpoints                     State            UpTime          Flapped      Sync
  mgmt-BngDirector                  connected        09:27:22        0            -
  mgmt-Scache-Client                reconciled       -               0            completed
  mgmt-Scache-Server                connected        09:27:22        0            continuous
```

show health subsystem <microservice-name> service

```
user@host> show health subsystem cpi-boston service gtp-proxy-service
Subsystem: cpi-boston
Cluster: -
Service: gtp-proxy-service
  Status: up
  State: ready
  Up-time: 01:45:23
  Memory Usage: 21364 kB
  CPU% (threads): 0.0% (3)
```

```
  Restarts: 0
  Endpoint-health: healthy
  Endpoint:Id                Flapped   State             Up-time
    GtpProxySmdLowIpc:1024   0         connected         01:45:10
    GtpProxySmdHighIpc:1024  0         connected         01:45:10
    GtpProxySmdMedIpc:1024   0         connected         01:45:10
```

show health subsystem <microservice-name> endpoint

```
user@host> show health subsystem cpi-boston endpoint apm
Subsystem: cpi-boston
Cluster: -
  Endpoint: Apm
  Status: connected
  Up-time: 01:45:04
```

show health user-plane

```
user@host> show health user-plane
Name     Address       CPi         State          Health           Up-time       Active/Backup-
sess
test1    192.32.6.32   cpi-boston  connected      unhealthy-minor  2d 03:10:44   31281/10400
test2    156.9.0.41    -           connecting     unhealthy-major  -             0/0
test3    178.3.65.9    cpi-boston  misconfig      healthy          16d 14:23:07  0/0
test4    77.100.1.19   -           disconnected   healthy          0d 00:00:00   0/0
test5    187.22.14.37  -           disconnecting  healthy          0d 00:00:00   0/0
```

show health user-plane

```
user@host> show health user-plane test123
User-plane: test123
Address: 192.32.6.32
Id: 1
CPi: cpi-boston
State: connected
Health: unhealthy-minor
Up-time: 2d 03:10:44
```

```
Active-sessions: 31281
Backup-sessions: 10400
Endpoints              Flapped    State            Up-time
  L2tpSmdIpc           0          connected        2d 03:10:44
  SmdL2tpIpc           0          connected        2d 03:10:44
  PppSmdIpc            0          connected        2d 03:10:44
  SmdPppIpc            0          connected        2d 03:10:44
  AuthSmdIpc           0          connected        2d 03:10:44
  SmdAuthIpc           0          connected        2d 03:10:44
  DhcpSmdIpc           0          connected        2d 03:10:44
  SmdDhcpIpc           0          connected        2d 03:10:44
  RepServerSS          0          connected        2d 03:10:44
  Cpri                 0          disconnected     -
  Sci                  0          connected        2d 03:10:44
  PfcpProxySmdIpc      0          connected        2d 03:10:44
  PfcpProxyStatsIpc    0          connected        2d 03:10:44
  SmdPfcpProxyIpc      0          connected        2d 03:10:44
  StatsPfcpProxyIpc    0          connected        2d 03:10:44
```

show health user-plane <user-plane-name> endpoint

```
user@host> show health user-plane test123 endpoint Cpri
User-plane: test123
Endpoint: Cpri
  Status: Connected
    High-priority   : Connected
    Medium-priority : Connected
    Low-priority    : Connected

  High-Priority
    Pkts client rx:             0
    Pkts terminated locally:    13242
    Pkts aggr rx:               13242
    Pkts enqueue rx fail:       0
    Client packets cp to up:    0
    Aggr packets cp to up:      26512
    Aggr packets cp to up fail: 0
    Pkts injected locally:      26512
    Last local seq num tx:      0
    Last local seq num rx:      0
    Last remote seq num rx:     13270
```

```
    Total local echo pkts rx:    0
    Total remote echo pkts rx:   13242
    Num of echo pkts lost:       13270

  Medium Priority:
    Pkts client rx:              0
    Pkts terminated locally:     13242
    Pkts aggr rx:                13242
    Pkts enqueue rx fail:        0
    Client packets cp to up:     0
    Aggr packets cp to up:       26512
    Aggr packets cp to up fail:  0
    Pkts injected locally:       26512
    Last local seq num tx:       0
    Last local seq num rx:       0
    Last remote seq num rx:      13270
    Total local echo pkts rx:    0
    Total remote echo pkts rx:   13242
    Num of echo pkts lost:       13270

  Low Priority:
    Pkts client rx:              0
    Pkts terminated locally:     13242
    Pkts aggr rx:                13242
    Pkts enqueue rx fail:        0
    Client packets cp to up:     0
    Aggr packets cp to up:       26512
    Aggr packets cp to up fail:  0
    Pkts injected locally:       26512
    Last local seq num tx:       0
    Last local seq num rx:       0
    Last remote seq num rx:      13270
    Total local echo pkts rx:    0
    Total remote echo pkts rx:   13242
    Num of echo pkts lost:       13270
```

show load-balancing-group

IN THIS SECTION

Syntax

```
show load-balancing-group
<group group-name>
```

Description

Displays information for the load-balancing group.

Options

Empty	If you do not enter a load-balancing group name, all load-balancing groups are listed.
`group group-name`	(Optional) Displays information about the listed load-balancing group.

Required Privilege Level

root

Output Fields

Table 6 on page 211 lists the output fields for the `show load-balancing-group` command.

Table 6: show load-balancing-group Output Fields

Field Name	Field Description
`Group Name`	The name of the load-balancing group.
`Logical-Port`	BNG User Plane logical port.
`% Usage`	The logical port's current load, represented as a percentage.
`CPU Exceeded`	Indicates whether the CPU load has been exceeded.
`Computed weight`	Current computed weight.
`Max weight`	Configured maximum weight.
CPi	The BNG CUPS Controller that the loading-balancing groups belong to.
Logical-port-List	List of BNG User Plane logical ports.

Sample Output

show load-balancing-group

```
user@host> show load-balancing-group
Name                CPi      Logical-Port-List
cohort              cpi-xxx  up:cheleb:xe-2/1/18 up:etrigan:xe-2/2/0
```

show load-balancing-group group *load-balancing-group-name*

```
user@host> show load-balancing-group group mygroup
Group Name     Logical-Port            % Usage  CPU Exceeded Computed weight       Max weight
mygroup        up:UP-example-1:ps0.30  80       Yes          6                     10
               up:UP-example-3:ps0.25  5        No           3                     20
```

```
                up:UP-example-2:ps0.22  30      No           2                     20
                up:UP-example-7:ps0.27  7       No           1                     20
```

show network-access address-assignment address-pool-manager status

IN THIS SECTION

Syntax

```
show network-access address-assignment address-pool-manager status
```

Description

Displays the status of Juniper Address Pool Manager (APM).

Required Privilege Level

root

Output Fields

Table 7 on page 213 lists the output fields for the `show network-access address-assignment address-pool-manager status` command. Output fields are listed in alphabetical order.

Table 7: show network-access address-assignment address-pool-manager status Output Fields

Field Name	Field Description
Address Pool Manager	IP address for APM
Status	Connection status of APM
Pool Count	Number of pools
Connect Timestamp	Time at which APM first connected to BNG CUPS Controller
Security	Connection status: secured or not secured
Appointment mode	One of the following appointment modes: • None • Remote • Local

Sample Output

show network-access address-assignment address-pool-manager status

```
user@host> show network-access address-assignment address-pool-manager status
Address Pool Manager: 10.9.160.19
Protocol: gRPC
Security: clear-text
Apportionment mode: Remote
```

show network-access address-assignment domain

IN THIS SECTION

Syntax

```
show network-access address-assignment domain
<name> domain-name
<routing-instance> routing-instance-name
```

Description

Displays the state of each pool domain (dynamic linked address pool) connected to APM and a count of the transmitted alarms for a specified routing instance.

Options

name ***domain-name / domain-profile***	(Optional) Displays information depending on which of the following variables are entered: • Empty—A summary of all domains for the routing instance that is entered. • *domain-name*—Displays the pool structure of the pool domain.
routing-instance ***routing-instance-name***	(Optional) Designate the routing instance to use. If left empty, the default routing instance is used.

Required Privilege Level

root

Output Fields

Table 8 on page 215 lists the output fields for the `show network-access address-assignment domain` command. Output fields are listed in alphabetical order.

Table 8: show network-access address-assignment domain Output Fields

Field Name	Field Description
Abatement (Abate)	The number of abatement alarms. An the alarm occurs when either of the following conditions changes, causing APM to disregard the original alarm: • The number of free addresses rises above the reclaim threshold. • The number of free addresses falls below the apportion threshold.
Active-Tag	The value of the route tag that is associated with the discard routes installed on the active BNG User Plane.
Addresses	Total number of addresses in the pool domain.
Apportion (Apport)	The number of apportion alarms. The alarm occurs when the number of free addresses falls below the apportion threshold.
Backup-Tag	The value of the route tag that is associated with the discard routes installed on the backup BNG User Plane.
Domain Name	Name of the pool domain.
Free	Number of addresses in the pool domain that are available for allocation.
Pool Count	Number of pools.
Pool Drain (Drain)	The number of pool drain alarms. The alarm occurs when a pool is completely drained.
Pool Name	Name of the pool.
Prefix	Subnetwork allocated to the address pool.
Programmed	The state of the pool state (discard routes, dhcp gateway address, and so on) programming to the BNG User Plane.

Table 8: show network-access address-assignment domain Output Fields *(Continued)*

Field Name	Field Description
Reclaim	The number of reclaim alarms. The alarm occurs when the number of free addresses for the pool domain on the BNG CUPS Controller rises above the reclaim threshold.
State	State of the pool domain.
Status	The pool is either active or in drain mode.
Type	The source of the pool prefix. It can be a local reserve partition or a remote (APM) partition.
Used	The number of addresses being used.
User-Plane	The BNG User Plane that is the target for the programmed pool state.

Sample Output

show network-access address-assignment domain

```
user@host> show network-access address-assignment domain
Domain Name          Active-Tag Backup-Tag Pool Count Addresses Free Apport Reclaim Abate Drain
v4pool-milan-default 44         55         2          510       268  2      0       0     0
v4pool-milan-foo     33         66         4          1020      137  4      0       0     0
```

show network-access address-assignment domain name (routing instance)

```
user@host> show network-access address-assignment domain name test1234 routing-instance default56
Domain Name            Pool Count Addresses  Free Apportion   Reclaim  Abatement PoolDrain
test1234-default56       1          1024      98   1            1        0         1
```

show network-access address-assignment domain name

```
user@host show network-access address-assignment domain name v4pool-default
Pool Name           Prefix       Addresses Used Type Status  User-Plane Programmed User-Plane
Programmed
v4pool-default-00001 10.19.0.0/24 254       24   Local Active  milan      added      rome
adding
v4pool-default-00002 10.19.2.0/24 254       0    Local Drained  milan    removed    rome
removing
```

show network-access address-assignment domain name

```
user@host show network-access address-assignment domain name genoa-default
Pool Name                  Prefix               Addresses  Used     Status     Mode
genoa-default              6.0.0.0/30           4          4        Active     Remote
genoa-default-0000         6.0.0.4/30           4          0        Active     Remote
genoa-default-00001        10.0.0.0/30          4          0        Active     Local
```

show network-access address-assignment domain-state

IN THIS SECTION

Syntax

```
show network-access address-assignment domain-state
<routing-instance> routing-instance-name
```

Description

Displays the alarm state (outstanding alarms) for each pool domain.

Options

routing-instance ***routing-instance-name***	(Optional) Designate the routing instance to use. If left empty, the default routing instance is used.

Required Privilege Level

root

Output Fields

Table 9 on page 218 lists the output fields for the `show network-access address-assignment domain-state` command. Output fields are listed in alphabetical order.

Table 9: show network-access address-assignment domain-state Output Fields

Field Name	Field Description
Domain Name	Name of the pool domain.

Table 9: show network-access address-assignment domain-state Output Fields *(Continued)*

Field Name	Field Description

Alarm	Name of the alarm.
	• reclaim—When the number of free addresses for the pool domain on BNG CUPS Controller rises above the reclaim threshold.
	• apportion—When the number of free addresses falls below the apportion threshold.
	• pool-drained—When a pool is completely drained.
	• abatement—BNG CUPS Controller sends an abatement alarm when either of the following conditions changes, causing APM to disregard the original alarm.
	  • The number of free addresses rises above the reclaim threshold.
	  • The number of free addresses falls below the apportion threshold.

Age	How long an alarm has been outstanding.
Logins	Whether logins are enabled.
State	State of the pool domain.

Sample Output

show network-access address-assignment domain

```
user@host> show network-access address-assignment domain
Domain Name         Pool Count Addresses  Free Apportion   Reclaim  Abatement PoolDrain
1232-default        3          507         120   1           0        0         0
test-default        2          1535        279   1           0        0         0
```

show network-access address-assignment domain name (using domain profile)

```
user@host> show network-access address-assignment domain name test1234 routing-instance default56
Domain Name           Pool Count Addresses  Free Apportion   Reclaim  Abatement PoolDrain
```

```
test1234-default56    1          1024         98  1            1         0          1
```

show network-access address-assignment domain name

```
user@host show network-access address-assignment domain name test1234-default56 routing-instance
default56
Pool Name                 Prefix           Addresses  Used  State
test1234-default56        192.0.2.1/24      255        253   Active
test1234-default56-000    192.0.2.8/24      254        0     Active
-
```

show routing-instances

IN THIS SECTION

Syntax

```
show routing instances <routing-instance-name>
```

Description

Displays a list of BNG User Planes that are using the listed routing instance.

Options

routing-instance-name	The routing instance name for which you want the list of BNG User Planes.

Required Privilege Level

view

Output Fields

Table 26 on page 298 lists the output fields for the `show routing-instances` command. Output fields are listed in the approximate order in which they appear.

Table 10: show routing-instances Output Fields

Field Name	Field Description
`User Plane Name`	Name of the BNG User Plane.
`Routing Instance State`	The routing instance state: • Connected—The node is connected to the network. • Isolated—The node is isolated from the rest of the network.

Sample Output

show routing-instances

```
user@host> show routing-instances example-1
Routing Instance: example-1
User Plane Name         Routing Instance State
test-2                  isolated
test-3                  isolated
example-1               connected
example-2               connected
```

show subscriber-group

IN THIS SECTION

Syntax

```
show subscriber-group subscriber-group-name
```

Description

Displays information for subscriber groups.

Options

`subscriber-group` *`subscriber-group-name`*	Subscriber group for which you want to display information. If you do not enter a *subscriber-group-name*, the command only displays information for subscriber groups that have logical-ports associated with the subscriber group. If you want to see information for a subscriber group that does not have any logical ports associated with it, you must include the *subscriber-group-name* in the command.

Required Privilege Level

root

Output Fields

Table 11 on page 223 lists the output fields for the `show subscriber-group` command.

Table 11: show subscriber-group Output Fields

Field Name	Field Description
Name	Subscriber group name.
ID	ID number for the subscriber group.
SGRP Mode	The operational mode of the device, either Control Plane or User Plane
SGRP State	Health status of the subscriber group. Either healthy or unhealthy.
User Plane	The BNG User Planes that belong to the subscriber group.
Active UP	The active BNG User Plane.
Mode	The operational mode of the device, either Control Plane or User Plane.
BB device	The Broadband device that is a member of the subscriber group.
Name (Logical port mapping)	Logical port mapping name.
Logical-port	BNG User Plane logical port.
Sessions	The number of subscriber sessions.
Name (Address domains)	Address Domain name.
Prefixes	The number of address prefixes assigned to the specified BNG User Plane for the subscriber group.

Table 11: show subscriber-group Output Fields *(Continued)*

Field Name	Field Description
User-Plane	The BNG User Plane that the specified subscriber group belongs to. The User-Plane field lists the following information: • active or backup—Indicates whether the BNG User Plane is configured as active or backup. • hot, warm-enhanced, or warm-scale—The backup mode for the subscriber group. • admin-disabled—Indicates if the BNG User Plane port for the subscriber group is disabled.
Programmed	The number of address prefixes programmed on the User Plane for the subscriber group
Name (Routing Instances)	Routing instances name.
User-Plane (Address domain)	BNG User Plane name.
CPi	The BNG CUPS Controller that the subscriber groups belong to.

Sample Output

show subscriber-group

```
user@host> show subscriber-group
Name          CPi      ID    SGRP Mode      SGRP State     User Plane       User Plane
Active UP
SGRP1_CP      cpi-xxx  2048  Control Plane  healthy        r14mx960wf       r7mx304wf
r14mx960wf
r7mx304wf     cpi-yyy  1025  Control Plane  healthy        r7mx304wf        ---
r7mx304wf
```

show subscriber-group *subscriber-group-name*

```
user@host> show subscriber-group SGRP1
Name: SGRP1
ID: 2048
User-Plane: jersey (active) (hot) (admin-disabled)
User-Plane: boston (backup) (hot)
Health status: healthy
Mode: User Plane
VMAC: 0A:0A:0A:0A:0A:0A
Logical port mapping:
  BB device   Name             Logical-port   Sessions   Logical-port         Sessions
  bb0.8       beta             up:ps0        0          up:boston:ps0        0
```

show subscriber-group *subscriber-group-name*

```
user@host> show subscriber-group example-1
Name: example-1
ID: 1
User-Plane: caelum (active) , (hot)
Health status: healthy
Mode: Control Plane
Logical port mapping:
  BB device   Name             Logical-port         Sessions   Logical-port         Sessions
  bb0.12      bb0.12           up:caelum:xe-2/1/0   13000      ---                  ---
  bb0.9       bb0.9            up:caelum:xe-2/0/0   12999      ---                  ---
  bb0.8       bb0.8            up:caelum:xe-1/1/0   16000      ---                  ---
  bb0.7       bb0.7            up:caelum:xe-1/0/0   16000      ---                  ---
  bb0.6       bb0.6            up:caelum:ge-2/3/0   12999      ---                  ---
  bb0.5       bb0.5            up:caelum:ge-2/2/0   13000      ---                  ---
Address domains:
  Name                               Prefixes    User-Plane       Programmed    User-Plane
Programmed
  v4pool:caelum:default              352         caelum           352           ---
---
```

show subscribers

IN THIS SECTION

Syntax

```
show subscribers
<display (detail | extensive | terse)>
<accounting-statistics>
<aci-interface-set-name address>
<address address>
<agent-circuit-identifier agent-circuit-identifier>
<agent-remote-identifier agent-remote-identifier>
<id> session-id
<mac-address mac-address>
<user-name user-name>
```

Description

Display information for active subscribers.

Options

`display (detail | extensive | terse)`	(Optional) Display the specified level of output.
`accounting-statistics`	(Optional) Display subscriber accounting statistics

aci-interface-set-name	(Optional) Display all the dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. You must use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) control packets.
address	(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.
agent-circuit-identifier	(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

agent-remote-identifier	(Optional) Display all dynamic subscriber sessions whose ARI value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.
`id` *session-id*	(Optional) Display a specific subscriber session whose session ID matches the specified subscriber ID. You can display subscriber IDs by using the `show subscribers extensive` command.
mac-address	(Optional) Display subscribers whose MAC address matches the specified MAC address.
user-name	(Optional) Display subscribers whose username matches the specified subscriber name.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 12 on page 228 lists the output fields for the `show subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 12: show subscribers Output Fields

Field Name	Field Description
`Interface`	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is `Tunnel-switched`.
`User Name`	Name of subscriber.
`LS:RI`	Logical system and routing instance associated with the subscriber.
`Type`	Subscriber client type (DHCP, FWA, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
`IP Address`	Subscriber IPv4 address.
`IP Netmask`	Subscriber IP netmask. This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Primary DNS Address`	IP address of primary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`Secondary DNS Address`	IP address of secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Primary DNS Address`	IPv6 address of primary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`IPv6 Secondary DNS Address`	IPv6 address of secondary DNS server. This field is displayed with the `extensive` option only when the address is provided by RADIUS.
`Domain name server inet`	IP addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Domain name server inet6`	IPv6 addresses for the DNS server, displayed in order of configuration. This field is displayed with the `extensive` option only when the addresses are derived from the access profile or the global access configuration.
`Primary WINS Address`	IP address of primary WINS server.
`Secondary WINS Address`	IP address of secondary WINS server.
`IPv6 Address`	Subscriber IPv6 address, or multiple addresses.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv6 Prefix`	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
`IPv6 User Prefix`	IPv6 prefix obtained through NDRA.
`IPv6 Address Pool`	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
`IPv6 Network Prefix Length`	Length of the network portion of the IPv6 address.
`IPv6 Prefix Length`	Length of the subscriber IPv6 prefix.
`Logical System`	Logical system associated with the subscriber.
`Routing Instance`	Routing instance associated with the subscriber.
`Interface`	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form `demux0.`*nnnn* (for example, `demux0.3221225472`), to which access-internal and framed subscriber routes are mapped.
`Interface Type`	Whether the subscriber interface is `Static` or `Dynamic`.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Interface Set`	Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line. The name of the interface set uses one of the following prefixes: • `aci`—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • `ari`—ARI; for example, ari-1033-demux0.3221225524. • `aci+ari`—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524. • `noids`—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. **NOTE**: ACI interface sets are configured with the `agent-circuit-identifier` autoconfiguration stanza. ALI interface sets are configured with the `line-identity` autoconfiguration stanza. Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable $junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.
`Interface Set Type`	Interface type of the ACI interface set: `Dynamic`. This is the only ACI interface set type currently supported.
`Interface Set Session ID`	Identifier of the dynamic ACI interface set entry in the session database.
`Underlying Interface`	Name of the underlying interface for the subscriber session.
`Dynamic Profile Name`	Dynamic profile used for the subscriber.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init, Configured, Active, Terminating, Tunneled).
L2TP State	Current state of the L2TP session, Tunneled or Tunnel-switched. When the value is Tunnel-switched, two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
PFE Flow ID	Forwarding flow identifier.
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Agent Circuit ID`	For the `dhcp` client type, option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the `vlan-oob` client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.
`Agent Remote ID`	For the `dhcp` client type, option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format. For the `vlan-oob` client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.
`Aggregation Interface-set Name`	Value of the $junos-aggregation-interface-set-name predefined variable; one of the following: • When the `hierarchical-access-network-detection` option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the ANCP Port Up message or PPPoE PADR IA tags begins with a `#` character, then the variable takes the value of the remainder of the string after the `#` character. • When the `hierarchical-access-network-detection` option is not configured, or if the sting does not begin with the `#` character, then the variable takes the value specified with the `predefined-variable-defaults` statement.
`Accounting Statistics`	Actual transmitted subscriber accounting statistics by session ID or interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the `actual-transmit-statistics` statement is included in the dynamic profile.
`DHCP Relay IP Address`	IP address used by the DHCP relay agent.
`Login Time`	Date and time at which the subscriber logged in.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 Input Service Set	Input service set in access dynamic profile.
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv4 rpf-check Fail Filter Name`	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
`IPv6 rpf-check Fail Filter Name`	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
`DHCP Options`	`len` = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
`Session ID`	ID number for a subscriber session.
`Underlying Session ID`	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
`Service Sessions`	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
`Service Session ID`	ID number for a subscriber service session.
`Service Session Name`	Service session profile name.
`Session Timeout (seconds)`	Number of seconds of access provided to the subscriber before the session is automatically terminated.
`Idle Timeout (seconds)`	Number of seconds subscriber can be idle before the session is automatically terminated.
`IPv6 Delegated Address Pool`	Name of the pool used for DHCPv6 prefix delegation.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute. This field is displayed only when the predefined variable $junos-ipv6-address is used in the dynamic profile.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv6 Output Filter Name`	Name assigned to the IPv6 output filter (client or service session).
`IFL Input Filter Name`	Name assigned to the logical interface input filter (client or service session).
`IFL Output Filter Name`	Name assigned to the logical interface output filter (client or service session).
`DSL type`	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: `ADSL`, `ADSL2`, `ADSL2+`, `OTHER`, `SDSL`, `VDSL`, or `VDSL2`.
`Frame/Cell Mode`	Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091): • `Cell`—When the DSL line type is one of the following: ADSL, ADSL2, or ADSL2+. • `Frame`—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. The value is stored in the subscriber session database.
`Overhead accounting bytes`	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
`Actual upstream data rate`	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).
`Actual downstream data rate`	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Adjusted downstream data rate`	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
`Adjusted upstream data rate`	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
`Local TEID-U`	Tunnel endpoint identifier on the BNG for the GTP-U user plane tunnel to the eNodeB. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPU Tunnel Local IP address` value.
`Local TEID-C`	Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the MME. The identifier is allocated by the BNG. A fully qualified local TEID-C consists of this identifier and the `GTPC Local IP address` value.
`Remote TEID-U`	Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB. A fully qualified remote TEID-U consists of this identifier and the `GTPU Tunnel Remote IP address` value.
`Remote TEID-C`	Tunnel endpoint identifier on the MME for the GTP-C control plane tunnel to the BNG. The identifier is allocated by the MME. A fully qualified remote TEID-C consists of this identifier and the `GTPC Remote IP address` value.
`GTPU Tunnel Remote IP address`	IP address of the S1-U interface on the eNodeB for the GTP-U tunnel endpoint. A fully qualified remote TEID-U consists of this address and the `Remote TEID-U` value.
`GTPU Tunnel Local IP address`	IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint. A fully qualified local TEID-U consists of this address and the `Local TEID-U` value

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`GTPC Remote IP address`	IP address of the S11 interface on the MME for the GTP-C tunnel endpoint. A fully qualified remote TEID-C consists of this address and the `Remote TEID-C` value.
`GTPC Local IP address`	IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint. A fully qualified local TEID-C consists of this address and the `Local TEID-C` value.
`Access Point Name`	Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.
`Tenant`	Name of the tenant system. You can create multiple tenant system administrators for a tenant system with different permission levels based on your requirements.
`User Plane id`	ID number for the BNG User Plane that the subscriber belongs to.
`User Plane Name`	Name of the BNG User Plane that the subscriber belongs to.
`User-plane:port`	The BNG User Plane that the subscriber belongs to with its port number and whether it is configured as active or backup.
`Routing instance`	Name of the routing instance. When a custom routing instance is created for a tenant system, all the interfaces defined in that tenant system are added to that routing instance.
`Dynamic Profile Version Alias`	Configured name for a specific variation of a base dynamic profile. IT's presence indicates that the profile configuration is different from that of the base profile. The value is conveyed to the RADIUS server during authentication in the Client-Profile-Name VSA (26–4874–174).
`CP-instance`	BNG CUPS Controller instance.
`SGRP`	The subscriber group that the subscriber belongs to.

Table 12: show subscribers Output Fields *(Continued)*

Field Name	Field Description
`Active-UP`	Lists the active BNG User Plane for the SGRP.

Sample Output

show subscribers user-name

```
user@host> show subscribers user-name user@host.com
SID       CP-instance    SGRP             Ports
234096    cpi-boston     alk-vest003      alkaid:xe-0/0/1,cora:xe-2/1/1
230077    cpi-boston     alk-west001      alkaid:xe-2/1/1,cora:xe-0/2/0
```

show subscribers user-name

```
user@host> show subscribers user-name user@host.com display detail
<device header for cpi-boston>
Type: DHCP
User Name: joe@sixpack.com
IP Address: 192.168.0.1
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225553
Interface-tag: foobar-tag
Interface type: Dynamic
Underlying Interface: demux0.3221225547
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:01:02:03:04:28
Idle Timeout (seconds): 1800
Idle Timeout Ingress Only: FALSE
State: Active
Radius Accounting ID: 4106
Session ID: 234096
SGRP: alk-vest003
```

```
User-plane:port: alkaid:xe-2/1/1 (active)
User-plane:port: cora:xe-2/0/1 (backup)
PFE Flow ID: 132
Stacked VLAN Id: 210
VLAN Id: 214
Login Time: 2023-04-24 07:44:46 PDT
DHCP Options: len 3
35 01 01
DHCP Header: len 44
01 01 06 00 84 76 db 36 00 00 80 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 01 02 03 04 28 00 00 00 00 00 00
00 00 00 00
```

show subscribers detail (DHCP)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 16.0.0.2
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
User Plane ID: 1
Interface: up:green-arrow:demux0.3221225474
Interface type: Dynamic
Underlying Interface: up:green-arrow:ge-0/3/5.2
Dynamic Profile Name: client-dhcp-demux
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 8
Session ID: 8
PFE Flow ID: 12
VLAN Id: 10
Login Time: 2022-02-23 22:35:35 UTC
DHCP Options: len 3
35 01 01
DHCP Header: len 44
01 01 06 00 dd 7d 5a 46 00 00 80 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 64 03 01 02 00 00 00 00 00 00
00 00 00 00

Type: DHCP
```

```
IPv6 Address: 1000::3
Logical System: default
Routing Instance: default
User Plane ID: 1
Interface: up:green-arrow:demux0.3221225475
Interface type: Dynamic
Underlying Interface: up:green-arrow:ge-0/3/5.2
Dynamic Profile Name: client-dhcp-demux
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 9
Session ID: 9
PFE Flow ID: 13
VLAN Id: 10
Login Time: 2022-02-23 22:35:44 UTC
DHCPV6 Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 03 00 03 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
DHCPV6 Header: len 4
01 03 00 00
```

show subscribers accounting-statistics id

```
user@host> show subscribers accounting-statistics id 206
Session ID: 206
Interface: pp0.3221225677
Accounting Statistics
   Input bytes: 0
   Input packets: 0
   Output bytes: 0
   Output packets: 0
IPv6
   Input bytes: 0
   Input packets: 0
   Output bytes: 0
   Output packets: 0
```

show subscribers client-type (PPPoE)

```
user@host> show subscribers client-type pppoe
Interface                       IP Address/VLAN ID        User Name         LS:RI
up:green-arrow:pp0.3221225473   100.16.0.2                user-example-1    default:default
*                               1000::2
```

show subscribers client-type (DHCP)

```
user@host> show subscribers client-type pppoe
Interface                          IP Address/VLAN ID      User Name        LS:RI
up:green-arrow:demux0.3221225474   16.0.0.2                                 default:default
up:green-arrow:demux0.3221225475   1000::3                                  default:default
```

show subscribers detail

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
User Plane ID: 2
SGRP ID: 1
SGRP Refcnt: 1
Interface: demux0.3221225472
Interface type: Dynamic
Underlying Interface: bb0
Dynamic Profile Name: ppp-dvlan
State: Active
Session ID: 469
PFE Flow ID: 419
Stacked VLAN Id: 0x8100.3500
VLAN Id: 0x8100.3500
Login Time: 2023-02-23 16:40:27 UTC

Type: VLAN
Logical System: default
Routing Instance: default
User Plane ID: 2
```

```
SGRP ID: 1
SGRP Refcnt: 1
Interface: ge-0/0/0
Interface Set: aci-1002-demux0.3221225472
Interface Set Session ID: 0
Underlying Interface: demux0.3221225472
Dynamic Profile Name: ACI-SET-NGN2
State: Active
Session ID: 470
Agent Circuit ID: ACI-Household-1
Login Time: 2023-02-23 16:40:27 UTC

Type: PPPoE
User Name: DEFAULTUSER
IP Address: 192.0.101.31
IP Netmask: 255.255.255.255
IPv6 User Prefix: 3000:0:0:119::/64
Logical System: default
Routing Instance: default
User Plane ID: 2
SGRP ID: 1
Interface: pp0.3221225473
Interface type: Dynamic
Interface Set: aci-1002-demux0.3221225472
Interface Set Session ID: 470
Underlying Interface: demux0.3221225472
Dynamic Profile Name: SOHO-NGN2-FTTH
MAC Address: 00:03:01:00:00:01
State: Active
Radius Accounting ID: 471
Session ID: 471
PFE Flow ID: 419
Stacked VLAN Id: 3500
VLAN Id: 3500
Agent Circuit ID: ACI-Household-1
Login Time: 2023-02-23 16:40:27 UTC
```

show subscribers agent-circuit-identifier *substring* detail

```
user@host> show subscribers agent-circuit-identifier ACI-Household-1 detail
Type: VLAN
```

```
Logical System: default
Routing Instance: default
User Plane ID: 2
SGRP ID: 1
SGRP Refcnt: 1
Interface: ge-0/0/0
Interface Set: aci-1002-demux0.3221225472
Interface Set Session ID: 0
Underlying Interface: demux0.3221225472
Dynamic Profile Name: ACI-SET-NGN2
State: Active
Session ID: 470
Agent Circuit ID: ACI-Household-1
Login Time: 2023-02-23 16:40:27 UTC

Type: PPPoE
User Name: DEFAULTUSER
IP Address: 192.0.101.31
IP Netmask: 255.255.255.255
IPv6 User Prefix: 3000:0:0:119::/64
Logical System: default
Routing Instance: default
User Plane ID: 2
SGRP ID: 1
Interface: pp0.3221225473
Interface type: Dynamic
Interface Set: aci-1002-demux0.3221225472
Interface Set Session ID: 470
Underlying Interface: demux0.3221225472
Dynamic Profile Name: SOHO-NGN2-FTTH
MAC Address: 00:03:01:00:00:01
State: Active
Radius Accounting ID: 471
Session ID: 471
PFE Flow ID: 419
Stacked VLAN Id: 3500
VLAN Id: 3500
Agent Circuit ID: ACI-Household-1
Login Time: 2023-02-23 16:40:27 UTC
```

show subscribers subscriber-group

IN THIS SECTION

Syntax

```
show subscribers subscriber-group sgrp-name subscriber-goup-name
<client-type client-type>
<broadband-device broadband-device-name>
<interface interface-name
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<user-name user-name>
<vlan-id vlan-id>
```

Description

Display information for subscribers as part of a subscriber group.

Options

broadband-device-name	Name of the broadband edge device.
client-type	(Optional) Display subscribers whose client type matches one of the following client types:

- `dhcp`—DHCP clients only.
- `dotlx`—Dotlx clients only.
- `essm`—ESSM clients only.
- `fixed-wireless-access`—Fixed wireless access clients only.
- `fwauth`—FwAuth (authenticated across a firewall) clients only.
- `l2tp`—L2TP clients only.
- `mlppp`—MLPPP clients only.
- `ppp`—PPP clients only.
- `pppoe`—PPPoE clients only.
- `static`—Static clients only.
- `vlan`—VLAN clients only.
- `vlan-oob`—VLAN out-of-band (ANCP-triggered) clients only.
- `vpls-pw`—VPLS pseudowire clients only.
- `xauth`—Xauth clients only.

interface (Optional) Display subscribers whose interface matches the specified interface.

profile-name (Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance (Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id (Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state (Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name (Optional) Display subscribers whose username matches the specified subscriber name.

vlan-id (Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN.

For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the `stacked-vlan-id` *`stacked-vlan-id`* option to match the outer VLAN tag.

NOTE: Because of display limitations, routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 13 on page 248 lists the output fields for the `show subscribers subscriber-group` command. Output fields are listed in the approximate order in which they appear.

Table 13: show subscribers subscriber-group Output Fields

Field Name	Field Description
`Control-Plane-instance`	The associated BNG CUPS Controller.
`Broadband-device`	The list of broadband edge devices.
`IP Address`	Subscriber IPv4 address.
`User Name`	Name of subscriber.
`RI`	Routing instance associated with the subscriber.
`SID`	ID number for a subscriber session.
`Resiliency-state`	Lists whether the device is configured as active or backup.

Sample Output

show subscribers subscriber-group sgrp-name

```
user@host> show subscribers subscriber-group sgrp-name SGRP2_UP
Broadband-device    IP Address              User Name              RI
SID             Resiliency-state
bb0.12                                                             default
6183            -
bb0.12              5.0.0.254               user@juniper.com       RI_2
6184            -
```

show subscribers summary

IN THIS SECTION

Syntax

```
show subscribers summary
<user-plane user-plane-name>
<subscriber-group subscriber-group-name>
<control-plane-instance control-plane-instance-name>
<port>
```

Description

Display summary information for subscribers.

Options

`user-plane` *`user-plane-name`*	Display subscriber information for the designated BNG User Plane.
`subscriber-group` *`subscriber-group-name`*	Display subscriber information for the designated subscriber group.
`control-plane-instance` *`control-plane-instance-name`*	Display subscriber information for the designated control plane instance.

`port` Display the number of active subscriber sessions for the BNG User Plane ports.

Required Privilege Level

view

Output Fields

Table 14 on page 250 lists the output fields for the `show subscribers summary` command. Output fields are listed in the approximate order in which they appear.

Table 14: show subscribers summary Output Fields

Field Name	Field Description
user-plane	BNG User Plane name.
control-plane-instance	BNG CUPS Controller name.
subscriber-group	Subscriber group name.

Table 14: show subscribers summary Output Fields *(Continued)*

Field Name	Field Description
`Subscribers by State`	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.
`Subscribers by Client Type`	Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, and VLAN. Also displays the total number of subscribers for all client types (Total).

Sample Output

show subscribers summary

```
user@host> show subscribers summary user-plane up-example-1
Subscribers by State
   Active: 1
   Total: 1

Subscribers by Client Type
   DNCP: 1
   Total: 1
```

show subscribers summary control-plane-instance

```
user@host> show subscribers summary control-plane-instance
Subscriber by State
```

```
control-plane-instance     init   configured    active      terminating terminated total
cpi-boston                 200    100           38665       610         425        40000
cpi-new-york               0      0             20000       0           0          20000



Subscribers by Client Type
control-plane-instance     vlan   dhcp    pppoe      l2tp  total
cpi-boston                 10     39990   0          0     40000
cpi-new-york               10     19990   0          0     20000
```

show subscribers summary control-plane-instance *control-plane-instance-name*

```
user@host> show subscribers summary control-plane-instance cpi-boston
Subscribers by State
   Init: 200
   Configured: 100
   Active: 38665
   Terminating: 610
   Terminated: 425
   Total: 40000

Subscribers by Client Type
   DHCP: 39990
   VLAN: 10
   Total: 40000
```

show subscribers summary user-plane

```
user@host> show subscribers summary user-plane
Subscriber by State
user-plane    control-plane-instance     init   configured    active      terminating
terminated  total
R0mx30401     cpi-boston                  0      0             36235       610
425        37270
R0mx3048a     cpi-boston                  200    100           2430        0
0           2730
R1mx1003a0    cpi-new-york                0      0             20000       0
0          20000

Subscribers by Client Type
```

```
user-plane    control-plane-instance     vlan   dhcp    pppoe      l2tp      total
R0mx30401     cpi-boston                 10     37260   0          0         37270
R0mx3048a     cpi-boston                 20     2410    0          0         2730
R1mx1003a0    cpi-new-york               1000   19000   0          0         20000
```

show subscribers summary user-plane *user-plane-name*

```
user@host> show subscribers summary user-plane up-example1
Subscribers by State
   Active: 36235
   Terminating: 610
   Terminated: 425
   Total: 37270

Subscribers by Client Type
   DHCP: 37260
   VLAN: 10
   Total: 37270
Subscribers by Resiliency State
   Active: 18630
   Backup: 18630
   Total: 37270
```

show subscribers summary subscriber-group

```
user@host> show subscribers summary subscriber-group
Subscriber by State
subscriber-group          init   configured    active      terminating terminated  total
mx-default                0      0             36235       610         425        37270
mx-def-nord               0      100           2430        0           0           2730
mx-sud                    0      0             20000       0           0          20000

Subscribers by Client Type
subscriber-group         vlan   dhcp    pppoe      l2tp     total
mx-default               10     37260   0          0        37270
mx-def-nord              20     2410    0          0        2730
mx-sud                   1000   19000   0          0        20000
```

show subscribers summary subscriber-group *subscriber-group-name*

```
user@host> show subscribers summary subscriber-group example-subscriber-group
Subscribers by State
   Active: 36235
   Terminating: 610
   Terminated: 425
   Total: 37270

Subscribers by Client Type
   DHCP: 37260
   VLAN: 10
   Total: 37270
```

show subscribers summary port

```
user@host> show subscribers summary subscriber-group example-subscriber-group
Interface
up:acubens-up:et-0/0/0:0            145096
up:acubens-up:et-0/1/1:0            2291
up:cassiopeia-up:et-0/0/0:1         7822
```

show subscribers user-plane

IN THIS SECTION

Syntax

```
show subscribers user-plane up-name user-plane-name
<client-type client-type>
<interface interface
<physical-interface physical-interface-name>
<profile-name profile-name>
<routing-instance routing-instance>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<user-name user-name>
<vlan-id vlan-id>
```

Description

Displays information for subscribers associated to a BNG User Plane.

Options

client-type (Optional) Display subscribers whose client type matches one of the following client types:

- `dhcp`—DHCP clients only.
- `dotlx`—Dotlx clients only.
- `essm`—ESSM clients only.
- `fixed-wireless-access`—Fixed wireless access clients only.
- `fwauth`—FwAuth (authenticated across a firewall) clients only.
- `l2tp`—L2TP clients only.
- `mlppp`—MLPPP clients only.
- `ppp`—PPP clients only.
- `pppoe`—PPPoE clients only.
- `static`—Static clients only.

- `vlan`—VLAN clients only.
- `vlan-oob`—VLAN out-of-band (ANCP-triggered) clients only.
- `vpls-pw`—VPLS pseudowire clients only.
- `xauth`—Xauth clients only.

interface	(Optional) Display subscribers whose interface matches the specified interface.
physical-interface-name	(Optional) Display subscribers whose physical interface matches the specified physical interface.
profile-name	(Optional) Display subscribers whose dynamic profile matches the specified profile name.
routing-instance	(Optional) Display subscribers whose routing instance matches the specified routing instance.
stacked-vlan-id	(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.
subscriber-state	(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).
user-name	(Optional) Display subscribers whose username matches the specified subscriber name.
vlan-id	(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the `stacked-vlan-id` *`stacked-vlan-id`* option to match the outer VLAN tag.

NOTE: Because of display limitations, routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 15 on page 257 lists the output fields for the `show subscribers user-plane` command. Output fields are listed in the approximate order in which they appear.

Table 15: show subscribers user-plane Output Fields

Field Name	Field Description
`Control-Plane-instance`	The associated BNG CUPS Controller.
`Broadband-device`	The list of broadband edge devices.
`IP Address`	Subscriber IPv4 address.
`User Name`	Name of subscriber.
`RI`	Routing instance associated with the subscriber.
`SID`	ID number for a subscriber session.
`Resiliency-state`	Lists whether the device is configured as active or backup.

Sample Output

show subscribers user-plane up-name

```
user@host> show subscribers user-plane up-name test123 client-type dhcp physical-interface ae1
Control-plane-instance: cpi-test1
Broadband-device    IP Address     User Name      RI      SID    Resiliency-state
bb0.1               192.168.0.0    user@host.com  default  1340   Active
bb0.1               192.168.0.5    user@host.com  default  1897   Active
bb0.2               192.168.0.6    user@host.com  default  2349   Backup
```

show user-plane igmp interface

IN THIS SECTION

Syntax

```
show user-plane igmp interface up-name user-plane-name
```

Description

Displays information about Internet Group Management Protocol (IGMP)-enabled interfaces on BNG User Planes.

Options

none	When you run this command on the BNG CUPS Controller, the output displays standard information about all IGMP-enabled interfaces on all BNG User Planes associated to the BNG CUPS Controller. When you run this command on a BNG User Plane, the output displays standard information about all IGMP-enabled interfaces on the BNG User Plane.
`up-name` *`user-plane-name`*	(Optional) Displays information about the IGMP-enabled interfaces on the specified BNG User Plane.

Required Privilege Level

view

Output Fields

Table 16 on page 259 describes the output fields for the `show user-plane igmp interface` command. Output fields are listed in the approximate order in which they appear.

Table 16: show user-plane igmp interface Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Querier`	Address of the routing device that has been elected to send membership queries.
`State`	State of the interface: **Up** or **Down**.
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.
`Version`	IGMP version being used on the interface: **1** , **2** , or **3**.
`Groups`	Number of groups on the interface.
`Group threshold`	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.
`Group log-interval`	Time (in seconds) between consecutive log messages.
`SSM map policy`	The SSM map policy applied to the IGMP interface..

Table 16: show user-plane igmp interface Output Fields *(Continued)*

Field Name	Field Description
`Immediate Leave`	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.
`Promiscuous Mode`	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.
`Distributed`	State of IGMP, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed IGMP is enabled.
`Passive`	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The `passive` statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic.

Table 16: show user-plane igmp interface Output Fields *(Continued)*

Field Name	Field Description
Group policy	The group policy applied to the IGMP interface.

Sample Output

show user-plane igmp interface up-name

```
user@host> show user-plane igmp interface up-name up-example-1
Interface: up:up-example-1:pp0.3221225481
    Querier: 0.0.0.0
    State:         Up Timeout:    None Version:   3
    Group threshold:   0
    SSM map policy: igmp-ssm-map-policy
    Immediate leave: On
    Promiscuous mode: On
    Passive: Off
    Group policy: igmp-group-policy
    Distributed: On
```

Release Information

Statement introduced in Juniper BNG CUPS Release.

show user-plane ipv6 router-advertisement

IN THIS SECTION

Syntax

```
show user-plane ipv6 router-advertisement
<interface interface>
<prefix prefix/prefix length>
<up-name user-plane-name>
```

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

The router advertisement module does not function in the backup Routing Engine as the Routing Engine does not send an acknowledgment message after receiving the packets.

Options

interface *interface*	(Optional) Display IPv6 router advertisement information for the specified interface.
prefix *prefix/prefix length*	(Optional) Display IPv6 router advertisement information for the specified prefix.
up-name *user-plane-name*	The BNG User Plane for which you want to view IPv6 router advertisement information.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

Output Fields

Table 17 on page 263 describes the output fields for the `show user-plane ipv6 router-advertisement` command. Output fields are listed in the approximate order in which they appear.

Table 17: show user-plane ipv6 router-advertisement Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Advertisements sent`	Number of router advertisements sent and the elapsed time since they were sent.
`Solicits received`	Number of solicitation messages received.
`Advertisements received`	Number of router advertisements received.
`Advertisements from`	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
`Managed`	Managed address configuration flag: `0` (stateless) or `1` (stateful).
`Other configuration`	Other stateful configuration flag: `0` (stateless) or `1` (stateful).
`Reachable time`	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
`Default lifetime`	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
`Retransmit timer`	Time between retransmitted Neighbor Solicitation messages, in milliseconds.

Table 17: show user-plane ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
`Current hop limit`	Configured current hop limit.
`Prefix`	Name and length of the prefix.
`Valid lifetime`	How long the prefix remains valid for onlink determination.
`Preferred lifetime`	How long the prefix generated by stateless autoconfiguration remains preferred.
`On link`	Onlink flag: `0` (not onlink) or `1` (onlink).
`Autonomous`	Autonomous address configuration flag: `0` (not autonomous) or `1` (autonomous).
`Upstream Mode`	Configured interface as upstream interface for RA proxy
`Downstream Mode`	Configured interface as downstream interface for RA proxy.
`Downstream`	Downstream interface for RA proxy.
`Passive Mode`	RA receive only mode is enabled.
`Proxy Blackout Timer`	Proxy blackout timer interval is the time interval for which the interface must not be used as a proxy interface. Proxy functionality is disabled on that interface.
`Parameter Preference`	Preference to select configured or proxied parameters for downstream interface
`error`	Displays the details of the error.

Sample Output

show user-plane ipv6 router-advertisement up-name

```
user@host> show user-plane ipv6 router-advertisement up-name up-test1
   Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

Release Information

show user-plane maintenance

IN THIS SECTION

Syntax

```
show user-plane maintenance up-name user-plane-name
```

Description

Displays the maintenance state for a BNG User Plane.

Options

user-plane-name Display the maintenance status of the specified BNG User Plane.

Required Privilege Level

view

Output Fields

Table 18 on page 266 lists the output fields for the `show user-plane maintenance up-name` command.

Table 18: show user-plane maintenance up-name Output Fields

Field Name	Field Description
`Maintenance Status`	The current maintenance status of the BNG User Plane.
`Serviced UP`	The BNG User Plane that is being serviced.
`Backup UP`	The backup BNG User Plane.
`BB device`	Broadband device for the logical port pair for the BNG User Plane being serviced and the backup BNG User Plane.
`Name (Logical port mapping)`	Redundancy interface name for the logical port pair for the BNG User Plane being serviced and the backup BNG User Plane.

Table 18: show user-plane maintenance up-name Output Fields *(Continued)*

Field Name	Field Description
`Logical-port`	The number of subscriber sessions configured on the BNG User Plane logical port.
`Sesions`	The number of subscriber sessions configured on the BNG User Plane logical port.
`Name (address Domains)`	Name of the address Domain.
`Prefixes`	The number of address prefixes assigned to the address domain.
`User-Plane`	The BNG User Plane name.
`programmed`	One of the following: • Address Domain—The number of address prefixes configured on the BNG User Plane. • Routing Instance—The programming state of the prefixes with the route tag on the BNG User Plane.
`Name (Routing Instance)`	Name of the routing instance.
`Tag`	The routing tag value used for the active BNG User Plane.
`Backup-Tag`	The routing tag value used for the backup BNG User Plane.

Sample Output

show user-plane maintenance up-name

```
user@host> show user-plane maintenance up-name up-exampl-1
Maintenance Status:- In-progress
Serviced UP: up-exampl-1 - Active/Synchronized
Backup UP: up-examp2-2- Backup/Sync-in-progress
Logical port mapping:
```

```
BB device  Name   Logical-port     Sessions   Logical-port          Sessions
bb0.1      alpha  up:NYC:xe-2/0/0  1125       up:Jersey:xe-2/0/0    1120
bb0.2      beta   up:NYC:xe-3/0/0  4588       up:Jersey:xe-2/0/0    4588
Address Domains:
   Name                     Prefixes User-Plane          Programmed   User-Plane  Programmed
   Domain-foo:NYC:default   125      up-exampl-1         125          Jersey      125
   Domain-bar:NYC:bar       255      up-exampl-1         255          Jersey      240

Routing Instances:
   Name         Tag    Backup-Tag   User-Plane           Programmed   User-Plane  Programmed
   default      55     77           up-exampl-1          added        Jersey          adding
   bar          277    314          up-exampl-1          removing     Jersey          removed
```

show user-plane mld interface

IN THIS SECTION

Syntax

```
show user-plane mld interface up-name user-plane-name
```

Description

Displays information about multipoint Listener Discovery (MLD)-enabled interfaces.

Options

none	When you run this command on the BNG CUPS Controller, the output displays standard information about all MLD-enabled interfaces on all BNG User Planes associated to the BNG CUPS Controller. When you run this command on a BNG User Plane, the output displays standard information about all MLD-enabled interfaces on the BNG User Plane.
up-name *user-plane-name*	(Optional) Displays information about the MLD-enabled interfaces on the specified BNG User Plane.

Required Privilege Level

view

Output Fields

Table 19 on page 269 describes the output fields for the `show user-plane mld interface` command. Output fields are listed in the approximate order in which they appear.

Table 19: show user-plane mld interface Output Fields

Field Name	Field Description
`Interface`	Name of the interface.
`Querier`	Address of the router that has been elected to send membership queries.
`State`	State of the interface: **Up** or **Down**.
`Up Timeout`	How long until the MLD querier is declared to be unreachable, in seconds.
`Version`	MLD version being used on the interface: **1** or **2**.
Groups	Number of groups on the interface.

Table 19: show user-plane mld interface Output Fields *(Continued)*

Field Name	Field Description
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. The `passive` statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • **Off**—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.
Distributed	State of MLD, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed MLD is enabled.

Table 19: show user-plane mld interface Output Fields *(Continued)*

Field Name	Field Description
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept MLD reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept MLD reports only from subnetworks that are associated with its interfaces.

Sample Output

show user-plane mld interface up-name

```
user@host> show user-plane mld interface up-name up-example-1
IInterface: pp0.3221225473
    Querier: ::
    State:         Up Timeout:    None Version:  2
    Group threshold:   0
    Immediate leave: Off
    Promiscuous mode: Off
    Passive: Off
    Distributed: On
```

Release Information

Statement introduced in Juniper BNG CUPS Release

show user-plane pppoe interfaces

IN THIS SECTION

Syntax

```
show user-plane pppoe interfaces
<brief | detail | extensive>
<up-name user-plane-name>
```

Description

Display session-specific information about PPPoE interfaces.

Options

none	Display interface information for all PPPoE interfaces.
brief | detail | extensive	(Optional) Display the specified level of output.
up-name ***user-plane-name***	The BNG User Plane for which you want to view PPPoE interface information.

Required Privilege Level

view

Output Fields

Table 20 on page 273 lists the output fields for the `show user-plane pppoe interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 20: show user-plane pppoe interfaces Output Fields

Field Name	Field Description	Level of Output
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail` `extensive` none
`State`	State of the logical interface: `up` or `down`.	All levels
`Session ID`	Session ID.	All levels
`Type`	Origin of the logical interface: `Static` or `Dynamic`. Indicates whether the interface was statically or dynamically created.	`detail` `extensive` none
`Service name`	Type of service required (can be used to indicate an ISP name or a class or quality of service).	`detail` `extensive` none
`Configured AC name`	Configured access concentrator name.	`detail` `extensive` none
`Session AC name`	Name of the access concentrator.	`detail` `extensive` none
`Remote MAC address` or `Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	All levels

Table 20: show user-plane pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Session uptime`	Length of time the session has been up, in *hh:mm:ss*.	`detail` `extensive` none
`Dynamic Profile`	Name of the dynamic profile that was used to create this interface. If the interface was statically created, this field is not displayed.	`detail` `extensive` none
`Underlying interface`	Interface on which PPPoE is running.	All levels
`Agent Circuit ID`	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. If the agent circuit ID is not configured, this field is not displayed.	`detail` `extensive` none
`Agent Remote ID`	Agent remote identifier that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. If the agent remote ID is not configured, this field is not displayed.	`detail` `extensive` none
`ACI Interface Set`	Internally-generated name of the dynamic ACI interface set, if configured, and the set index number of the ACI entry in the session database.	`detail` `extensive` none

Table 20: show user-plane pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Packet Type	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • PADI—PPPoE Active Discovery Initiation packets. • PADO—PPPoE Active Discovery Offer packets. • PADR—PPPoE Active Discovery Request packets. • PADS—PPPoE Active Discovery Session-Confirmation packets. • PADT—PPPoE Active Discovery Termination packets. • Service name error—Packets for which the Service-Name request could not be honored. • AC system error—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • Generic error—Packets that indicate an unrecoverable error occurred. • Malformed packets—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • Unknown packets—Unrecognized packets.	extensive

Sample Output

show user-plane pppoe interfaces up-name

```
user@host> show user-plane pppoe interfaces up-name up-test1
up:green-arrow:pp0.3221225473 Index 536870923
  State: Session Up, Session ID: 1, Type: Dynamic,
  Service name: AGILENT, Remote MAC address: 00:00:64:02:01:02,
  Session AC name: bng-controller,
  Session uptime: 00:00:19 ago,
```

```
Dynamic Profile: ppp-dp-pp0,
Underlying interface: up:green-arrow:ge-0/3/5.1 Index 3
```

show user-plane pppoe lockout

IN THIS SECTION

Syntax

```
show user-plane pppoe lockout
<underlying-interface-name>
<up-name user-plane-name>
```

Description

Display summary information about PPPoE clients currently undergoing lockout or currently in a lockout grace period on all PPPoE underlying logical interfaces or on a specified PPPoE underlying logical interface, for the specified BNG User Plane. You can configure PPPoE subscriber session lockout, also known as short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

none	Display information about the lockout condition and the lockout grace period for PPPoE clients on all PPPoE underlying logical interfaces.

underlying-interface-name	(Optional) Name of the PPPoE underlying logical interface. If you do not specify an underlying interface, the router iteratively displays output for all existing clients undergoing lockout per PPPoE underlying logical interface.
up-name *user-plane-name*	The BNG User Plane for which you want to view PPPoE information.

Required Privilege Level

view

Output Fields

Table 21 on page 277 lists the output fields for the `show user-plane pppoe lockout` command. Output fields are listed in the approximate order in which they appear.

Table 21: show user-plane pppoe lockout Output Fields

Field Name	Field Description
`underlying-interface-name`	Name of the PPPoE underlying logical interface.
`Index`	Index number of the logical interface, which reflects its initialization sequence.
`Device`	Name of the physical interface or aggregated Ethernet bundle.
`SVLAN`	Stacked VLAN ID, also known as the *outer tag*.
`VLAN`	VLAN ID, also know as the *inner tag*.
`VPI`	Virtual path identifier value for the PPPoE client.
`VCI`	Virtual circuit identifier value for the PPPoE client.

Table 21: show user-plane pppoe lockout Output Fields *(Continued)*

Field Name	Field Description
`Short-Cycle Protection`	State of PPPoE short-cycle protection, also known as PPPoE subscriber session lockout, on the underlying interface: • `circuit-id`—Filters PPPoE client sessions by their agent circuit identifier (ACI) value when configured for short-cycle protection • `mac-address`—Filters PPPoE client sessions by their unique media access control (MAC) address when configured for short-cycle protection • `off`—Short-cycle protection not configured for PPPoE client sessions Enabling short-cycle protection temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time.
`Lockout Time (seconds)`	Displays the PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period: • `Min`—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • `Max`—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • `Total clients in lockout`—Number of PPPoE clients currently undergoing lockout. • `Total clients in lockout grace period`—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.
`Client Address`	MAC source address or agent circuit idenfiier (ACI) value of the PPPoE client.
`Current`	Current lockout time, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout.
`Elapsed`	Time elapsed into the lockout period, in seconds; displays `0` (zero) if the PPPoE client is not undergoing lockout
`Next`	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.

Sample Output

show user-plane pppoe lockout up-name

```
user@host> show user-plane pppoe lockout at-1.0.0.30 up-name test-up1
at-1/0/0.30 Index 10305
Device: at-1/0/0, VPI: 1, VCI: 30
Short Cycle Protection: circuit-id,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 1
    Total clients in lockout grace period: 1

  Client Address                      Current   Elapsed    Next
    Relay-identifier atm 3/0:100.33        64        22     128
      00:00:5e:00:53:ab
      00:00:5e:00:53:21
```

show user-plane pppoe lockout up-name

```
user@host> show user-plane pppoe lockout demux0.100 up-name test-up1
demux0.100 Index 10305
Device: xe-1/0/0, SVLAN: 100, VLAN: 100,
  Short-Cycle Protection: mac-address,
  Lockout Time (seconds): Min: 1, Max: 300
    Total clients in lockout: 3
    Total clients in lockout grace period: 1

  Client Address                      Current   Elapsed    Next
    00:00:5e:00:53:15                      16        10      32
    00:00:5e:00:53:ab                     256       168     300
    00:00:5e:00:53:23                       0         0       8
```

show user-plane pppoe service-name-tables

IN THIS SECTION

Syntax

```
show user-plane pppoe service-name-tables
<table-name>
<up-name user-plane-name>
```

Description

Display configuration information about PPPoE service name tables, for the specified BNG User Plane.

Options

none	Display the names of configured PPPoE service name tables.
table-name	(Optional) Name of a configured PPPoE service name table.
up-name *user-plane-name*	The BNG User Plane for which you want to view PPPoE service name table information.

Required Privilege Level

view

Output Fields

Table 22 on page 281 lists the output fields for the `show user-plane pppoe service-name-tables` command. Output fields are listed in the approximate order in which they appear.

Table 22: show user-plane pppoe service-name-tables Output Fields

Field Name	Field Description	Level of Output
`Service Name Table`	Name of the PPPoE service name table.	none
`Service Name`	Name of a configured service in the PPPoE service name table: • `<empty>`—Service of zero length that represents an unspecified service • `<any>`—Default service for non-empty service entries that do not match the configured empty or named service entries • *`service-name`*—Named service entry	none
`Action`	Action taken when the PPPoE underlying interface interface receives a PPPoE Active Discovery Initiation (PADI) packet with the specified named service, `empty` service, `any` service, or ACI/ARI pair: • `Delay` *`seconds`*—Number of seconds that the interface delays before responding with a PPPoE Active Discovery Offer (PADO) packet • `Drop`—Interface drops (ignores) the packet. • `Terminate`—Interface responds immediately with a PADO packet	none
`Dynamic Profile`	Name of the dynamic profile with which the router creates a dynamic PPPoE subscriber interface. A dynamic profile can be assigned to a named service, `empty` service, `any` service, or ACI/ARI pair.	none
`Routing Instance`	Name of the routing instance in which to instantiate the dynamic PPPoE subscriber interface. A routing instance can be assigned to a named service, `empty` service, `any` service, or ACI/ARI pair.	none

Table 22: show user-plane pppoe service-name-tables Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Max Sessions`	Maximum number of active PPPoE sessions that the router can establish with the specified named service, `empty` service, or `any` service.	none
`Active Sessions`	Current count of active PPPoE sessions created using the specified named service, `empty` service, or `any` service. The Active Sessions value cannot exceed the Max Sessions value.	none
`ACI`	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. An ACI can be configured as part of an ACI/ARI pair for a named service, `empty` service, or `any` service.	none
`ARI`	Agent remote identifier (ARI) that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. An ARI can be configured as part of an ACI/ARI pair for a named service, `empty` service, or `any` service.	none
`Static Interface`	Name of the static PPPoE interface reserved for exclusive use by the PPPoE client with matching ACI/ARI information. A static interface can be configured only for an ACI/ARI pair.	none

Sample Output

show user-plane pppoe service-name-tables up-name

```
user@host> show user-plane pppoe service-name-tables up-name test-up1
Service Name Table: test1
Service Name Table: test2
Service Name Table: test3
```

show user-plane pppoe service-name-tables up-name

```
user@host> show user-plane pppoe service-name-tables Table1 up-name test-up1
Service Name Table: Table1
Service Name: <empty>
Action: Terminate
Dynamic Profile: BasicPppoeProfile
Max Sessions: 100
Active Sessions: 3
Service Name: <any>
Action: Drop
ACI: velorum-ge-2/0/3
ARI: westford
Action: Terminate
Static Interface: pp0.100
ACI: volantis-ge-5/0/5
ARI: sunnyvale
Action: Terminate
Static Interface: pp0.101
Service Name: Wholesale
Action: Terminate
Dynamic Profile: WholesalePppoeProfile
Routing Instance: WholesaleRI
Max Sessions: 16000
Active Sessions: 4
```

show user-plane pppoe sessions

IN THIS SECTION

Syntax

```
show user-plane pppoe sessions
<aci circuit-id-string>
<ari remote-id-string>
<service service-name>
<up-name user-plane-name>
```

Description

Display information about all active PPPoE sessions on the router, or about the active PPPoE sessions established for a specified service name, agent circuit identifier (ACI), or agent remote identifier (ARI).

Options

none	Display information for all active PPPoE sessions on the router.
aci *circuit-id-string*	(Optional) Display information only for active PPPoE sessions established with the specified agent circuit identifier. The agent circuit identifier corresponds to the DSLAM interface that initiated the service request.
ari *remote-id-string*	(Optional) Display information only for active PPPoE sessions established with the specified agent remote identifier. The agent remote identifier corresponds to the subscriber associated with the DSLAM interface that initiated the service request.
service *service-name*	(Optional) Display information only for active PPPoE sessions established with the specified service, where *service-name* can be `empty`, `any`, or a named service.
up-name *user-plane-name*	The BNG User Plane for which you want to view active PPPoE sessions.

Required Privilege Level

view

Output Fields

Table 23 on page 285 lists the output fields for the `show user-plane pppoe sessions` command. Output fields are listed in the approximate order in which they appear.

Table 23: show user-plane pppoe sessions Output Fields

Field Name	Field Description	Level of Output
`Interface`	Name of the statically-created or dynamically-created PPPoE interface for the active PPPoE session.	none
`Underlying interface`	Interface on which PPPoE is running.	none
`State`	State of the PPPoE session; displays `Session Up` for active PPPoE sessions.	none
`Session ID`	PPPoE session identifier.	none
`Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	none

Sample Output

show user-plane pppoe sessions up-name

```
user@host> show user-plane pppoe sessions up-name test-up1
Interface Underlying State Session
Remote
interface ID MAC
up:green-arrow:pp0.3221225473 up:green-arrow:ge-0/3/5.1 Session Up 1 00:00:64:02:01:02
```

show user-plane pppoe statistics

IN THIS SECTION

Syntax

```
show user-plane pppoe statistics
<logical-interface-name>
<up-name user-plane-name>
```

Description

Display statistics information about PPPoE interfaces.

Options

none	Display PPPoE statistics for all interfaces.
logical-interface-name	(Optional) Name of a PPPoE underlying logical interface.
up-name *user-plane-name*	The BNG User Plane for which you want to view PPPoE statistics.

Required Privilege Level

view

Output Fields

Table 24 on page 287 lists the output fields for the `show user-plane pppoe statistics` command. Output fields are listed in the approximate order in which they appear.

Table 24: show user-plane pppoe statistics Output Fields

Field Name	Field Description
`Active PPPoE sessions`	Total number of active PPPoE sessions and the number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.
`Timeouts`	Information about timeouts that occurred during the PPPoE session (not displayed for M120, M320, and MX Series routers): • `PADI`—No PADR packet has been received within the timeout period. (This value is always zero and is not supported.) • `PADO`—No PPPoE Active Discovery Offer packet has been received within the timeout period. • `PADR`—No PADS packet has been received within the timeout period.

Sample Output

show user-plane pppoe statistics up-name

```
user@host> show user-plane pppoe statistics up-name test-up1
Active PPPoE sessions: 32000
PacketType Sent Received
PADI 0 60216
PADO 60216 0
PADR 0 60216
PADS 60216 0
PADT 0 28178
Service name error 0 0
AC system error 0 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0
Active PPPoE sessions: 53326
PacketType Sent Received
PADI 0 244012
PADO 244012 0
PADR 0 244287
PADS 244287 0
PADT 1 187851
Service name error 0 0
AC system error 275 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0
Active PPPoE sessions: 54598
PacketType Sent Received
PADI 0 242606
PADO 242606 0
PADR 0 242774
PADS 242774 0
PADT 0 185503
Service name error 0 0
AC system error 168 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0
```

show user-plane pppoe underlying-interfaces

IN THIS SECTION

Syntax

```
show user-plane pppoe underlying-interfaces
<brief | detail | extensive>
<lockout>
<logical-interface-name>
<up-name user-plane-name>
```

Description

Display information about PPPoE underlying interfaces.

Options

brief | detail | extensive	(Optional) Display the specified level of output.
lockout	(Optional) Display summary information about the lockout condition and the lockout grace period for PPPoE clients on the PPPoE underlying interface.
logical-interface-name	(Optional) Name of a PPPoE underlying logical interface.
up-name *user-plane-name*	The BNG User Plane for which you want to view PPPoE underlying interfaces information.

Required Privilege Level

view

Output Fields

Table 25 on page 290 lists the output fields for the `show user-plane pppoe underlying-interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 25: show user-plane pppoe underlying-interfaces Output Fields

Field Name	Field Description	Level of Output
Underlying Interface	Name of the PPPoE underlying logical interface.	All levels
Service Name Table	Name of the service name table.	All levels
Dynamic Profile	Name of the dynamic profile that was used to create this interface. If the interface was statically created, then the value is **none**.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive`
State	Origin of the logical interface: **Static** or **Dynamic**. Indicates whether the interface was statically or dynamically created.	`detail extensive`
Operational States	Fields in this block are actual operational values rather than simply the configured values. The operational values can be the result of RADIUS-initiated changes.	`detail extensive`
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface. When this number of logical interfaces has been established, all subsequent PPPoE Active Discovery Initiation (PADI) packets are dropped and all subsequent PPPoE Active Discovery Request (PADR) packets trigger PPPoE Active Discovery Session (PADS) error responses.	`detail extensive`

Table 25: show user-plane pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Max Sessions VSA Ignore	Whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configure with the `max-sessions` statement: **Off** (default) or **On**.	`detail` `extensive` none
Active Sessions	Number of active PPPoE sessions on the underlying interface. If a dynamic profile is listed, then it is the number of active PPPoE sessions on the underlying interface that are using this profile. The Active Sessions value must not exceed the Max Sessions value.	`detail` `extensive`
Agent Circuit Identifier	Whether the underlying interface is configured with the `agent-circuit-identifier` statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information. `Autosensing` indicates that creation of ACI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ACI-based dynamic VLANs is not configured on the underlying interface, this field does not appear. **NOTE**: The Agent Circuit Identifier field is replaced with the Line Identity field when an ALI interface set is configured with the `line-identity` autoconfiguration stanza.	`detail` `extensive` none
`Line Identity`	Whether the underlying interface is configured with the `line-identity` statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on the specified trusted option: ACI, ARI, both, or neither. `Autosensing` indicates that creation of ALI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ALI dynamic VLANs based on trusted options is not configured on the underlying interface, this field does not appear. **NOTE**: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the `agent-circuit-id` autoconfiguration stanza.	`detail` `extensive` none

Table 25: show user-plane pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client. The uniqueness of the PPPoE client is determined by the client's MAC address.	`detail extensive`
Short Cycle Protection	State of PPPoE short cycle protection: **mac-address**, **circuit-id**, or **Off**. Enabling short cycle protection, also known as PPPoE lockout, on the PPPoE underlying interface temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time. PPPoE client sessions are identified by their unique media access control (MAC) source address or agent circuit identifier (ACI) value.	`detail extensive`
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	`detail extensive` none
AC Name	Name of the access concentrator.	`detail extensive`

Table 25: show user-plane pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
PacketType	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.	`detail` `extensive`

Table 25: show user-plane pppoe underlying-interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Lockout Time (sec)	The PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period if **Short Cycle Protection** is enabled (**On**): • **Min**—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • **Max**—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • **Total clients in lockout**—Number of PPPoE clients currently undergoing lockout. • **Total clients in lockout grace period**—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.	`extensive`
Client Address	MAC source address of the PPPoE client.	`extensive`
Current	Current lockout time, in seconds; displays **0** (zero) if the PPPoE client is not undergoing lockout.	`extensive`
Elapsed	Time elapsed into the lockout period, in seconds; displays `0` if the PPPoE client is not undergoing lockout	`extensive`
Next	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.	`extensive`

Sample Output

show user-plane pppoe underlying-interfaces brief up-name

```
user@host> show use-plane pppoe underlying-interfaces brief up-name test-up1
Underlying Interface Service Name Table Dynamic Profile
```

```
ge-4/0/3.1 Premium None
ge-4/0/3.2 None PppoeProfile
```

show user-plane pppoe underlying-interfaces detail up-name

```
user@host> show use-plane pppoe underlying-interfaces detail up-name test-up1ge-4/0/3.1 Index 73
Operational States:
State: Static, Dynamic Profile: None,
Max Sessions: 4000, Max Sessions VSA Ignore: Off,
Active Sessions: 0,
Service Name Table: Premium,
Direct Connect: Off,
AC Name: velorum, Duplicate Protection: On,
Short Cycle Protection: Off
ge-4/0/3.2 Index 78
Operational States:
State: Dynamic, Dynamic Profile: PppoeProfile,
Max Sessions: 500, Max Sessions VSA Ignore: Off,
Active Sessions: 3,
Service Name Table: None,
Direct Connect: Off,
AC Name: velorum, Duplicate Protection: On,
Short Cycle Protection: Off
```

show user-plane pppoe underlying-interfaces extensive up-name

```
user@host> show use-plane pppoe underlying-interfaces extensive up-name test-up1
ge-4/0/3.1 Index 73ge-4/0/3.1 Index 73
Operational States:
1053
State: Static, Dynamic Profile: None,
Max Sessions: 4000, Max Sessions VSA Ignore Off,
Active Sessions: 0,
Service Name Table: None,
Direct Connect: Off,
AC Name: velorum, Duplicate Protection: Off,
Short Cycle Protection: Off
PacketType Sent Received
PADI 0 0
PADO 0 0
```

```
PADR 0 0
PADS 0 0
PADT 0 0
Service name error 0 0
AC system error 0 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0
ge-4/0/3.2 Index 78
Operational States:
State: Dynamic, Dynamic Profile: PppoeProfile,
Max Sessions: 4000, Max Sessions VSA Ignore: Off
Active Sessions: 3,
Service Name Table: None,
Direct Connect: Off,
AC Name: velorum, Duplicate Protection: Off,
Short Cycle Protection: Off
PacketType Sent Received
PADI 0 5
PADO 5 0
PADR 0 5
PADS 4 0
PADT 0 1
Service name error 0 0
AC system error 0 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0f
```

show user-plane route

IN THIS SECTION

Syntax

```
show user-plane route
<family family>
<incomplete>
<level (brief | detail)>
<next-hop index>
<prefix>
<routing-instance name>>
<route-type type>
<rrt-index index>
<summary> user-plane-name
<up-name> user-plane-name
```

Description

Display information about how routes are mapped to specific enhanced subscriber management interfaces. You can customize the output by including one or more optional filters in the command. With the exception of the `summary` option, all filter options can be combined in a single command.

NOTE: This command is only run on BNG User Planes.

Options

`family family`	(Optional) Display route mapping information for the specified protocol family: `inet` (IPv4) or `inet6` (IPv6).
`incomplete`	(Optional) Display route mapping information for incomplete routes that are missing elements required to add the routes to the routing table.
`level (brief | detail)`	(Optional) Display the specified level of output: `brief` or `detail`.

`next-hop` *index*	(Optional) Display the next hop associated with the route entry with the specified next-hop index, in the range 1 through 65535.
`prefix` *address*	(Optional) Use the same prefix and prefix length as the subscriber host address. Output includes attributes that originate in the Famed-Route record of an upstream RADIUS server (Tag, Metric, Preference).
`route-type` *type*	(Optional) Display route mapping information for the specified route type: `access`, `access-internal`, `kernel`, or `local`.
`routing-instance` *name*	(Optional) Display route mapping information for the specified routing-instance
`rrt-index` *index*	(Optional) Display mapping information for the specified routing table index, in the range 0 through 65535. An `rtt-index` value of 0 (zero) denotes routes in the default routing table managed by enhanced subscriber management.
`summary` *user-plane-name*	(Optional) Display summary information about the routes managed by enhanced subscriber management for the specified BNG User Plane.
`up-name` *user-plane-name*	The BNG User Plane for which to display information about the routes managed by enhanced subscriber management.

Required Privilege Level

view

Output Fields

Table 26 on page 298 lists the output fields for the `show user-plane route` command. Output fields are listed in the approximate order in which they appear.

Table 26: show user-plane route Output Fields

Field Name	Field Description	Level of Output
address	IPv4 or IPv6 address associated with the route entry.	All levels
Route	IPv4 or IPv6 address associated with the route entry.	All levels

Table 26: show user-plane route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Route Type`	One of the following route types: • `Access` • `Access-internal` • `Framed` • `Kernel` • `Local`	All levels
`Interface`	Name of the enhanced subscriber management interface associated with the route entry.	All levels
`Next-hop`	Next-hop associated with the route entry.	All levels
`Tag`	Reflects the Tag attribute used in the RADIUS Framed-Route type record.	All levels
`Metric`	Reflects the Metric attribute used in the RADIUS Framed-Route type record.	All levels
`Preference`	Reflects the Preference attribute used in the RADIUS Framed-Route type record.	All levels
`Rtt-index`	Value of the routing table index. A value of `0` (zero) denotes a route in the default routing table managed by enhanced subscriber management.	`detail`
`Bbe index`	Value of the interface index for the control plane.	`detail`
`Flow id`	Value of the route object index.	`detail`
`Reference Count`	Used for internal accounting.	`detail`

Table 26: show user-plane route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Discard route count	Number of discard routes.	Summary
Discard route gateway	Number of gateway routes.	Summary
`Dirty Flags`	Used for internal accounting.	`detail`
`Flags`	Used for internal accounting.	`detail`
`Family`	One of the following protocol families: • `AF_INET`—IPv4 • `AF_INET6`—IPv6	`detail`
`UP Name`	Name of the BNG User Plane.	All levels
`Kernel rt-table id-instance`	The kernel routing table ID number.	Summary
`Local route count`	The number of local routes.	Summary
`Access route count`	The number of access routes.	Summary
`Access internal route count`	The number of access internal routes.	Summary
`Kernel route count`	The number of kernel routes.	Summary
`Dirty local route count`	The number of local routes that have not been fully installed. It is always 0 for the active RE for the BNG User Plane. It can be non-zero for the standby RE (representing a transient condition).	Summary

Table 26: show user-plane route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Dirty access route count`	The number of access routes that have not been fully installed. It is always 0 for the active RE for the BNG User Plane. It can be non-zero for the standby RE (representing a transient condition).	Summary
`Dirty access internal route count`	The number of access-internal routes that have not been fully installed. It is always 0 for the active RE for the User Plane. It can be non-zero for the standby RE (representing a transient condition).	Summary
`Dirty kernel route count`	The number of kernel routes that have not been fully installed. It is always 0 for the active RE for the BNG User Plane. It can be non-zero for the standby RE (representing a transient condition).	Summary
`Dirty inflight route count`	The number of inflight routes that have not been fully installed. It is always 0 for the active RE for the BNG User Plane. It can be non-zero for the standby RE (representing a transient condition).	Summary

Sample Output

show user-plane route up-name

```
user@host> show user-plane route up-name up-example-1
Route:  193.0.21.0/32
     Route Type:              Access-internal
     Interface:               pp0.3221324082
     Next-Hop index:          0
Route:  193.0.21.1/32
     Route Type:              Access-internal
     Interface:               pp0.3221324088
     Next-Hop index:          0
Route:  193.0.21.2/32
     Route Type:              Access-internal
     Interface:               pp0.3221324092
     Next-Hop index:          0
Route:  193.0.21.3/32
```

```
    Route Type:                 Access-internal
    Interface:                  pp0.3221324094
    Next-Hop index:             0
```

show user-plane route summary up-name up-exampl-1

```
user@host> show user-plane route summary up-name up-example-1
UP Name:  up-example-1
  Routing-instance:  default:default
     Kernel rt-table id:                  0
     Family:                              AF_INET
     Local route count:                   1
     Access route count:                  0
     Access internal route count:         31985
     Kernel route count:                  0
     Dirty local route count:             0
     Dirty access route count:            0
     Dirty access internal route count:   0
     Dirty kernel route count:            0
     Dirty inflight route count:          0
```

show user-plane route summary

```
user@host> show user-plane route summary
Routing-instance: default:default
Kernel rt-table id: 0
Family: AF_INET
Local route count: 0
Access route count: 4
Access internal route count: 604
Kernel route count: 0
Discard route count: 20
Gateway route count: 1
Dirty local route count: 0
Dirty access route count: 0
Dirty access internal route count: 0
Dirty kernel route count: 0
Dirty inflight route count: 0
```

show user-plane route route-type discard

```
user@host> show route route-type discard
Route: 173.162.0.0/24
Route Type: Discard
Tag: 33
Next-Hop index: 0
Route: 173.162.0.0/24
Route Type: Discard
Tag: 33
Next-Hop index: 0
Route: 173.162.0.0/24
Route Type: Discard
Tag: 33
Next-Hop index:
```

show user-plane route prefix <address>

`rtt-index 0`

```
user@host> show user-plane route prefix 10.10.0.1/32
Route:  10.10.0.1/32
     Routing-instance:          default:default
     Kernel rt-table id :       0
     Family:                    AF_INET
     Route Type:                Framed
     Protocol Type:             Unspecified
     Interface:                 pp0.3221225491
     Interface index:           26
     Internal Interface index:  26
     Route index:               20
     Next-Hop:                  684
     Tag:                       9999
     Metric:                    56
     Preference:                10
     Reference-count:           1
     L2 Address:                00:00:5e:00:53:0b
     Flags:                     0x0
     Dirty Flags:               0x0
```

show user-plane route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show user-plane route family inet6 route-type access rtt-index 0 level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

show user-plane route family route-type rtt-index level detail

The following example displays detailed information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show user-plane route family inet6 route-type access rtt-index 0 level detail
2001:db8::/64
     Route Type:      Access
     Interface:       pp0.3221225479
     Next-hop:        721
```

```
        Rtt-index:       0
        Bbe index:       9
        Flow id:         1
        Reference Count: 1
        Dirty Flags:     0
        Flags:           0x10082
        Family:          AF_INET6
    2001:db8:0:0:1::/64
        Route Type:      Access
        Interface:       pp0.3221225477
        Next-hop:        721
        Rtt-index:       0
        Bbe index:       9
        Flow id:         1
        Reference Count: 1
        Dirty Flags:     0
        Flags:           0x10082
        Family:          AF_INET6
    2001:db8:0:0:2::/64
        Route Type:      Access
        Interface:       pp0.3221225478
        Next-hop:        721
        Rtt-index:       0
        Bbe index:       9
        Flow id:         1
        Reference Count: 1
        Dirty Flags:     0
        Flags:           0x10082
        Family:          AF_INET6
    2001:db8:0:0:3::/64
        Route Type:      Access
        Interface:       pp0.3221225480
        Next-hop:        721
        Rtt-index:       0
        Bbe index:       9
        Flow id:         1
        Reference Count: 1
        Dirty Flags:     0
        Flags:           0x10082
        Family:          AF_INET6
```

show user-plane route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show user-plane route family inet6 route-type access rtt-index 0 level brief
2001:db8::/64
     Route Type: Access
     Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
     Route Type: Access
     Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
     Route Type: Access
     Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
     Route Type: Access
     Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
     Route Type: Access
     Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

show user-plane routing-instances

IN THIS SECTION

Syntax

```
show user-plane routing instances up-name <bng-user-plane-name>
```

Description

Displays routing instances in use by a particular BNG User Plane.

Options

bng-user-plane The BNG User Plane for which you want to know the routing instance that are being used.

Required Privilege Level

view

Output Fields

Table 26 on page 298 lists the output fields for the `show user-plane routing-instances` command. Output fields are listed in the approximate order in which they appear.

Table 27: show user-plane routing-instances Output Fields

Field Name	Field Description
`User Plane`	Name of the BNG User Plane.

Table 27: show user-plane routing-instances Output Fields *(Continued)*

Field Name	Field Description
Routing Instance Name	Name of the routing instance.
`State`	The routing instance state: • Connected—The node is connected to the network. • Isolated—The node is isolated from the rest of the network.

Sample Output

show user-plane routing-instances

```
user@host> show user-plane routing-instances example-1
User-plane: example-1
Routing Instance Name            State
default                          connected
RETAILER33                       connected
RETAILER0                        connected
```

show user-plane statistics

IN THIS SECTION

Syntax

```
show user-plane statistics
<all>
<dhcp>
<dvlan>
<l2tp>
<ppp>
<pppoe>
<up-name user-plane-name>
```

Description

Display statistics for the specified BNG User Plane. You can customize the output by including one or more optional filters in the command.

Options

all	(Optional) Display packet statistics for all protocols.
dhcp	(Optional) Display DHCP packet statistics.
dvlan	(Optional) Display DVLAN packet statistics.
l2tp	(Optional) Display L2TP packet statistics.
ppp	(Optional) Display PPP packet statistics.
pppoe	(Optional) Display PPPoE packet statistics.
up-name *user-plane-name*	The BNG User Plane for which you want to view packet statistics.

Required Privilege Level

view

Output Fields

Table 28 on page 310 lists the output fields for the `show user-plan statistics` command. Output fields are listed in the approximate order in which they appear.

Table 28: show user-plane statistics Output Fields

Field Name	Field Description
`User-plane`	The BNG User Plane for which the information is being displayed.
`Rx Statistics`	Statistics for packets received.
`Tx Statistics`	Statistics for packets sent.
`Enhanced I/O Statistics`	Statistics for visibility into packet drops from the queue.
`Error Statistics`	Includes connection packets, flow control, and messages and packets sent to and received from the daemon.
`ERA discards`	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
`Layer 3 Statistics`	Statistics for Layer 3 packets.
`padis`	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
`padrs`	PPPoE Active Discovery Request packets.
`ppp`	Point-to-Point Protocol packets.
`router solicitations`	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.

Table 28: show user-plane statistics Output Fields *(Continued)*

Field Name	Field Description
`router advertisements`	Number of router advertisements sent or received.

Sample Output

show user-plane statistics up-name

```
user@host> show user-plane statistics up-name up-example-1
User-plane : up-example-1
I/O Statistics:
    Rx Statistics
        packets                          : 3059637
    Tx Statistics
        packets                          : 2837485
  Layer 3 Statistics
    Rx Statistics
        packets                          : 0
    Tx Statistics
        packets                          : 0
```

show user-plane statistics pppoe up-name

```
user@host> show user-plane statistics pppoe up-name up-example-1
User-plane : up-example-1
I/O Statistics:
    Rx Statistics
        packets                          : 3059637
    Tx Statistics
        packets                          : 2837485
  Layer 3 Statistics
    Rx Statistics
        packets                          : 0
    Tx Statistics
        packets                          : 0
PPPoE Statistics:
```

```
    Rx Statistics
        packets                           : 369141
        padis                             : 32027
        padrs                             : 32000
        ppp packets                       : 241057
```

5

CHAPTER

Juniper BNG User Plane CLI Configuration Statements

IN THIS CHAPTER

Juniper BNG User Plane CLI Configuration Statements

IN THIS SECTION

This topic provides an overview of configuration commands, including syntax and option descriptions, that you use with a Juniper router performing the role of a Juniper BNG User Plane.

A Juniper router that is performing the role of a BNG User Plane can be configured and managed like any standalone Junos OS router (see Junos CLI Reference).

Be aware that most of the configurations related to subscriber management are not applicable. This is because for Juniper BNG CUPS, subscriber management configurations are performed on the Juniper BNG CUPS Controller.

This topic focuses on configurations relevant to the following for a BNG User Plane:

- Establishing the role of a BNG User Plane
- Establishing communication with the BNG CUPS Controller

For any additional configurations that must be align with the configuration on the BNG CUPS Controller, see "Juniper BNG CUPS Controller CLI Configuration Statements" on page 124.

captive-portal-content-delivery

IN THIS SECTION

Syntax

```
captive-portal-content-delivery {
    profile name
        dynamic;
    }
    traceoptions {
        file <filename> <files number> <match match> <size size> <(world-readable | no-world-
readable)>;
        flag name;
        no-remote-trace no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber. Use the statement at the `[edit services]` hierarchy level for dynamic CPCD.

Options

`profile` *profile-name*	Configure the service profile for HTTP redirect services, also known as Captive Portal and Content Delivery (CPCD) services. The profile contains rules or rule sets that specify the details of the service. The CPCD service profile is included in a service set that applies the service to a service interface.
`dynamic`	Indicate that the service is a dynamic, converged service.
`traceoptions`	Define tracing operations for captive portal content delivery processes.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

no-usage-report

IN THIS SECTION

Syntax

```
no-usage-report;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Disable subscriber physical interface usage reporting to the BNG CUPS Controller. This command runs on the BNG User Planes.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

oversubscription-limit

IN THIS SECTION

Syntax

```
oversubscription-limit limit;
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type any fpc slot-number]
```

Description

Configure the maximum number of warm-enhanced oversubscribed subscribers on the backup BNG User Plane. As a backup BNG User Plane, if the `oversubscription-limit` is reached, the subsequent subscribers are logged out.

> **NOTE**: The oversubscription limit cannot be configured unless the `subscriber-limit` *`limit`* is configured.

Options

limit Number of oversubscription subscribers allowed (1 through 256000).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

policy-options

IN THIS SECTION

Syntax

```
policy-options
      policy-statement policy-name {
        term term-name {
            from {
                family family-name;
                match-conditions;
                policy subroutine-policy-name;
                prefix-list prefix-list-name;
                prefix-list-filter prefix-list-name match-type <actions>;
                route-filter destination-prefix match-type <actions>;
                source-address-filter source-prefix match-type <actions>;
            }
            to {
                match-conditions;
                policy subroutine-policy-name;
            }
            then actions;
        }
    }
```

Hierarchy Level

```
[edit]
```

Description

Configure options such as application maps for DCBX application protocol exchange and policy statements. This command runs on the BNG User Planes.

Required Privilege Level

storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

policy-statement

IN THIS SECTION

Syntax

```
policy-statement policy-name {
    term term-name {
        from {
            as-path-neighbors (as-list | as-list-group);
            as-path-origins (as-list | as-list-group);
            as-path-transits (as-list | as-list-group);
            as-path-unique-count count (equal | orhigher | orlower);
            as-path-calc-length count (equal | orhigher | orlower);
            family family-name;
            match-conditions;
            policy subroutine-policy-name;
            prefix-list prefix-list-name;
            prefix-list-filter prefix-list-name match-type <actions>;
            programmed;
            protocol protocol-name;
            route-filter destination-prefix match-type <actions>;
            validation-database-instance {
                                database <database-name> state (valid|invalid|unknown);
                                state (valid|invalid|unknown);
                                }
            source-address-filter source-prefix match-type <actions>;
            tag value;
            traffic-engineering;
        }
```

```
        to {
            match-conditions;
            policy subroutine-policy-name;
        }
        then actions;
    }
    then {
        advertise-locator;

        aggregate-bandwidth;
        dynamic-tunnel-attributes dynamic-tunnel-attributes;
        limit-bandwidth limit-bandwidth;
        multipath-resolve;
        no-entropy-label-capability;
        prefix-attribute-flags;
        prefix-segment {
            index index;
            node-segment;
        }
        priority (high | medium | low);
        resolution-map map-name;
        set-down-bit
    }
}
```

Hierarchy Level

```
[edit policy-options]
```

Description

Define a routing policy, including subroutine policies. This command runs on the BNG User Planes.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of `accept` or `reject` is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either `accept` or `reject`) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the `[edit policy-options]` hierarchy level by `policy-statement` *`policy-name`* in alphabetical order, enter the `show policy-options` configuration command.

The statements are explained separately.

Options

`actions`—(Optional) One or more actions to take if the conditions match.

`family` *`family-name`*—(Optional) Specify an address family protocol. Specify `inet` for IPv4. Specify `inet6` for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify `iso`. For IPv4 multicast VPN traffic, specify `inet-mvpn`. For IPv6 multicast VPN traffic, specify `inet6-mvpn`. For multicast-distribution-tree (MDT) IPv4 traffic, specify `inet-mdt`. For BGP route target VPN traffic, specify `route-target`. For traffic engineering, specify `traffic-engineering`.

> **NOTE**: When `family` is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the family statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

`from`—(Optional) Match a route based on its source address.

`as-path-neighbors (as-list | as-list-group)`—Compares the AS that originated the route. Evaluates if the right most AS number on the AS path belongs to the `as-list` or `as-list-group` specified in the `as-path-origins` configuration statement. In the case where the route has been aggregated, and the location of the originating AS contains an AS-set, the `as-path-origins` operator evaluates to true if any AS contained in the AS-set belongs to the `as-list` or `as-list-group` specified in the `as-path-origins` configuration statement.

`as-path-origins (as-list | as-list-group)`—Compares the neighbor AS in the AS path. Evaluates if the first AS number on the AS path matches the `as-list` or `as-list-group` specified in the `as-path-neighbors` configuration statement. If the neighboring AS location happens to be an AS-set, the `as-path-neighbors` operator evaluates to true if any AS contained in the AS-set belongs to the `as-list` or `as-list-group` specified in the `as-path-neighbors` configuration statement.

`as-path-transits (as-list | as-list-group)`—Compares any AS in the AS-Path. Evaluates when any AS belongs to the `as-list` or `as-list-group` specified in the `as-path-transit` configuration statement. In the case of AS-set, the `as-path-transit` operator compares all the ASes in the AS-set.

`as-path-calc-length` *`count`* `(equal | orhigher | orlower)`—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of calculated autonomous systems (ASs) in the AS path.

NOTE:

- ASs in a sequence count as 1.
- AS sets count as 1.
- BGP confederation segments count as 0.

`as-path-unique-count` *`count`* `(equal | orhigher | orlower)`—(Optional) Specify a number from 0 through 1024 to filter routes based on the total number of unique non-BGP confederation autonomous systems (ASs) in the AS path.

NOTE: Duplicate AS numbers are ignored for the count.

`advertise-locator`—(Optional) Enable IS-IS to summarize and advertise locator prefixes.

Range: *0-255*

`aggregate-bandwidth`—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing.

`dynamic-tunnel-attributes` *`dynamic-tunnel-attributes`*—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

`match-conditions`—(Optional in `from` statement; required in `to` statement) One or more conditions to use to make a match. The qualifiers are described in Routing Policy Match Conditions.

`multipath-resolve` *multipath-resolve*–(Optional) Enable the use of all paths for resolution over the specified prefix.

`limit-bandwidth` *limit-bandwidth*—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

- **Range:** 0 through 4,294,967,295 bytes

`no-entropy-label-capability`—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

`priority (high | medium | low)`—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

`policy` *subroutine-policy-name*—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see Understanding Policy Subroutines in Routing Policy Match Conditions.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

`prefix-list` *prefix-list-name*—Name of a list of IPv4 or IPv6 prefixes.

`prefix-list-filter` *prefix-list-name*—Name of a prefix list to evaluate using qualifiers; *match-type* is the type of match, and *actions* is the action to take if the prefixes match.

`programmed`—(Optional) Allow policy matches for routes injected by JET APIs.

`protocol` *protocol-name*—Name of the protocol used to control traffic engineering database import at the originating point.

You can specify options to match label IS-IS and label OSPF routes using the `l-isis` and `l-ospf` options, respectively. The `isis` options matches all IS-IS routes, excluding labelled IS-IS routes. The `ospf` option matches all OSPF routes, including OSPFv2, OSPFv3 and labelled OSPF routes.

`resolution-map`—(Optional) Set resolution map modes. A given resolution-map can be shared across multiple policy-statements.

`route-filter` *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see Configuring Route Lists), and *actions* is the action to take if the *destination-prefix* matches.

`source-address-filter` *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see Configuring Route Lists), and *actions* is the action to take if the *source-prefix* matches.

`tag` *`value`*—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

`term` *`term-name`*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ ”). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the `term` statement when defining match conditions and actions.

`to`—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

`then`—(Optional) Actions to take on matching routes. The actions are described in Configuring Flow Control Actions and Configuring Actions That Manipulate Route Characteristics.

`set-down-bit`—(Optional) Configure this option to aggregate leaked locator routes using routing policies.

`validation-database-instance`—(Optional) Name to identify a validation-state with database name.`database-name <database-name>`—(Optional) Route Validation Database name to be looked at. `state (valid|invalid|unknown)`—(Optional) Name to identify a validation-state

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

query-interval

IN THIS SECTION

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit protocols igmp]
```

```
{edit protocols mld]
```

Description

Specify how often the querier routing device sends general host-query messages. This command runs on the BNG User Planes.

Options

seconds—Time interval.

- **Range:** 1 through 1024
- **Default:** 125 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

query-last-member-interval

IN THIS SECTION

Syntax

```
query-last-member-interval seconds;
```

Hierarchy Level

```
[edit protocols igmp]
```

```
[edit protocols mld]
```

Description

Specify how often the querier routing device sends group-specific query messages. This command runs on the BNG User Planes.

Options

seconds—Time interval, in fractions of a second or seconds.

- **Range:** 0.1 through 0.9, then in 1-second intervals 1 through 1024
- **Default:** 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

query-response-interval

IN THIS SECTION

Syntax

```
query-response-interval seconds;
```

Hierarchy Level

```
[edit protocols igmp]
```

```
[edit protocols mld]
```

Description

Specify how long the querier routing device waits to receive a response to a host-query message from a host. This command runs on the BNG User Planes.

Options

seconds—The query response interval must be less than the query interval.

- **Range:** 1 through 1024
- **Default:** 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

robust-count

IN THIS SECTION

Syntax

```
robust-count number;
```

Hierarchy Level

```
[edit protocols igmp]
```

```
{edit protocols mld]
```

Description

Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count. This command runs on the BNG User Planes.

Options

number—Robustness variable.

- **Range:** 2 through 10
- **Default:** 2

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

routing-instance

IN THIS SECTION

Syntax

```
routing-instance routing-instance-name}
```

Hierarchy Level

```
[edit system services subscriber-management mode user-plane user-plane-name user-plane-name
transport]
```

Description

(Optional) Designate the routing instance for the BNG User Plane to communicate with the BNG CUPS Controller. If not specified, the default routing instance is used by the BNG User Plane to communicate with the BNG CUPS Controller.

Options

routing-instance-name	The name of the routing instance to use.

Required Privilege Level

root—To view this statement in the configuration.

root—To add this statement to the configuration.

selection-function

IN THIS SECTION

Syntax

```
selection-function {
   cluster cluster-name, cluster-name;
   service-group service-group-name, service-group-name;
}
```

Hierarchy Level

```
[edit system services subscriber-management mode user-plane]
```

Description

Sets the clusters in which the BNG User Plane is a member. Also, you can set the service class that the BNG User Plane supports within each cluster.

Options

`cluster` *`cluster-name`*	The name or names of the cluster to which the BNG User Plane belongs. You can enter one or more names.
`service-group` *`service-group-names`*	The names of the service classes that the BNG User Plane supports within each cluster. You can enter one or more names.

Required Privilege Level

root—To view this statement in the configuration.

root—To add this statement to the configuration.

user-plane

IN THIS SECTION

Syntax

```
user-plane {
    user-plane-name bng-user-plane-name;
    transport {
        inet ip-address | inet6 <varname>ip-address</varname>
    }
```

```
    control-plane {
        bng-controller-name bng-cups-controller-name;
        transport {
            inet ip-address | inet6 ip-address;
        }
    }
    pfcp {
        retransmission-timer seconds;
        retries number;
        heartbeat-interval seconds;
    }
    selection-function {
        cluster cluster-name, cluster-name;
        service-group service-group-name, service-group-name;
}
```

Hierarchy Level

```
[edit system services subscriber-management enable mode]
```

Description

Sets the system to take on the role of a BNG User Plane.

Options

`user-plane-name` ***`bng-user-plane-name`***

The *user-plane-name* attribute is described in the following:

- Is mandatory and identifies the BNG User Plane
- Must be unique within the domain of the BNG CUPS Controller
- Is a reference to the local system and can be 1 to 12 characters long, containing upper and lowercase letters, numerals, hyphens, and periods
- Must not start or end with a hyphen

`transport (user-plane)`

The transport stanza is a mandatory stanza that defines the source address from which the BNG User Plane initiates associations.

- inet *ip-address*—Specify the IPv4 address of the BNG User Plane.
- inet6 *ip-address*—Specify the IPv6 address of the BNG User Plane.

bng-controller-name *bng-cups-controller-name*

The `bng-controller-name` is a reference to the local system and can be 1 to 12 characters long. You can combine uppercase letters and lowercase letters, numbers, hyphens, and periods in this reference but cannot start or end it with a hyphen.

> **NOTE**: The name that is entered here for the `bng-controller-name` setting, must match the name that is configured for the `bng-controller-name` setting in the `bng-controller` configuration (see "bng-controller" on page 127).

transport (control-plane)

Defines the IPv4 or IPv6 address and port number of the BNG CUPS Controller with which the BNG User Plane attempts to make an association. The address family that you choose must match the family in the BNG User Plane's transport stanza.

- inet *ip-address*—Specify the IPv4 address of the BNG CUPS Controller.
- inet6 *ip-address*—Specify the IPv6 address of the BNG CUPS Controller.

pfcp

Specify the Packet Forwarding Control Protocol protocol attributes for the control plane manager and any other daemons using Packet Forwarding Control Protocol to communicate with their peer.

- `retransmission-timer`—Defines the retransmission interval in seconds.
  - **Default:** 5 seconds
  - **Range:** 3 through 30

    seconds
- `retries`—Defines the number of retransmission attempts.
  - **Default:** 5
  - **Range:** 5 through 10
- `heartbeat-interval`—Defines the interval in seconds between keep-alive messages.
  - **Default:** 60

    seconds
  - **Range:** 60 through 600

    seconds

auto-configure

IN THIS SECTION

Syntax

```
auto-configure {
    static-vlan-identity {
        interface-tag name
    }
    vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types  [packet-types];
            password password-string;
            username-include{
                circuit-id;
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                remote-id;
                user-prefix user-prefix-string;
```

```
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            accept  (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            accept-out-of-band protocol;
            ranges (any | low-tag)-(any | high-tag);
        }
        override;
    }
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            packet-types  [packet-types];
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-18;
                option-37;
                option-82 <circuit-id> <remote-id>;
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
            ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
        }
        override;
    }
    remove-when-no-subscribers;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Enable the configuration of dynamic, auto-sensed VLANs.

Options

static-vlan-identity (Only for Junos OS Evolved that support subscriber management) Static VLAN identity configuration. This configuration creates dynamic VLAN for static IFL. The `remove-when-no-subscribers` option is not applicable when you configure static-vlan-identity.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

captive-portal-content-delivery-profile (Services)

IN THIS SECTION

Syntax

```
captive-portal-content-delivery-profile profile-name
    interface-service {
        service-interface name;
    }
    next-hop-service {
        inside-service-interface interface-name.unit-number;
        outside-service-interface interface-name.unit-number;
        outside-service-interface-type local;
    }
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. Starting in Junos OS Release 19.3R2, you can configure converged services for SPC3s if you have enabled Next Gen Services on the MX Series router. At the `edit service-set` *`service set name`* `captive-portal-content-delivery-profile` *`profile-name`* `interface-service` hierarchy level, you can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the `service-interface ms-`*`fpc/pic/port`* statement, and configure captive portal content delivery (CPCD) profiles for SPC3s by including the `service-interface vms-`*`fpc/pic/port`* statement.

Options

`service-interface` *`name`*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 17.2.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

dynamic-profile

dynamic-profile (Stacked VLAN)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
    accept  (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
    access-profile vlan-dynamic-profile-name;
    ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Description

Configure a dynamic profile for use when configuring dynamic stacked VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a Basic Dynamic Profile

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

hierarchical-scheduler

IN THIS SECTION

Syntax

```
hierarchical-scheduler {
    implicit-hierarchy;
    maximum-hierarchy-levels number;
}
```

Hierarchy Level

```
[edit interfaces]
[edit dynamic-profiles name interfaces],
[edit dynamic-profiles name logical-systems name interfaces]
```

Description

Enable the use of hierarchical schedulers. If you do not include this statement, the interfaces on the router cannot use hierarchical interfaces.

> **NOTE**: When configuring hierarchical scheduling for network slices on an interface, you do not need to enable vlan-tagging on that interface.

NOTE: To enable hierarchical scheduling on MX80 and MX104 routers, configure the `hierarchical-scheduler` statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

implicit-hierarchy — For MPC/MIC subscriber interfaces and interface sets running over aggregated Ethernet on MX Series routers, form a hierarchical relationship between the CoS scheduler nodes at level 1, level 2, and level 3.

maximum-hierarchy-levels ***number*** — Set the maximum number of hierarchy levels.

- **Range:** 2 through 4
- **Default:** 3

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Configuring Hierarchical Schedulers for CoS

Understanding Hierarchical CoS for Subscriber Interfaces

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

ranges (Dynamic Stacked VLAN)

IN THIS SECTION

Syntax

```
ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name]
```

Description

Configure VLAN ranges for dynamic, auto-sensed stacked VLANs.

Options

`any`—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.

- **Range:** 1 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

service-set (Dynamic Service Sets)

IN THIS SECTION

Syntax

```
service-set service-set-name {
    service-filter filter-name;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service input],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service output],
[edit dynamic-profiles profile-name interfaces pp0 unit $junos-interface-unit family family service
input],
[edit dynamic-profiles profile-name interfaces pp0 unit  $junos-interface-unit family family service
output]
```

Description

Define one or more service sets in a dynamic profile. Service sets are applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration. You can use the predefined dynamic interface variables `$junos-input-service-set`, `$junos-output-service-set`, `$junos-input-ipv6-service-set`, and `$junos-output-ipv6-service-set`.

NOTE: Starting in Junos OS Release 17.2R1, you can configure converged services at the `edit dynamic-profiles http-redirect-converged` hierarchy level. CPCD rules can also be configured under the dynamic profiles stanza to achieve parameterization of the rules. This mechanism provides additional flexibility to customize the different rules on a per subscriber basis through service attachment.

Options

service-set-name—Name of the service set.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the `[edit dynamic-profiles` *profile-name* `interfaces pp0 unit “$junos-interface-unit” family` *family* `service input]` and `[edit dynamic-profiles` *profile-name* `interfaces pp0 unit “$junos-interface-unit” family` *family* `service output]` hierarchy levels introduced in Junos OS Release 10.1.

From 17.2R1 onwards, you can configure converged services at the `edit dynamic-profiles http-redirect-converged` hierarchy level.

RELATED DOCUMENTATION

Dynamic Service Sets Overview

Associating Service Sets with Interfaces in a Dynamic Profile

service-set-options

IN THIS SECTION

Syntax

```
service-set-options {
    bypass-traffic-on-exceeding-flow-limits;
    bypass-traffic-on-pic-failure;
    enable-asymmetric-traffic-processing;
    enable-descriptive-session-syslog;
    header-integrity-check;
    routing-engine-services;
    static-subscriber-application;
    subscriber-awareness;
    support-uni-directional-traffic;
    tcp-fast-open  {
        disabled;
        drop;
    }
    tcp-non-syn  {
        drop-flow;
        drop-flow-send-rst;
    }
 {
        input;
        output;
```

```
    }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Description

Specify the service set options to apply to a service set.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`enable-asymmetric-traffic-processing` and `support-uni-directional-traffic` options added in Junos OS Release 11.2.

`routing-engine-services` option added in Junos OS Release 15.1.

`enable-change-on-ams-redistribution` option added in Junos OS Release 15.1.

`subscriber-awareness` option added in Junos OS Release 17.1.

`tcp-fast-open` option added in Junos OS Release 17.2.

`enable-descriptive-session-syslog` option added in Junos OS Release 20.3.

`static-subscriber-application` option added in Junos OS Release 21.2.

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

Configuring APPID Support for Unidirectional Traffic

services (Captive Portal Content Delivery)

IN THIS SECTION

Syntax

```
services {
    ...
    captive-portal-content-delivery {
        auto-deactivate value;
        profile name
            cpcd-rule-sets rule-set-name;
            cpcd-rules rule-name;
            dynamic;
            http-redirect-options url;
            ipda-rewrite-options {
                destination-address destination-address;
                destination-port destination-port;
```

```
            }
        }
        rule rule-name {
            match-direction (input | output | input-output);
            from {
                destination-address address <except>;
            }
            Term term-name {
                then {
                    accept;
                    insert tag tag-name tag-value tag-value;
                    redirect url;
                    rewrite destination-address address <destination-port port-number>;
                    syslog;
                }
            }
        }
        rule-set rule-set-name {
            [rule rule-name];
        }
      traceoptions {
            file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
            flag name;
            no-remote-trace no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit],
[edit dynamic-profiles profile-name]
```

Description

CAUTION: Enabling tracing can adversely impact scale and performance and may increase security risk. We strongly recommend using the trace, tracing, or traceoptions commands only under the guidance of a JTAC support engineer. After collecting the debug information, immediately disable tracing to minimize risk and restore normal system performance.

Define the captive portal content delivery set of the rules statements to be applied to traffic. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles` *`profile-name`* `services...]` hierarchy level for converged services CPCD.

The `profile`, `rule-set`, and `traceoptions` stanzas are not supported at the `[edit dynamic-profiles` *`profile-name`* `hierarchy level]`.

Configure the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber. Use the statement at the `[edit services...]` hierarchy level for static CPCD. Use the statement at the `[edit dynamic-profiles` *`profile-name`* `services...]` hierarchy level for converged services CPCD.

The `profile`, `rule-set`, and `traceoptions` stanzas are not supported at the `[edit dynamic-profiles` *`profile-name`* `hierarchy level]`.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Options

auto-deactivate *value*

(Optional) Specify one of the following values to determine whether the redirect service is removed automatically when the router receives the subscriber’s initial HTTP-GET message:

- *user-defined-variable*—To use this option in a dynamic profile, you must create a user-defined variable with a name of your choice. The value of the variable can be supplied by the RADIUS server or PCRF. You can also define a default value that is used when the external servers do not supply it. Use the *variables* statement in the dynamic profile to define the default value. Whether supplied by the external server or by the CLI, the value must be either `initial-get` or `never`.

- `initial-get`—Receipt of the subscriber's initial HTTP-GET message triggers removal of the redirect service.

- `never`—Removal of the redirect service is triggered only by the external server, such as by a CoA message from a RADIUS server.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

services—To view this statement in the configuration.

services–control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles` *`profile-name`* `services]` hierarchy level added in Junos OS Release 17.2R1.

Statement introduced in Junos OS Release 10.4.

Support at the `[edit dynamic-profiles` *`profile-name`* `services]` hierarchy level added in Junos OS Release 17.2R1.

Support for Next Gen Services introduced in Junos OS Release 19.3R2.

`auto-deactivate` option added in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

HTTP Redirect Service Overview

Configuring MS-MPC-Based or MX-SPC3-Based Static HTTP Redirect Services

Configuring MS-MPC-Based or MX-SPC3-Based Converged HTTP Redirect Services

Configuring Routing Engine-Based, Static HTTP Redirect Services

Configuring Routing Engine-Based, Converged HTTP Redirect Services

Adding Subscriber Information to HTTP Redirect URLs

stacked-vlan-tagging

IN THIS SECTION

Syntax

```
stacked-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.

For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

6

CHAPTER

Juniper BNG User Plane CLI Operational Commands

IN THIS CHAPTER

Juniper BNG User Plane CLI Operational Commands

IN THIS SECTION

This topic provides an overview of `clear`, `request`, `restart`, and `show` commands, including syntax, option descriptions, and sample output. You use these commands with Juniper BNG User Planes.

restart bbe-cpcdd-dbng

IN THIS SECTION

Syntax

```
restart bbe-cpcdd-dbng
```

Description

Restarts the captive portal content delivery process for the routing engine.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart bbe-cpcdd-dbng

```
user@host> restart bbe-cpcdd-dbng
Captive-portal-content-delivery process for RE started, pid 67476
```

restart bbe-upm-daemon

IN THIS SECTION

Syntax

```
restart bbe-upm-daemon
```

Description

Restarts the BNG User Plane manager daemon.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart bbe-upm-daemon

```
user@host> restart bbe-upm-daemon
Control Plane Manager for dBNG started, pid <process-id>
```

restart gtp-proxy-service

IN THIS SECTION

Syntax

```
restart gtp-proxy-service
```

Description

Restarts the BNG User Plane's GPRS tunneling protocol (GTP) proxy service.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart gtp-proxy-service

```
user@host> restart gtp-proxy-service
User Plane GTP Protocol Proxy started, pid 18832
```

restart pfcp-proxy-service

IN THIS SECTION

Syntax

```
restart pfcp-proxy-service
```

Description

Restarts the BNG User Plane's Packet Forwarding Control Protocol (PFCP) proxy service.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart pfcp-proxy-service

```
user@host> restart pfcp-proxy-service
User Plane PFCP Protocol Proxy started, pid 18840
```

restart subscriber-helper-service

IN THIS SECTION

Syntax

```
restart subscriber-helper-service
```

Description

Restarts the enhanced BBE helper BNG User Plane process. This command runs on the BNG User Plane.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart subscriber-helper-service

```
user@host> restart subscriber-helper-service
Enhanced BBE helper user plane process started, pid 18750
```

restart subscriber-statistics-service

IN THIS SECTION

Syntax

```
restart subscriber-statistics-service
```

Description

Restarts the statistics service daemon for the BNG CUPS Controller. This command runs on the BNG User Plane.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart subscriber-statistics-service

```
user@host> restart subscriber-statistics-service
Statistics Services Daemon for dBNG started, pid 18764
```

restart up-smg-service

IN THIS SECTION

Syntax

```
restart up-smg-service
```

Description

Restarts the enhanced session management BNG User Plane process. This command runs on the BNG User Plane.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart up-smg-service

```
user@host> restart up-smg-service
Enhanced Session Management user plane process started, pid 18585
```

restart upm-service

IN THIS SECTION

Syntax

```
restart upm-service
```

Description

Restarts the BNG User Plane manager. This command runs on the BNG User Plane.

Options

This command does not have any options.

Required Privilege Level

root

Output Fields

When you enter this command, you receive feedback on the status of your request.

Sample Output

restart upm-service

```
user@host> restart upm-service
User Plane Manager for dBNG started, pid 18571
```

show igmp group

IN THIS SECTION

Syntax

```
show igmp group
```

Description

Display Internet Group Management Protocol (IGMP) group membership information. This command runs on BNG User Planes.

Required Privilege Level

view

Output Fields

Table 29 on page 369 describes the output fields for the `show igmp group` command. Output fields are listed in the approximate order in which they appear.

Table 29: show igmp group Output Fields

Field Name	Field Description
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.
Group	Group address.
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.
Source	Source address.
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.
Last reported by	Address of the host that last reported membership in this group.

Table 29: show igmp group Output Fields *(Continued)*

Field Name	Field Description
Timeout	Time remaining until the group membership is removed.
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.

Sample Output

show igmp group

```
user@host> show igmp group
Interface: pp0.3221225481, Groups: 1
    Group: 225.0.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 100.1.1.2
        Timeout:     232 Type: Dynamic
Interface: demux0.2147483652, Groups: 1
    Group: 225.0.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: ROUTE
Interface: local, Groups: 2
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
```

```
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp statistics

IN THIS SECTION

Syntax

```
show igmp statistics
<continuous>
```

Description

Display Internet Group Management Protocol (IGMP) statistics.

By default, Junos OS multicast devices collect statistics of received and transmitted IGMP control messages that reflect currently active multicast group subscribers.

Some devices also automatically maintain *continuous* IGMP statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted IGMP control packets that account for both past and current multicast group subscriptions processed on the device. With continuous statistics, you can see the total count of IGMP control packets the device processed since the last device reboot or `clear igmp statistics continuous` command. The device collects and displays continuous statistics only for the fields shown in the `IGMP packet statistics` output section of this command, and does not display the `IGMP Global statistics` section.

Devices that support continuous statistics maintain this information in a shared database and copy it to the backup Routing Engine at a configurable interval to avoid too much processing overhead on the Routing Engine. These actions preserve statistics counts across the following events or operations (which doesn't happen for the default active subscriber statistics):

- Routing daemon restart
- Graceful Routing Engine switchover (GRES)
- In-service software upgrade (ISSU)
- Line card reboot

You can change the default interval (300 seconds) using the `cont-stats-collection-interval` configuration statement at the `[edit routing-options multicast]` hierarchy level.

You can display either the default currently active subscriber statistics or continuous subscriber statistics (if supported), but not both at the same time. Include the `continuous` option to display continuous statistics, otherwise the command displays the statistics only for active subscribers.

Run the `clear igmp statistics` command to clear the currently active subscriber statistics. On devices that support continuous statistics, run the clear command with the `continuous` option to clear all continuous statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

NOTE: The `show igmp statistics` command runs on BNG User Planes.

Options

none — Display IGMP statistics for all interfaces. These statistics represent currently active subscribers.

brief | detail — (Optional) Display the specified level of output.

continuous — (Optional) Display continuous IGMP statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers.

Required Privilege Level

view

Output Fields

Table 30 on page 373 describes the output fields for the `show igmp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 30: show igmp statistics Output Fields

Field Name	Field Description
`IGMP packet statistics`	Heading for IGMP packet statistics for all interfaces or for the specified interface name. **NOTE**: Shows currently active subscriber statistics in this section by default, or when the command includes the `continuous` option, shows continuous, persistent statistics that account for all IGMP control packets processed on the device.

Table 30: show igmp statistics Output Fields *(Continued)*

Field Name	Field Description
`IGMP Message type`	Summary of IGMP statistics: • `Membership Query`—Number of membership queries sent and received. • `V1 Membership Report`—Number of version 1 membership reports sent and received. • `DVMRP`—Number of DVMRP messages sent or received. • `PIM V1`—Number of PIM version 1 messages sent or received. • `Cisco Trace`—Number of Cisco trace messages sent or received. • `V2 Membership Report`—Number of version 2 membership reports sent or received. • `Group Leave`—Number of group leave messages sent or received. • `Mtrace Response`—Number of Mtrace response messages sent or received. • `Mtrace Request`—Number of Mtrace request messages sent or received. • `Domain Wide Report`—Number of domain-wide reports sent or received. • `V3 Membership Report`—Number of version 3 membership reports sent or received. • `Other Unknown types`—Number of unknown message types received. • `IGMP v3 unsupported type`—Number of messages received with unknown and unsupported IGMP version 3 message types. • `IGMP v3 source required for SSM`—Number of IGMP version 3 messages received that contained no source. • `IGMP v3 mode not applicable for SSM`—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). Beginning with certain releases, this type includes records received for groups in the SSM range of addresses and in which the mode is MODE_IS_EXCLUDE or CHANGE_TO_EXCLUDE_MODE. This includes records with a non-empty source list.
`Received`	Number of messages received.
`Sent`	Number of messages sent.

Table 30: show igmp statistics Output Fields *(Continued)*

Field Name	Field Description
`Rx errors`	Number of received packets that contained errors.
`Max Rx rate (pps)`	Maximum number of IGMP packets received during 1 second interval.
`IGMP Global Statistics`	Summary of IGMP statistics for all interfaces. **NOTE**: These statistics are not supported or displayed with the `continuous` option. • `Bad Length`—Number of messages received with length errors so severe that further classification could not occur. • `Bad Checksum`—Number of messages received with a bad IP checksum. No further classification was performed. • `Bad Receive If`—Number of messages received on an interface not enabled for IGMP. • `Rx non-local`—Number of messages received from senders that are not local. • `Timed out`—Number of groups that timed out as a result of not receiving an explicit leave message. • `Rejected Report`—Number of reports dropped because of the IGMP group policy. • `Total Interfaces`—Number of interfaces configured to support IGMP.

Sample Output

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          2          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
```

```
Group Leave                 0          0          0
Mtrace Response             0          0          0
Mtrace Request              0          0          0
Domain Wide Report          0          0          0
V3 Membership Report        2          0          0
Other Unknown types                               0
IGMP v3 unsupported type                          0
IGMP v3 source required for SSM                   0
IGMP v3 mode not applicable for SSM               0

IGMP Global Statistics
Bad Length                  2
Bad Checksum                0
Bad Receive If           4878
Rx non-local                6
Timed out                   6
Rejected Report             0
Total Interfaces            2
Max Rx rate (pps)          58
```

show igmp statistics continuous

```
user@host> show igmp statistics continuous
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query            0       6932          0
V1 Membership Report        0          0          0
DVMRP                       0          0          0
PIM V1                      0          0          0
Cisco Trace                 0          0          0
V2 Membership Report        0          0          0
Group Leave                 0          0          0
Mtrace Response             0          0          0
Mtrace Request              0          0          0
Domain Wide Report          0          0          0
V3 Membership Report        6          0          0
Other Unknown types                               0
IGMP v3 unsupported type                          0
IGMP v3 source required for SSM                   0
IGMP v3 mode not applicable for SSM               0
```

show mld group

IN THIS SECTION

Syntax

```
show mld group
```

Description

Display information about Multicast Listener Discovery (MLD) group membership. This command runs on BNG User Planes.

Required Privilege Level

view

Output Fields

Table 31 on page 377 describes the output fields for the `show mld group` command. Output fields are listed in the approximate order in which they appear.

Table 31: show mld group Output Fields

Field Name	Field Description
Interface	Name of the interface that received the MLD membership report; **local** means that the local router joined the group itself.

Table 31: show mld group Output Fields *(Continued)*

Field Name	Field Description
Group	Group address.
Source	Source address.
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.
Last reported by	Address of the host that last reported membership in this group.
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.
Timeout	Time remaining until the group membership is removed.
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.

Sample Output

show mld group

```
user@host> show mld group
Interface: pp0.3221225483, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
```

```
        Timeout:     243 Type: Dynamic
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: fe80::e
        Timeout:     249 Type: Dynamic
Interface: demux0.2147483653, Groups: 2
    Group: ff1e::1
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
    Group: ff1e::2
        Group mode: Exclude
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: ROUTE
Interface: local, Groups: 4
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::1:2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff05::1:3
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld statistics

IN THIS SECTION

Syntax

```
show mld statistics
<continuous>
```

Description

Display information about Multicast Listener Discovery (MLD) statistics.

By default, Junos OS multicast devices collect statistics of received and transmitted MLD control messages that reflect currently active multicast group subscribers.

Some devices also automatically maintain *continuous* MLD statistics globally on the device in addition to the default active subscriber statistics—these are persistent, continuous statistics of received and transmitted MLD control packets that account for both past and current multicast group subscriptions processed on the device. With continuous statistics, you can see the total count of MLD control packets the device processed since the last device reboot or `clear mld statistics continuous` command. The device collects and displays continuous statistics only for the fields shown in the `MLD packet statistics...` output section of this command, and does not display the `MLD Global statistics` section.

Devices that support continuous statistics maintain this information in a shared database and copy it to the backup Routing Engine at a configurable interval to avoid too much processing overhead on the Routing Engine. These actions preserve statistics counts across the following events or operations (which doesn't happen for the default active subscriber statistics):

- Routing daemon restart

- Graceful Routing Engine switchover (GRES)
- In-service software upgrade (ISSU)
- Line card reboot

You can change the default interval (300 seconds) using the `cont-stats-collection-interval` configuration statement at the `[edit routing-options multicast]` hierarchy level.

You can display either the default currently active subscriber statistics or continuous subscriber statistics (if supported), but not both at the same time. Include the `continuous` option to display continuous statistics, otherwise the command displays the statistics only for currently active subscribers.

Run the `clear mld statistics` command to clear the currently active subscriber statistics. On devices that support continuous statistics, run the clear command with the `continuous` option to clear all continuous statistics. You must run these commands separately to clear both types of statistics because the device maintains and clears the two types of statistics separately.

NOTE: The `show mld statistics` command runs on BNG User Planes.

Options

none — Display MLD statistics for all interfaces. These statistics represent currently active subscribers.

continuous — (Optional) Display continuous MLD statistics that account for both past and current multicast group subscribers instead of the default statistics that only reflect currently active subscribers. This option is not available with the `interface` option for interface-specific statistics.

Required Privilege Level

view

Output Fields

Table 32 on page 382 describes the output fields for the `show mld statistics` command. Output fields are listed in the approximate order in which they appear.

Table 32: show mld statistics Output Fields

Field Name	Field Description
MLD Packet Statistics...	Heading for MLD packet statistics for all interfaces or for the specified interface name. **NOTE**: Shows currently active subscriber statistics in this section by default, or when the command includes the `continuous` option, shows continuous, persistent statistics that account for all MLD control packets processed on the device.
Received	Number of received packets.
Sent	Number of transmitted packets.
Rx errors	Number of received packets that contained errors.
MLD Message type	Summary of MLD statistics. • **Listener Query (v1/v2)**—Number of membership queries sent and received. • **Listener Report (v1)**—Number of version 1 membership reports sent and received. • **Listener Done (v1/v2)**—Number of Listener Done messages sent and received. • **Listener Report (v2)**—Number of version 2 membership reports sent and received. • **Other Unknown types**—Number of unknown message types received. • **MLD v2 source required for SSM**—Number of MLD version 2 messages received that contained no source. • **MLD v2 mode not applicable for SSM**—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).

Table 32: show mld statistics Output Fields *(Continued)*

Field Name	Field Description
MLD Global Statistics	Summary of MLD statistics for all interfaces. **NOTE**: These statistics are not supported or displayed with the `continuous` option. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with an invalid IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for MLD. • **Rx non-local**—Number of messages received from nonlocal senders. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the MLD group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show mld statistics

```
user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          3          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               7          0          0
Other Unknown types                                      0
MLD v2 unsupported type                                  0
MLD v2 source required for SSM                           0
MLD v2 mode not applicable for SSM                       0

MLD Global Statistics
```

```
Bad Length                    1
Bad Checksum                  0
Bad Receive If               26
Rx non-local                  0
Timed out                     4
Rejected Report               0
Max Rx rate (pps)             4
Total Interfaces              2
```

show mld statistics continuous

```
user@host> show mld statistics continuous
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          5          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               9          0          0
Other Unknown types                                      0
MLD v2 unsupported type                                  0
MLD v2 source required for SSM                           0
MLD v2 mode not applicable for SSM                       0
```

show system subscriber-management arp

IN THIS SECTION

Syntax

```
show system subscriber-management arp
<address ip-address>
<interface interface-name>
```

Description

Displays IPv4 to MAC address bindings (Address Resolution Protocol).

Options

none	Displays the entries in the ARP table.
`address` *`ip-address`*	(Optional) Displays information about ARP for the specified IPv4 address.
`interface` *`interface-name`*	(Optional) Displays information about ARP for the specified interface.

Required Privilege Level

view

Output Fields

Table 33 on page 385 lists the output fields for the `show system subscriber-management arp` command. Output fields are listed in the approximate order in which they appear.

Table 33: show system subscriber-management arp Output Fields

Field Name	Field Description
`MAC Address`	Media access control (MAC) address that corresponds to th
`Address`	IP address that corresponds to the hostname.
`Name`	Host name.

Table 33: show system subscriber-management arp Output Fields *(Continued)*

Field Name	Field Description
interface	Interface name.
Flags	Indicates how mappings between IP and MAC addresses a • permanent—Static mapping. The ARP entry never times • Blank.

Sample Output

show system subscriber-management arp

```
user@host> show system subscriber-management arp
MAC Address       Address         Name                      Interface              Flags
00:00:11:11:11:11 100.16.0.1      100.16.0.1                demux0.3222274078      permanent
00:00:11:11:11:14 100.16.0.2      100.16.0.2                demux0.3222274079      permanent
```

show system subscriber-management arp address *ip-address*

```
user@host> show system subscriber-management arp address 100.16.0.1
MAC Address       Address         Name                      Interface              Flags
00:00:11:11:11:11 100.16.0.1      100.16.0.1                demux0.3222274078      permanent
```

show system subscriber-management arp interface *interface-name*

```
user@host> show system subscriber-management arp interface demux0.3222274078
MAC Address       Address         Name                      Interface              Flags
00:00:11:11:11:11 100.16.0.1      100.16.0.1                demux0.3222274078      permanent
```

show system subscriber-management ipv6-neighbors

IN THIS SECTION

Syntax

```
show system subscriber-management ipv6-neighbors
<address ip-address>
<interface interface-name>
```

Description

Displays IPv6 to MAC address bindings (neighbor cache).

Options

none	Displays the entries in the IPv6 table.
`address` *`ip-address`*	(Optional) Display information about IPv6 for the specified IPv6 address.
`interface` *`interface-name`*	(Optional) Display information about IPv6 for the specified interface.

Required Privilege Level

view

Output Fields

Table 34 on page 388 lists the output fields for the `show system subscriber-management ipv6-neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 34: show system subscriber-management ipv6-neighbors Output Fields

Field Name	Field Description
`IPv6 Address`	The IPv6 address.
`Linklayer Address`	Link-layer address.
`State`	State of the link: • `up` • `down` • `incomplete` • `reachable` • `stale` • `unreachable`
`Exp`	Number of seconds until the entry expires. It will be zero f
`Rtr`	Whether the neighbor is a routing, device (`yes` or `no`).
`Secure`	Whether this entry was created using the Secure Neighbor
`interface`	Interface name.

Sample Output

show system subscriber-management ipv6-neighbors

```
user@host> show system subscriber-management ipv6-neighbors
IPv6 Address                      Linklayer Address  State       Exp   Rtr  Secure  Interface
1000::e3                          00:00:11:11:11:ea  reachable   0     no   no
demux0.3222274295
1000::1e6                         00:00:11:11:12:ec  reachable   0     no   no
demux0.3222274551
```

show system subscriber-management ipv6-neighbors address *ip-address*

```
user@host> show system subscriber-management ipv6-neighbors address 1000::e3
IPv6 Address                      Linklayer Address  State       Exp   Rtr  Secure  Interface
1000::e3                          00:00:11:11:11:ea  reachable   0     no   no
demux0.3222274295
1000::1e6                         00:00:11:11:12:ec  reachable   0     no   no
demux0.3222274551
```

show system subscriber-management ipv6-neighbors interface *interface-name*

```
user@host> show system subscriber-management ipv6-neighbors interface demux0.3222274295
IPv6 Address                       Linklayer Address  State       Exp   Rtr  Secure
Interface
1000::e3                            00:00:11:11:11:ea  reachable   0     no   no
demux0.3222274295
```

show system subscriber-management routing-instance

IN THIS SECTION

Syntax

```
show system subscriber-management routing-instance
<routing-instance-name>
```

Description

Displays the state of the routing instances. The state of the routing instance can be either connected or isolated.

Options

routing-instance-name	(Optional) Displays the state of the specified routing instance.

Required Privilege Level

view

Output Fields

Table 35 on page 391 lists the output fields for the `show system subscriber-management routing-instance` command. Output fields are listed in the approximate order in which they appear.

Table 35: show system subscriber-management routing-instance Output Fields

Field Name	Field Description
`Routing Intance Name`	Name of the routing instance.
`State`	The state of the routing instance. The state can be either **connected** or **isolated**.

Sample Output

show system subscriber-management routing-instance

```
user@host> show system subscriber-management routing-instance
Routing Instance Name                State
default                              connected
```

show system subscriber-management routing-instance

```
user@host> show system subscriber-management routing-instance example-1
Routing Instance Name                State
example-1                            connected
```

show system subscriber-management user-plane

IN THIS SECTION

Syntax

```
show system subscriber-management user-plane
<associations>
<clients>
<control-plane bng-cups-controller-name> <endpoint endpoint-name>
<services> <service service-name>
<routing-instance routing-instance-name>
<upm>
```

Description

Displays statistics for the specified option. This command runs on the BNG User Plane.

Options

`associations`

Displays the BNG CUPS Controller that is associated to the BNG User Plane. Lists information about the BNG CUPS Controller association including the control plane instance name, its address, state, the time the association was formed and the number of times the association was updated.

`clients`

Displays information about each client that interacts with the user plane manager.

`control-plane` *`bng-cups-controller-name`* `endpoint` *`endpoint-name`*

Options:

- *`bng-cups-controller-name`*—Displays information about all endpoints related to the BNG CUPS Controller association.
- `endpoint` *`endpoint-name`*—Displays details about a specific endpoint related to the BNG CUPS Controller association, along with its statistics.

`services`

Displays a list of BNG User Plane services.

Option:

- `service` *`service-name`*—Displays details of the specific service.

`upm` Displays summary information about the user plan manager.

Required Privilege Level

root

Output Fields

Table 36 on page 393 lists the output fields for the `show system subscriber-management user-plane` command.

Table 36: show system subscriber-management user-plane Output Fields

Field Name	Field Description
`User-plane name`	The name of the BNG User Plane.
`Control-plane`	The name, IP address, port, and protocol of the BNG CUPS Controller that the BNG User Plane is associated to.
`Attached clients`	The number of clients attach to the BNG User Plane.
`CP Events`	The number of associations, disassociations, and association updates that have occurred to the associated BNG CUPS Controller.
`CP Name`	Name of the BNG CUPS Controller that the BNG User Plane is associated with.
`Address`	IP address of the BNG CUPS Controller that the BNG User Plane is associated with.
`Assoc Time`	The time that the BNG User Plane started its association with the BNG CUPS Controller.

Table 36: show system subscriber-management user-plane Output Fields *(Continued)*

Field Name	Field Description
`Assoc-updates`	The number of times the association has been updated.
`Client-name`	Name of the BNG User Plane client.
`Status`	The status of the service (either up or down).
`Last Heartbeat Time`	The time when the client's last heartbeat was recorded.
`Services`	The services that information is being reported for.
`State (associations)`	The state of the BNG CUPS Controller that the BNG User Plane is associated to. The state can be the following: • configured • ready • connected • disconnecting • connected-pause

Table 36: show system subscriber-management user-plane Output Fields *(Continued)*

Field Name	Field Description
`State` (`control-plane` *`control-plane-name`*)	The state of the endpoint CUPS Controller that the BNG User Plane is associated to. The state can be the following: • init • configured • ready • connected • timed out • disconnecting • disconnected • warminit • connected-pause • connected-switched
`State` (`services`)	The state of the BNG User Plane services. The state can be the following: • starting • initialized • reconciling • reconciled • ready • terminating • purging-db • fatal-error

Table 36: show system subscriber-management user-plane Output Fields *(Continued)*

Field Name	Field Description
`State` (`services service` *`sevice-name`*)	The state of the endpoint that is managed by the service. The state can be the following: • disabled • disconnected • connected • timed-out • resyncing • synchronized • reconciled • reconciling • connected(backPressureOn)
`Endpoints` (`control-plane`)	The list of endpoints related BG CUPS Controller association service.
`Endpoints` (`services`)	The number of endpoints managed by the service.
`Restarted`	The number of times that the service has been restarted.
`Last Updated`	The last time the service or endpoint state was refreshed by the system.
`Healthy-Endpoints`	The number of healthy endpoints.
`Flapped`	The number of times the endpoint has transitioned from a connected state.

Sample Output

show system subscriber-management user-plane associations

```
user@host show system subscriber-management user-plane associations
CP Name       Address       State            Assoc Time                Assoc-updates
cpi-example    192.0.2.1     Connected         Thu Sep 24 10:44:31 2020  0
```

show system subscriber-management user-plane control-plane *control-plane-name*

```
user@host> show system subscriber-management user-plane control-plane cpi-example
Ip address: 10.216.173.57
Endpoints                Flapped         State                     Last Updated
  Sci                       2          timed out                    0s ago
  Cpri                      11         disconnected                 0s ago
  PfcpProxyStatsIpc         1          connected                    0s ago
  StatsPfcpProxyIpc         0          connected                    0s ago
  PfcpProxySmgIpc           1          connected                    0s ago
  SmgPfcpProxyIpc           0          connected                    0s ago
```

show system subscriber-management user-plane control-plane *control-plane-name* endpoint *endpoint-name*

```
user@host> show system subscriber-management user-plane control-plane cpi-example endpoint Sci
Sci
  Status: timed out
  Rx:                          87
    Invalid packets:           1
    HeartbeatReq                 27
    HeartbeatResp                27
    AssociationSetupReq          4
    AssociationUpdateReq         23
    AssociationUpdateRespAck     1
    NodeReportRespAck            2
    SessionEstablishmentReq      3
  Tx:                          253109
    Avg Round Trip Time:       0
    Transmit errors:           0
```

```
    Max Retransmissions:       0
    Retransmissions:           0
    HeartbeatReq                 253054
    HeartbeatResp                28
    AssociationSetupRespAck      3
    AssociationUpdateReq         1
    AssociationUpdateRespAck     18
    NodeReportReq                2
    SessionEstablishmentRespAck  3
```

show system subscriber-management user-plane services

```
user@host show system subscriber-management user-plane services
Services                           Status        State            Endpoints     Healthy-Endpoints
pfcp-proxy-service                  up           ready              2              2
up-smg-service                      up           ready              5              4
subscriber-helper-service           up           ready              0              0
replication-server-service          up           ready              0              0
replication-client-service          up           ready              0              0
upm-service                         up           ready              0              0
subscriber-statistics-service       up           ready              1              1
gtp-proxy-service                   up           ready              3              3
```

show system subscriber-management user-plane services service *service-name*

```
user@host show system subscriber-management user-plane services service pfcp-proxy-service
pfcp-proxy-service
  Status: up
  State: ready
  Restarted: 1 times
  Last Updated: 27s ago
  Endpoints:Id           Flapped       State            Last Updated
  PfcpProxySmgIpc:1        1         connected            0s ago
  PfcpProxyStatsIpc:0      1         connected            0s ago
```

show system subscriber-management user-plane upm

```
user@host> show system subscriber-management user-plane upm
Start time: Thu Sep 24 10:44:31 2020
User-plane name: vm123
Control-plane:
    Name:            test12
    IP Address:      192.0.2.1
    Port:            8805
    Protocol:        pfcp
Attached clients:    2
CP Events:
    Associations:    1
    Disassociations: 0
    Assoc Updates:   0
```

show user-plane firewall filter

IN THIS SECTION

Syntax

```
show firewall
<counter> counter-name
<filter filter-name>
<log>
```

```
<prefix-action-stats>
<terse>
```

Description

Displays statistics about configured firewall filters on the BNG User Plane.

Options

none	Displays statistics about configured firewall filters.
`counter` *`counter-name`*	(Optional) Name of a filter counter.
`filter` *`filter-name`*	(Optional) Displays firewall filter information for the specified configured filter.
`log`	(Optional) Displays log entries for firewall filters.
`prefix-action-stats`	(Optional) Display prefix action statistics for firewall filters.
`terse`	(Optional) Displays firewall filter names only.

Required Privilege Level

view

Output Fields

Table 37 on page 401 lists the output fields for the `show user-plane firewall filter` command. Output fields are listed in the approximate order in which they appear.

Table 37: show user-plane firewall filter Output Fields

Field Name	Field Description
`Filter`	Name of a filter that has been configured with the `filter` at the `[edit firewall]` hierarchy level. When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either `-i` for an input filter or `-o` for an output filter. When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either `-in` for an input filter or `-out` for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, `__ls1/filter1`).
`Counters`	Display filter counter information: • `Name`—Name of a filter counter that has been configured with the `counter` firewall filter action. • `Bytes`—Number of bytes that match the filter term under which the `counter` action is specified. • `Packets`—Number of packets that matched the filter term under which the `counter` action is specified.
`Policers`	Display policer information: • `Name`—Name of policer. • `Bytes`—Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. • `Packets`—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Sample Output

show user-plane firewall

```
user@host> show user-plane firewall
Filter: in-filter-2-et-0/0/0.99-i
Counters:
Name                                              Bytes              Packets
count-2-et-0/0/0.99-I                                 0                    0
Policers:
Name                                              Bytes              Packets
upstream-policer-term-1-et-0/0/0.99-i                 0                    0

Filter: in-filter_UID1075-pp0.3221225500-in
Counters:
Name                                              Bytes              Packets
P0-Aggregate-count-pp0.3221225500-in        25669737936             46757264
P1-Aggregate-count-pp0.3221225500-in        50120712891             91294559
P2-Aggregate-count-pp0.3221225500-in        50120448273             91294077
final-count-pp0.3221225500-in               50120511408             91294192

Filter: out-filter_UID1077-pp0.3221225500-out
Counters:
Name                                              Bytes              Packets
out-filter-count-pp0.3221225500-out                   0                    0
                                                      0                    0
```

show user-plane firewall filter

```
user@host> show user-plane firewall filter in-filter_UID1075-pp0.3221225500-in
Filter: in-filter_UID1075-pp0.3221225500-in
Counters:
Name                                              Bytes              Packets
P0-Aggregate-count-pp0.3221225500-in        25675095078             46767022
P1-Aggregate-count-pp0.3221225500-in        50131401372             91314028
P2-Aggregate-count-pp0.3221225500-in        50131463958             91314142
final-count-pp0.3221225500-in
```

show user-plane firewall filter *filter-name* counter

```
user@host> show user-plane firewall filter in-filter_UID1075-pp0.3221225500-in counter P0-
Aggregate-count-pp0.3221225500-in
Filter: in-filter_UID1075-pp0.3221225500-in
Counters:
Name                                                Bytes              Packets
P0-Aggregate-count-pp0.3221225500-in          25688770119             46791931
```

show user-plane firewall terse

```
user@host> show user-plane firewall terse

Filter: in-filter-2-et-0/0/0.99-i

Filter: in-filter_UID1075-pp0.3221225500-in

Filter: out-filter_UID1077-pp0.3221225500-out
```

show user-plane firewall templates-in-use

IN THIS SECTION

Syntax

```
show user-plane firewall templates-in-use up-name user-plane-name
```

Description

Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced.

Options

user-plane-name Display the configured filter templates for the specified BNG User Plane.

Required Privilege Level

view

Output Fields

Table 38 on page 404 lists the output fields for the `show user-plane firewall templates-in-use` command. Output fields are listed in the approximate order in which they appear.

Table 38: show user-plane firewall templates-in-use Output Fields

Field Name	Field Description
`Filter Template`	Name of a filter that has been configured using the `filter` statement at either the `[edit firewall]` or `[edit dynamic-profiles` *profile-name* `firewall]` hierarchy and is being used as a template for dynamic subscriber filtering.
`User Plane`	BNG User Plane name.
`Reference Count`	Number of times the filter has been referenced by subscribers accessing the network.

Sample Output

show user-plane firewall templates-in-use up-name up-example-1

```
user@host> show user-plane firewall templates-in-use up-name up-example-1
                         Dynamic Subscribers Reference Counts
Filter Template                 User Plane         Reference Count
----------------                ----------         ----------------
static-filter                   up-example-1       2
dynamic-filter_UID1080          up-example-1       2
```

show system subscriber-management redundancy-state interface

IN THIS SECTION

Syntax

```
show system subscriber-management redundancy-state interface interface-name
```

Description

Displays the status of the subscriber management redundancy service of the pseudowire interface.

Options

interface *interface-name*	Name of the pseudowire interface.

Required Privilege Level

view

Output Fields

Table 39 on page 406 lists the output fields for the `show system subscriber-management redundancy-state interface` *`interface-name`* command.

Table 39: show system subscriber-management redundancy-state interface Output Fields

Field Name	Field Description
Forwarding state	Status of the packet forwarding.
Service Activation Programming	Status of the service activation programming.
Standby-mode	Status of the standby mode.

Sample Output

Primary BNG Interface Status

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status in Normal Operation

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Secondary BNG Interface Status Immediately on Failover

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: In-progress
```

Secondary BNG Interface Status on Failover Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: hot-standby
Forwarding state: Active
Service Activation Programming: Completed
```

Secondary BNG Interface Status Immediately on Failover Reversal

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: In-rollback
```

Secondary BNG Interface Status Immediately on Failover Reversal Completion

```
user@host> show system subscriber-management redundancy-state interface ps1.0
```

```
Interface: ps1.0
Standby-mode: service-activation-on-failover
Forwarding state: Inactive
Service Activation Programming: Not-applicable
```

Release Information

Statement introduced in Junos OS Release 22.2R1.

show system subscriber-management resiliency

IN THIS SECTION

Syntax

```
show system subscriber-management resiliency
<detail>
<extensive>
<summary>
```

Description

Display information that indicates the health and relationship of session database replication between the primary and standby Routing Engines.

Options

detail (Optional) Displays brief information about the shared memory state for the primary and standby Routing Engines.

extensive (Optional) Displays very detailed statistics for the SDB components in shared memory for the primary and standby Routing Engines, enabling you to evaluate the state of replication between the two.

summary (Optional) Displays only an indication of whether the system is okay (replication is normal) or has some unexpected condition.

Required Privilege Level

system

Output Fields

Table 40 on page 410 lists the output fields for the `show system subscriber-management resiliency` command. Output fields are listed in the approximate order in which they appear.

Table 40: show system subscriber-management resiliency Output

Field Name	Field Description	Level
`Overall Status`	Indicates the condition of the system: • `Ok`—The system is functioning normally. • `Not-Ok`—An unexpected condition has been discovered. This status may require investigation by the Juniper Networks Technical Assistance Center (JTAC) to confirm whether anything is wrong and the root cause of the status.	`summary`

Table 40: show system subscriber-management resiliency Output *(Continued)*

Field Name	Field Description	Level
shared memory type	One of the following types of shared memory objects: • mmap—Memory-mapped file that stores the hash or entry data for an MMDB. • mmap Database (MMDB)–Memory-mapped database that uses memory-mapped files to store the MMDB hash and entry data. Each MMDB typically stores a type of statistic, such as statistics related logical interfaces, logical interface sets, or subscribers. • Shared Memory Segment—An operating system object that is a chunk of contiguous shared memory. Total—Number of memory objects of all types.	detail
count	Number of shared memory instances of a type.	detail
mapped bytes	Number of bytes mapped into process space.	detail
mmfs	Memory-mapped file information.	extensive
Name	File path including the filename of the shared memory object. For MMFs, the filename is the name of its associated MMDB and a suffix to indicate whether it stores hash or data. For MMDBs, the filename indicates the type of statistics stored in the database.	extensive
Current Bytes	Current total size of the shared memory object.	extensive
Maximum Bytes	Maximum size of the shared memory object.	extensive
Mapped Bytes	Number of bytes mapped into process space.	extensive
Lock Count	Number of times the shared memory object has been locked by a global, inter-process lock.	extensive

Table 40: show system subscriber-management resiliency Output *(Continued)*

Field Name	Field Description	Level
`Contention Count`	Number of times that a process or thread object waited to lock a shared memory object because a different process or thread already has the lock. This is a global, inter-process lock.	`extensive`
`Lock Wait Secs`	How long a process or thread taking a global, inter-process lock waited because a different process or thread already had the lock.	`extensive`
`mmap Count`	Number of times that parts of the overall memory mapped data have been mapped.	`extensive`
`Shared Memory Segments`	Information about the shared memory segments; each segment is a chunk of contiguous shared memory.	`extensive`
`Size in Bytes`	Number of bytes in the shared memory segment.	`extensive`
`MMDBs`	Information about the memory-mapped file databases that use memory-mapped files to store data (typically statistics associated with interfaces and subscribers).	`extensive`
`Hash Entries`	Number of different hash entries a key could be hashed to in this table.	`extensive`
`PLock Count`	Number of times the MMDB shared memory object has been locked by a process-level, intra-process lock.	`extensive`
`PLock Contention Count`	Number of times that a process or thread object waited to lock a shared memory object because a different process or thread already has the lock. This is a process-level, intra-process lock.	`extensive`
`PLock Wait Secs`	How long a process or thread taking a process-level, intra-process lock waited because a different process or thread already had the lock.	`extensive`

Sample Output

show system subscriber-management resiliency (Summary)

```
user@host> show system subscriber-management resiliency summary
Overall Status: Ok
```

show system subscriber-management resiliency (Detail)

```
user@host> show system subscriber-management resiliency detail
Master:
shared memory type            count mapped bytes
mmap                             43    195027200
mmap Database (MMDB)              9    (in mmap)
Shared Memory Segment             6     39163504
Total                            58    234190704

Standby:
shared memory type            count mapped bytes
mmap                             41    192930048
mmap Database (MMDB)              9    (in mmap)
Shared Memory Segment             6     39163504
Total                            56    232093552
```

show system subscriber-management resiliency (Extensive)

```
user@host> show system subscriber-management resiliency extensive
Master:


   mmfs:
      Name                                                    Current Bytes Maximum Bytes
Mapped bytes   Lock Count Lock Contention Count   Lock Wait Secs   mmap Count
      /mfs/var/smm_accounting-stats-db_hash                        15736832
15736832     15736832          17                     0         0.000000            0
      /mfs/var/smm_accounting-stats-db_data                      1139015680
9112125440      2097152          17                     0         0.000000           18
      /mfs/var/mmcq/mmdb_rep_mmcq                                   1048576
104857600      1048576          25                     1         0.011021            0
```

```
     /mfs/var/smm_accounting-ifl-db_hash                                  28672
28672        28672         17                     0        0.000000          0
     /mfs/var/smm_accounting-ifl-db_data                                  33554432
536870912      4194304          17                    0        0.000000          18
     /mfs/var/smm_accounting-iflset-db_hash                               28672
28672        28672         17                     0        0.000000          0
     /mfs/var/smm_accounting-iflset-db_data                               33554432
536870912      4194304          17                    0        0.000000          18
     /mfs/var/sdb/shmem/sdb.head                                          7680256
7680256      7680256      384006                   0        0.000000          0
     /mfs/var/sdb/shmem/sdb.lts.data                                      1620049920
8589934592     20971520          41                     0         0.000000          60
     /mfs/var/sdb/shmem/sdb_sts_data                                      51216384
51216384     51216384        20012                    0        0.000000          0
     /mfs/var/sdb/shmem/sdb_intf.db                                       409600
409600       409600           0                      0        0.000000          0
     /mfs/var/sdb/shmem/subscriber_hash                                   2408448
2408448      2408448          21                     0        0.000000          0
     /mfs/var/sdb/shmem/subscriber_data                                   33554432
536870912      2097152          22                    0        0.000000          22
     /mfs/var/sdb/shmem/service_hash                                      2408448
2408448      2408448          21                     0        0.000000          0
     /mfs/var/sdb/shmem/service_data                                      33554432
536870912      2097152          22                    0        0.000000          22
     /mfs/var/sdb/shmem/interface_hash                                    28672
28672        28672          21                     0        0.000000          0
     /mfs/var/sdb/shmem/interface_data                                    33554432
536870912      4194304         109                    0        0.000000          22
     /mfs/var/sdb/shmem/interface_set_hash                                28672
28672        28672          21                     0        0.000000          0
     /mfs/var/sdb/shmem/interface_set_data                                33554432
536870912      4194304          22                    0        0.000000          22
     /mfs/var/sdb/shmem/mobile_subs_location_hash                         1208320
1208320      1208320          21                     0        0.000000          0
     /mfs/var/sdb/shmem/mobile_subs_location_data                         33554432
536870912      2097152          22                    0        0.000000          22
     /mfs/var/sdb/shmem/mobile_subscriber_hash                            1208320
1208320      1208320          21                     0        0.000000          0
     /mfs/var/sdb/shmem/mobile_subscriber_data                            33554432
536870912      2097152          21                    0        0.000000          22
     /mfs/var/mmq/mmq_queue                                               126976
126976       126976           5                      0        0.000000          0
     /mfs/var/mmq/mmq_heap                                                5120000
```

```
5120000      5120000           4                     0        0.000000           0
      /mfs/var/mmcq/sdb_bbe_mmcq                                        25165824
318767104     25165824          21                    0        0.000000           0
      /mfs/var/mmcq/authdRxQueue                                         1048576
20971520      1048576           6                    0        0.000000           0
      /mfs/var/mmcq/pppdRxQueue                                          1048576
20971520      1048576           4                    0        0.000000           0
      /mfs/var/mmcq/bbeStatsdGetCollector                                1048576
20971520      1048576          16                    0        0.000000           0
      /mfs/var/mmdb/mmdb_ack_registry                                       8192
8192         8192          141                    0        0.000000           0
      /mfs/var/mmcq/mmdb_ackq_bbe-statsd                                 1048576
67108864      1048576           2                    0        0.000000           0
      /mfs/var/mmcq/jdchpdAccountingClientApp                            1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/ss/domain.0.data                                         16777216
2147483648      4194304          262                    0        0.000000          18
      /mfs/var/tmp/bbe_throttle_control                                     8192
8192         8192            7                    0        0.000000           0
      /mfs/var/mmcq/statsPluginGCClient                                  1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/sdb/shmem/sdb_reg_info                                       8192
8192         8192            2                    0        0.000000           0
      /mfs/var/mmcq/sdb_reg_q_bbe-statsd                                16777216
16777216     16777216           2                    0        0.000000           0
      /mfs/var/mmcq/jl2tpdCliRxQ                                         1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/mmcq/jl2tpdSnmpRxQ                                        1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/mmcq/authd                                                1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/mmcq/jpppdAccountingClientApp                             1048576
20971520      1048576           2                    0        0.000000           0
      /mfs/var/mmcq/mmdb_mmcq_0                                          1048576
104857600      1048576           42                    0        0.000000           0
      /mfs/var/ss/domain.0                                                409600
4294967295       409600      6400000                 3037         0.002642           0

   Shared Memory Segments:
      Name                                                          Size in Bytes
      /mfs/var/shmlog/shmlog                                            39071744
      sdb_rsmon_shared_memory                                              22536
      sdb_rsmon_ae_table                                                    4096
```

```
      sdb_rsmon_ps_table                                              60008
      sdb_rsmon_rlt_table                                              1024
      sdb_bbe_rep_mailbox                                              4096

   MMDBs:
      Name                                                        Hash Entries   Lock Count
Lock Contention Count   Lock Wait Secs   PLock Count PLock Contention Count   PLock Wait Secs
      /mfs/var/smm_accounting-stats-db                                 655360
7208990                    0         0.000000       1966111                      0
0.000000
      /mfs/var/smm_accounting-ifl-db                                     1000
11024                    0         0.000000         3025                       0
0.000000
      /mfs/var/smm_accounting-iflset-db                                  1000
11024                    0         0.000000         3025                       0
0.000000
      /mfs/var/sdb/shmem/subscriber                                    100000
1400010                    2         0.043705       400012                       0
0.000000
      /mfs/var/sdb/shmem/service                                       100000
1400010                    0         0.000000       400012                       0
0.000000
      /mfs/var/sdb/shmem/interface                                       1000
14430                    0         0.000000         4427                       0
0.000000
      /mfs/var/sdb/shmem/interface_set                                   1000
14010                    0         0.000000         4012                       0
0.000000
      /mfs/var/sdb/shmem/mobile_subs_location                           50000
700018                    0         0.000000       200020                      0
0.000000
      /mfs/var/sdb/shmem/mobile_subscriber                              50000
700010                    0         0.000000       200012                      0
0.000000

   Total Mapped Bytes                                         234190704

Standby:

   mmfs:
      Name                                                      Current Bytes Maxiumum Bytes
Mapped bytes   Lock Count Lock Contention Count   Lock Wait Secs   mmap Count
      /mfs/var/smm_accounting-stats-db_hash                           15736832
```

```
15736832     15736832          13                     0        0.000000         0
      /mfs/var/smm_accounting-stats-db_data                          1139015680
9112125440     2097152          13                     0        0.000000         14
      /mfs/var/mmcq/mmdb_rep_mmcq                                       1048576
104857600      1048576          15                     0        0.000000         0
      /mfs/var/smm_accounting-ifl-db_hash                                 28672
28672        28672          13                     0        0.000000         0
      /mfs/var/smm_accounting-ifl-db_data                              33554432
536870912      4194304          13                     0        0.000000         14
      /mfs/var/smm_accounting-iflset-db_hash                              28672
28672        28672          13                     0        0.000000         0
      /mfs/var/smm_accounting-iflset-db_data                           33554432
536870912      4194304          13                     0        0.000000         14
      /mfs/var/sdb/shmem/sdb.head                                       7680256
7680256      7680256       384005                     0        0.000000         0
      /mfs/var/sdb/shmem/sdb.lts.data                                1620049920
8589934592     20971520          11                     0        0.000000         20
      /mfs/var/sdb/shmem/sdb_sts_data                                  51216384
51216384     51216384        17510                     0        0.000000         0
      /mfs/var/sdb/shmem/sdb_intf.db                                     409600
409600       409600            0                     0        0.000000         0
      /mfs/var/sdb/shmem/subscriber_hash                                2408448
2408448      2408448            5                     0        0.000000         0
      /mfs/var/sdb/shmem/subscriber_data                               33554432
536870912      2097152            4                     0        0.000000         6
      /mfs/var/sdb/shmem/service_hash                                   2408448
2408448      2408448            5                     0        0.000000         0
      /mfs/var/sdb/shmem/service_data                                  33554432
536870912      2097152            4                     0        0.000000         6
      /mfs/var/sdb/shmem/interface_hash                                   28672
28672        28672            5                     0        0.000000         0
      /mfs/var/sdb/shmem/interface_data                                33554432
536870912      4194304            4                     0        0.000000         6
      /mfs/var/sdb/shmem/interface_set_hash                               28672
28672        28672            5                     0        0.000000         0
      /mfs/var/sdb/shmem/interface_set_data                            33554432
536870912      4194304            4                     0        0.000000         6
      /mfs/var/sdb/shmem/mobile_subs_location_hash                      1208320
1208320      1208320            5                     0        0.000000         0
      /mfs/var/sdb/shmem/mobile_subs_location_data                     33554432
536870912      2097152            4                     0        0.000000         6
      /mfs/var/sdb/shmem/mobile_subscriber_hash                         1208320
1208320      1208320            5                     0        0.000000         0
```

```
      /mfs/var/sdb/shmem/mobile_subscriber_data                               33554432
536870912     2097152          4                     0         0.000000           6
      /mfs/var/mmq/mmq_queue                                                    126976
126976      126976          4                     0         0.000000           0
      /mfs/var/mmq/mmq_heap                                                    5120000
5120000     5120000          3                     0         0.000000           0
      /mfs/var/mmcq/sdb_bbe_mmcq                                              25165824
318767104    25165824         11                     0         0.000000           0
      /mfs/var/mmcq/authdRxQueue                                               1048576
20971520     1048576          6                     0         0.000000           0
      /mfs/var/mmcq/pppdRxQueue                                                1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/mmcq/bbeStatsdGetCollector                                      1048576
20971520     1048576         14                     0         0.000000           0
      /mfs/var/mmdb/mmdb_ack_registry                                             8192
8192         8192           2                     0         0.000000           0
      /mfs/var/mmcq/mmdb_ackq_bbe-statsd                                       1048576
67108864     1048576          2                     0         0.000000           0
      /mfs/var/mmcq/jdchpdAccountingClientApp                                  1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/ss/domain.0.data                                               16777216
2147483648      4194304        261                     0         0.000000          16
      /mfs/var/tmp/bbe_throttle_control                                           8192
8192         8192           6                     0         0.000000           0
      /mfs/var/mmcq/statsPluginGCClient                                        1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/sdb/shmem/sdb_reg_info                                             8192
8192         8192           2                     0         0.000000           0
      /mfs/var/mmcq/sdb_reg_q_bbe-statsd                                      16777216
16777216    16777216          2                     0         0.000000           0
      /mfs/var/mmcq/jl2tpdCliRxQ                                               1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/mmcq/jl2tpdSnmpRxQ                                              1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/mmcq/authd                                                      1048576
20971520     1048576          2                     0         0.000000           0
      /mfs/var/ss/domain.0                                                      409600
4294967295       409600      8000000                  4044         0.002962           0

   Shared Memory Segments:
      Name                                                              Size in Bytes
      /mfs/var/shmlog/shmlog                                                  39071744
      sdb_rsmon_shared_memory                                                    22536
```

```
      sdb_rsmon_ae_table                                                      4096
      sdb_rsmon_ps_table                                                     60008
      sdb_rsmon_rlt_table                                                     1024
      sdb_bbe_rep_mailbox                                                     4096

   MMDBs:
      Name                                                            Hash Entries   Lock Count
Lock Contention Count   Lock Wait Secs   PLock Count PLock Contention Count   PLock Wait Secs
      /mfs/var/smm_accounting-stats-db                                       655360
5898264                    0         0.000000       1966105                      0
0.000000
      /mfs/var/smm_accounting-ifl-db                                           1000
9020                    0         0.000000          3021                      0
0.000000
      /mfs/var/smm_accounting-iflset-db                                        1000
9020                    0         0.000000          3021                      0
0.000000
      /mfs/var/sdb/shmem/subscriber                                          100000
300002                    0         0.000000       100003                      0
0.000000
      /mfs/var/sdb/shmem/service                                             100000
300002                    0         0.000000       100003                      0
0.000000
      /mfs/var/sdb/shmem/interface                                             1000
3002                    0         0.000000          1003                      0
0.000000
      /mfs/var/sdb/shmem/interface_set                                         1000
3002                    0         0.000000          1003                      0
0.000000
      /mfs/var/sdb/shmem/mobile_subs_location                                 50000
150002                    0         0.000000        50003                      0
0.000000
      /mfs/var/sdb/shmem/mobile_subscriber                                    50000
150002                    0         0.000000        50003                      0
0.000000

   Total Mapped Bytes                                                 232093552
```

Release Information

Command introduced in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

Junos OS Enhanced Subscriber Management

show system subscriber-management route

IN THIS SECTION

Syntax

```
show system subscriber-management route
<family family>
<incomplete>
<level (brief | detail)>
<next-hop index>
<prefix>
<routing-instance name>>
<route-type type>
```

```
<rrt-index index>
<summary>
```

Description

Display information about how routes are mapped to specific enhanced subscriber management interfaces. You can customize the output by including one or more optional filters in the command. With the exception of the `summary` option, all filter options can be combined in a single command.

Options

family *family*	(Optional) Display route mapping information for the specified protocol family: `inet` (IPv4) or `inet6` (IPv6).
incomplete	(Optional) Display route mapping information for incomplete routes that are missing elements required to add the routes to the routing table.
level (brief | detail)	(Optional) Display the specified level of output: `brief` or `detail`.
next-hop *index*	(Optional) Display the next hop associated with the route entry with the specified next-hop index, in the range 1 through 65535.
prefix *address*	(Optional) Use the same prefix and prefix length as the subscriber host address. Output includes attributes that originate in the Famed-Route record of an upstream RADIUS server (Tag, Metric, Preference).
route-type *type*	(Optional) Display route mapping information for the specified route type: `access`, `access-internal`, `kernel`, or `local`.
routing-instance *name*	(Optional) Display route mapping information for the specified routing-instance
rrt-index *index*	(Optional) Display mapping information for the specified routing table index, in the range 0 through 65535. An `rtt-index` value of 0 (zero) denotes routes in the default routing table managed by enhanced subscriber management.
summary	(Optional) Display summary information about the routes managed by enhanced subscriber management.

Required Privilege Level

view

Output Fields

Table 41 on page 422 lists the output fields for the `show system subscriber-management route` command. Output fields are listed in the approximate order in which they appear.

Table 41: show system subscriber-management route Output Fields

Field Name	Field Description	Level of Output
address	IPv4 or IPv6 address associated with the route entry.	All levels
`Route Type`	One of the following route types: • `Access` • `Access-internal` • `Framed` • `Kernel` • `Local`	All levels
`Interface`	Name of the enhanced subscriber management interface associated with the route entry.	All levels
`Next-hop`	Next-hop associated with the route entry.	All levels
`Tag`	Reflects the Tag attribute used in the RADIUS Framed-Route type record.	All levels
`Metric`	Reflects the Metric attribute used in the RADIUS Framed-Route type record.	All levels

Table 41: show system subscriber-management route Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Preference`	Reflects the Preference attribute used in the RADIUS Framed-Route type record.	All levels
`Rtt-index`	Value of the routing table index. A value of `0` (zero) denotes a route in the default routing table managed by enhanced subscriber management.	`detail`
`Bbe index`	Value of the interface index for the control plane.	`detail`
`Flow id`	Value of the route object index.	`detail`
`Reference Count`	Used for internal accounting.	`detail`
`Dirty Flags`	Used for internal accounting.	`detail`
`Flags`	Used for internal accounting.	`detail`
`Family`	One of the following protocol families: • `AF_INET`—IPv4 • `AF_INET6`—IPv6	`detail`

Sample Output

show system subscriber-management route prefix <address>

rtt-index 0

```
user@host> show system subscriber-management route prefix 10.10.0.1/32
Route:  10.10.0.1/32
     Routing-instance:         default:default
```

```
    Kernel rt-table id :      0
    Family:                   AF_INET
    Route Type:               Framed
    Protocol Type:            Unspecified
    Interface:                pp0.3221225491
    Interface index:          26
    Internal Interface index: 26
    Route index:              20
    Next-Hop:                 684
    Tag:                      9999
    Metric:                   56
    Preference:               10
    Reference-count:          1
    L2 Address:               00:00:5e:00:53:0b
    Flags:                    0x0
    Dirty Flags:              0x0
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
    Route Type: Access
    Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
    Route Type: Access
    Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
    Route Type: Access
    Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
    Route Type: Access
    Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
    Route Type: Access
    Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
    Route Type: Access
```

```
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

show system subscriber-management route family route-type rtt-index level detail

The following example displays detailed information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level detail
2001:db8::/64
     Route Type:      Access
     Interface:       pp0.3221225479
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:1::/64
     Route Type:      Access
     Interface:       pp0.3221225477
     Next-hop:        721
     Rtt-index:       0
     Bbe index:       9
     Flow id:         1
     Reference Count: 1
     Dirty Flags:     0
     Flags:           0x10082
     Family:          AF_INET6
2001:db8:0:0:2::/64
     Route Type:      Access
     Interface:       pp0.3221225478
     Next-hop:        721
     Rtt-index:       0
```

```
    Bbe index:       9
    Flow id:         1
    Reference Count: 1
    Dirty Flags:     0
    Flags:           0x10082
    Family:          AF_INET6
2001:db8:0:0:3::/64
    Route Type:      Access
    Interface:       pp0.3221225480
    Next-hop:        721
    Rtt-index:       0
    Bbe index:       9
    Flow id:         1
    Reference Count: 1
    Dirty Flags:     0
    Flags:           0x10082
    Family:          AF_INET6
```

show system subscriber-management route family route-type rtt-index level brief

The following example displays abbreviated information about IPv6 access routes in the default routing table (`rtt-index 0`) managed by enhanced subscriber management.

```
user@host> show system subscriber-management route family inet6 route-type access rtt-index 0
level brief
2001:db8::/64
    Route Type: Access
    Interface: pp0.3221225479, Next-hop:721
2001:db8:0:0:1::/64
    Route Type: Access
    Interface: pp0.3221225477, Next-hop:721
2001:db8:0:0:2::/64
    Route Type: Access
    Interface: pp0.3221225478, Next-hop:721
2001:db8:0:0:3::/64
    Route Type: Access
    Interface: pp0.3221225480, Next-hop:721
2001:db8:0:0:4::/64
    Route Type: Access
    Interface: pp0.3221225481, Next-hop:721
2001:db8:2002::/84
```

```
     Route Type: Access
     Interface: demux0.3221225492, Next-hop:721
2001:db8:0:0:5::/64
     Route Type: Access
     Interface: pp0.3221225487, Next-hop:721
2001:db8:0:0:6::/64
     Route Type: Access
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Support for passing **Framed-Route** attributes from a RADIUS server to the router was added in Junos OS Release 17.2 on MX Series routers for enhanced subscriber management. This allows the tagged subscriber host routes to be imported to the routing table and advertised by BGP.

RELATED DOCUMENTATION

Verifying and Managing Junos OS Enhanced Subscriber Management

show system subscriber-management statistics

IN THIS SECTION

Syntax

```
show system subscriber-management statistics
<all>
<dhcp>
<dvlan>
<fixed-wireless-access>
<l2tp>
<ppp>
<pppoe>
```

Description

Display statistics for the specified option. You can customize the output by including one or more optional filters in the command. With the exception of the `extensive` option, all filter options can be combined in a single command.

Options

all	(Optional) Display packet statistics for all protocols.
dhcp	(Optional) Display DHCP packet statistics.
dvlan	(Optional) Display DVLAN packet statistics.
fixed-wireless-access	(Optional) Display fixed wireless access packet statistics.
l2tp	(Optional) Display L2TP packet statistics.
ppp	(Optional) Display PPP packet statistics.
pppoe	(Optional) Display PPPoE packet statistics.

Required Privilege Level

view

Output Fields

Table 42 on page 429 lists the output fields for the `show system subscriber-management statistics` command. Output fields are listed in the approximate order in which they appear.

Table 42: show system subscriber-management statistics Output Fields

Field Name	Field Description
`Rx Statistics`	Statistics for packets received.
`Tx Statistics`	Statistics for packets sent.
`Enhanced I/O Statistics`	Statistics for visibility into packet drops from the queue.
`Error Statistics`	Includes connection packets, flow control, and messages and packets sent to and received from the daemon.
`ERA discards`	Event Rate Analyzer discards. For DHCP and PPPoE in advanced subscriber management, ERA packet discard counts are included for Discover, Solicit, and PADI packets .
`Layer 3 Statistics`	Statistics for Layer 3 packets.
`padis`	PPPoE Active Discovery Initiation (PADI) packets. PADI is the first step in the PPPoE establishment protocol.
`padrs`	PPPoE Active Discovery Request packets.
`ppp`	Point-to-Point Protocol packets.

Table 42: show system subscriber-management statistics Output Fields *(Continued)*

Field Name	Field Description
`router solicitations`	Number of router solicitations sent or received. Router solicitations are sent to prompt all on-link routers to send it router advertisements.
`router advertisements`	Number of router advertisements sent or received.
`route solicit response packet`	Number of router solicitation responses sent or received.

Sample Output

The following examples displays packet statistics accumulated for DHCP, hybrid access, and PPPoE since the last time the session manager was cleared.

show system subscriber-management statistics all

```
user@host> show system subscriber-management statistics all
  user@host> show system subscriber-management statistics all
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
-------------------------------------------------------------
                     Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660
```

```
DHCP Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                          : 320008
        ERA discards                     : 6274
    Tx Statistics
        transmit request packets         : 320482
        sent packets                     : 320482
    Error Statistics
    Connection Statistics
        no connection packets            : 0

PPPoE Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                          : 486165
        padis                            : 36768
        padrs                            : 35421
        ppp packets                      : 341787
        ERA discards                     : 8249
    Tx Statistics
        packets                          : 70842
        send failures                    : 6240
```

show system subscriber-management statistics dhcp

```
user@host> show system subscriber-management statistics dhcp
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
--------------------------------------------------------------
                     Packet Statistics
--------------------------------------------------------------
I/O Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
  Layer 3 Statistics
    Rx Statistics
```

```
        packets                         : 356218
    Tx Statistics
        packets                         : 6604660

DHCP Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                         : 320008
        ERA discards                    : 6274
    Tx Statistics
        transmit request packets        : 320482
        sent packets                    : 320482
    Error Statistics
    Connection Statistics
        no connection packets           : 0
```

show system subscriber-management statistics dhcp extensive

```
user@host> show system subscriber-management statistics dhcp extensive
Session Manager started @ Mon Dec  6 06:14:27 2021
Session Manager cleared @ Mon Dec  6 06:14:27 2021
-------------------------------------------------------------
                     Packet Statistics
-------------------------------------------------------------
I/O Statistics:
-------------------------------------------------------------
    Rx Statistics
        packets                         : 7986
    Tx Statistics
        packets                         : 15
          l2 inject                     : 15
          l3 inject                     : 0
    Buffer Statistics
        allocations                     : 7990
        frees                           : 7990
        allocation failures             : 0
  Layer 3 Statistics
    Rx Statistics
        packets                         : 29
    Tx Statistics
        packets                         : 15
```

```
    PFE Event Statistics
        packets                          : 7957
--------------------------------------------------------------
Enhanced I/O Statistics:
--------------------------------------------------------------
    bbe_io_rcv l2                        : 0
    bbe_io_rcv l3                        : 29
    bbe_io_rcv l3 v4                     : 17
    bbe_io_rcv l3 v6                     : 12
    bbe_io_rcv l3 unspec                 : 0
    bbe_io_rcv l3 unknown af             : 0
    bbe_io_rcv routed                    : 0
    bbe_io_rcv routed_v4                 : 0
    bbe_io_rcv routed_v6                 : 0
    bbe_io_rcv routed_no_route           : 0
    bbe_io_rcv routed_default            : 0
    bbe_io_rcv resolve v4                : 0
    bbe_io_rcv resolve v6                : 0
    bbe_io_rcv resolve default           : 0
    bbe_io_rcv pfe event                 : 0
    bbe_io_rcv default                   : 0
    ----------------------------------------------------------
    rx inet                              : 17
    rx inet6                             : 12
    rx inet6 icmp6                       : 2
    rx inet igmp                         : 0
    rx inet gre                          : 0
    rx inet agf                          : 0
    rx inet icmp                         : 0
    rx inet udp v4                       : 17
    rx inet udp v6                       : 10
          rx inet udp l2tp               : 0
          rx inet udp jdhcp v4           : 17
          rx inet udp jdhcp v6           : 10
          rx inet udp bfd v4             : 0
          rx inet udp bfd v6             : 0
    rx inet tcp proxy v4                 : 0
    rx inet tcp proxy v6                 : 0
    rx v4 ip frag reassm proc cnt        : 0
    rx v4 ip frag reassm pkt cnt         : 0
    rx v4 ip frag reassm alloc cntx cnt  : 0
    rx v4 ip frag reassm free cntx cnt   : 0
    tx l3 forward                        : 15
```

```
    tx udp v4                           : 5
    tx udp v6                           : 6
    tx tcp proxy client v4              : 0
    tx tcp proxy client v4 error        : 0
    tx tcp proxy v4 drop                : 0
    tx tcp proxy v6 drop                : 0
    tx igmp                             : 0
    tx v6 l3 route forwards             : 0
    tx v6 l3 route drops                : 0
    tx v6 kernel forwards               : 0
    tx v6 kernel forward drops          : 0
    tx v6 l2 forwards                   : 10
    tx v6 l2 drops                      : 0
    tx v4 l3 route forwards             : 0
    tx v4 l3 route drops                : 0
    tx v4 kernel forwards               : 0
    tx v4 kernel forward drops          : 0
    tx v4 l2 forwards                   : 5
    tx v4 l2 drops                      : 0
    tx v4 ip fragment forward           : 0
    tx v4 ip fragment DF drops          : 0
    tx v4 ip fragment drops             : 0
    tx v4 ip fragment failed            : 0
    bbe_ifl_output tx failed            : 0
    bbe_io_send tx failed               : 0
    bbe_io_send tx partial failed       : 0
    io_queue low_prio_packets_dropped   : 0
    io_queue mlow_prio_packets_dropped  : 0
    io_queue med_prio_packets_dropped   : 0
    io_queue high_prio_packets_dropped  : 0
--------------------------------------------------------------
DHCP Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                         : 27
    Tx Statistics
        transmit request packets        : 11
        sent packets                    : 11
    DHCPv4 Rx Statistics
        total packets                   : 17
        boot request                    : 15
        boot reply                      : 2
        discover                        : 11
```

```
        offer                            : 1
        request                          : 3
        ack                              : 1
        release                          : 1
    DHCPv4 Tx Statistics
        total packets                    : 5
        boot reply                       : 5
        offer                            : 2
        ack                              : 3
    DHCPv6 Rx Statistics
        total packets                    : 10
        solicit                          : 4
        advertise                        : 1
        request                          : 2
        reply                            : 1
        renew                            : 1
        release                          : 1
        relay repl                       : 2
    DHCPv6 Tx Statistics
        total packets                    : 6
        advertise                        : 2
        reply                            : 4
    Error Statistics
    Connection Statistics
        no connection packets            : 0
        connection down events           : 1
        connection up events             : 2
        flow control invoked             : 0
        flow control released            : 0
        packets sent to daemon           : 27
        packets received from daemon     : 11
        messages sent to daemon          : 0
        messages received from daemon    : 1481
        notifies while not connected     : 0
NET Statistics:
--------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                         : 2
        neighbor advertisements          : 2
    Tx Statistics
        packets:                         : 4
        neighbor solicitations           : 4
```

```
Management Statistics:
---------------------------------------------------------------
    ipdemux                             : 4
    ipdemux add                         : 4
    ipdemux delete                      : 2
    ifl flow                            : 7
    ifl flow adds                       : 6
    ifl flow changes                    : 1
    ifl flow deletes                    : 3
    ip flow                             : 9
    ip flow add                         : 6
    ip flow delete                      : 3
    service                             : 27
    service adds                        : 24
    service deletes                     : 3
Management Config Status:
---------------------------------------------------------------
    gres state enabled state            : 0
    shmlog disabled state               : 0
    shmlog filtering state              : 0
    vc backup member local switch state : 0
    dscp code point value               : 0x30
```

show system subscriber-management statistics pppoe

```
user@host> show system subscriber-management statistics pppoe
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
---------------------------------------------------------------
                     Packet Statistics
---------------------------------------------------------------
I/O Statistics:
---------------------------------------------------------------
    Rx Statistics
        packets                         : 784711
    Tx Statistics
        packets                         : 7013122
  Layer 3 Statistics
    Rx Statistics
        packets                         : 356218
    Tx Statistics
```

```
        packets                          : 6604660

PPPoE Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                          : 486165
        padis                            : 36768
        padrs                            : 35421
        ppp packets                      : 341787
        ERA discards                     : 8249
    Tx Statistics
        packets                          : 70842
        send failures                    : 6240
```

show system subscriber-management statistics extensive

```
user@host> show system subscriber-management statistics extensive
Session Manager started @ Tue Nov  3 10:00:57 2015
Session Manager cleared @ Tue Nov  3 11:10:01 2015
--------------------------------------------------------------
                     Packet Statistics
--------------------------------------------------------------
I/O Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                          : 784711
    Tx Statistics
        packets                          : 7013122
    Buffer Statistics
        allocations                      : 7032618
        frees                            : 7032624
        allocation failures              : 0
  Layer 3 Statistics
    Rx Statistics
        packets                          : 356218
    Tx Statistics
        packets                          : 6604660
    PFE Event Statistics
        packets                          : 0

--------------------------------------------------------------
```

```
Enhanced I/O Statistics:
--------------------------------------------------------------
    bbe_io_rcv l2                         : 0
    bbe_io_rcv l3                         : 0
    bbe_io_rcv l3 v4                      : 0

    io low queue drops                    :12
    io mlow queue drops                   :0
    io medium queue drops                 :0
    io high queue drops                   :0
```

show system subscriber-management statistics ppp (LCP Vendor-Specific Counters)

```
user@host> show system subscriber-management statistics ppp
Session Manager started @ Thu Feb 11 00:37:43 2020
Session Manager cleared @ Thu Feb 11 00:37:43 2020
--------------------------------------------------------------
                     Packet Statistics
--------------------------------------------------------------
I/O Statistics:
--------------------------------------------------------------
    Rx Statistics
        packets                           : 486783
    Tx Statistics
        packets                           : 144
  Layer 3 Statistics
    Rx Statistics
        packets                           : 8
    Tx Statistics
        packets                           : 0
PPP Statistics:
--------------------------------------------------------------
    Rx Statistics
        network packets                   : 123
        plugin packets                    : 123
        lcp config requests               : 18
        lcp config acks                   : 18
        lcp conf nacks                    : 8
        lcp conf rejects                  : 6
        lcp termination requests          : 4
        lcp termination acks              : 13
```

```
        lcp code rejects                    : 2
        lcp vendor-specific acks       : 10
        pap requests                      : 8
        ipcp requests                     : 27
        ipcp acks                         : 9
        ipv6cp requests                   : 11
        ipv6cp acks                       : 1
    Tx Statistics
        packets                           : 101
        lcp config requests               : 32
        lcp config acks                   : 18
        lcp termination requests          : 13
        lcp termination acks              : 4
        lcp vendor-specific requests   : 10
        pap acks                          : 8
        ipcp requests                     : 9
        ipcp acks                         : 5
        ipcp nacks                        : 9
        ipv6cp requests                   : 1
        ipv6cp acks                       : 1
        ipv6cp nacks                      : 1
NET Statistics:
--------------------------------------------------------------
  ICMP6 Statistics
    Rx Statistics
        packets:                          : 8
        router solicitations              : 8
    Tx Statistics
        packets:                          : 0
```

Release Information

Command introduced in Junos OS Release 15.1R3.

Enhanced I/O Statistics introduced as part of Extensive output in Junos OS Release 15.1R4 on MX Series routers for enhanced subscriber management.

RELATED DOCUMENTATION

Understanding Dropped Packets and Untransmitted Traffic Using show Commands

show system subscriber-management subscriber-group

IN THIS SECTION

Syntax

```
show system subscriber-management subscriber-group
<subscriber-group-name>
<redundancy-service service-id>
```

Description

Displays subscriber group information .

Options

`subscriber-group-name`	(Optional) Displays information for the specified subscriber group.
`redundancy-service service-id`	(Optional) Displays information for the specified redundancy service ID.

Required Privilege Level

view

Output Fields

Table 43 on page 441 lists the output fields for the `show system subscriber-management subscriber-group` command. Output fields are listed in the approximate order in which they appear.

Table 43: show system subscriber-management subscriber-group Output Fields

Field Name	Field Description
`Name`	Name of the subscriber group.
`ID`	ID for the subscriber group.
`Mode`	The subscriber group mode. The mode can be either `Control Plane` or `User Plane`.
`Backup Mode`	The backup mode for the subscriber group. The backup mode can be either `Hot`, `Warm-enhanced`, or `Warm-scale`.
`Oper State`	The operational state of the subscriber group. The state can be either `Active` or `Backup`.

Table 43: show system subscriber-management subscriber-group Output Fields *(Continued)*

Field Name	Field Description
`Vmac`	The virtual MAC address for the subscriber group.
`Sessions`	The number of subscriber sessions for the subscriber group.
`Ports`	The number of ports associated with the subscriber group. The ports are listed and state whether they are `admin-disabled`.
`Redundancy Services`	Grouping of subscriber sessions for full state restoration to become active. Lists the number of redundancy services.
ID	Redundancy service ID number.
`PFE-Subscriber-State`	The PFE subscriber state. The state can be the following: • `Flow Only`—Partial state subscribers, while in `warm-enhanced` backup mode. • `RE Only`—Routing engine only subscribers. • `Full Service`—Full state subscribers. • `Activating`—Displays the percentage of subscribers that are activate during the switchover process (0 to 99). • `Deactivating`—Displays the percentage of subscribers that have been deactivated during the switchover process (0 to 99). • `Exceeded Limit`—Switchover failed because the RSMON subscriber limit was reached. Subscriber logout occurs.

Table 43: show system subscriber-management subscriber-group Output Fields *(Continued)*

Field Name	Field Description
Switchover-%	The percent of the sessions that have been fully restored on switchover.

Sample Output

show system subscriber-management subscriber-group

```
user@host> show system subscriber-management subscriber-group
Name      ID    SGRP Mode      Oper State Ports Sessions
exampl-1  1024  Control Plane  Active      1     2000
```

show system subscriber-management subscriber-group

```
user@host> show system subscriber-management subscriber-group SGRP1
Name:                   SGRP1
ID:                     2048
Mode:                   User Plane
Backup Mode:            Hot
Oper State:             Backup
Vmac:                   0a:0a:0a:0a:0a:0a
Sessions:               0
Ports:                  1
  ps0              (admin-disabled)
```

show system subscriber-management subscriber-group

```
user@host> show system subscriber-management subscriber-group exampl-1
Name:                   example-1
ID:                     2048
Mode:                   Control Plane
Backup Mode:            Hot
```

```
Oper State:             Backup
Vmac:                   22:0:0:0:0:22
Sessions:               10000
Ports:                  1
  ae1

Redundancy Services:    10
ID          Sessions PFE-Subscriber-State Switchover-%
524322      1000     Flow Only            -
524323      1000     Flow Only            -
524324      1000     Flow Only            -
524325      1000     Flow Only            -
```

show system subscriber-management subscriber-group *subscriber-group-name* redundancy-service

```
user@host> show system subscriber-management subscriber-group SGRP1 redundancy-service 524322
Name:                   SGRP1
ID:                     2048
Mode:                   Control Plane
Backup Mode:            Warm-enhanced
Oper State:             Backup
Vmac:                   22:0:0:0:0:22
Sessions:               10000
Ports:                  1
  ae1

Redundancy Service:

Service Id:           524322
Sessions:             1006
PFE Subscriber State: Flow Only
Switchover %:         -
Session Id
221
222
223
```

show system subscriber-management summary

IN THIS SECTION

Syntax

```
show system subscriber-management summary
```

Description

Display complete subscriber management database summary information.

Options

none   This command has no options.

Required Privilege Level

view

Output Fields

Table 44 on page 446 lists the output fields for the `show system subscriber-management summary` command. Output fields are listed in the approximate order in which they appear.

Table 44: show system subscriber-management summary Output Fields

Field Name	Field Description
`Graceful Restart`	State of graceful Routing Engine switchover (GRES): • `Enabled` • `Disabled` (Enhanced subscriber management for MX Series routers) The name of this field is `Graceful Switchover`.
`Mastership`	State of the Routing Engine: • `Master` • `Standby`
`Database`	State of the subscriber management database: • `Available` • `Init` • `Not-available`
`Standby`	(Enhanced subscriber management for MX Series routers) State of the standby Routing Engine: • `Connected`—Connected but not synchronized • `Disconnected`—Not connected • `Resync (`*nn*`%)`—Connected and *nn* percent synchronized with the primary Routing Engine • `Synchronized`—Synchronized with the primary Routing Engine

Table 44: show system subscriber-management summary Output Fields *(Continued)*

Field Name	Field Description
`Disconnection Reason`	Reason why both Routing Engines are disconnected when there is a DRAM mismatch. • Primary/Standy RE DRAM Size Mismatch—Displayed when the amount of memory is different on the primary and standby Routing Engines.
`Chassisd ISSU State`	State of unified ISSU chassis daemon: • `ABORT` • `DAEMON_ISSU_PREPARE` • `DAEMON_ISSU_PREPARE_DONE` • `DAEMON_SWITCHOVER_PREPARE` • `DAEMON_SWITCHOVER_PREPARE_DONE` • `FRU_ISSU` • `FRU_ISSU_DONE` • `IDLE` • `UNKNOWN`
`ISSU State`	State of unified ISSU aggregate daemon: • `ABORT` • `IDLE` • `PREPARE` • `READY` • `SWITCHOVER_PREPARE` • `SWITCHOVER_READY` • `UNKNOWN`

Table 44: show system subscriber-management summary Output Fields *(Continued)*

Field Name	Field Description
`ISSU Wait`	Amount of time, in seconds, requested by a daemon to perform cleanup. If multiple daemons request time, the displayed value is the highest wait time requested by a daemon.

Sample Output

show system subscriber-management summary

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

show system subscriber-management summary (Enhanced Subscriber Management)

```
user@host> show system subscriber-management summary
General:
    Graceful Switchover     Enabled
    Mastership              Master
    Database                Available
    Standby                   Resync (75%)
    Chassisd ISSU State     IDLE
    ISSU State              IDLE
    ISSU Wait               0
```

show system subscriber-management summary (DRAM Size Mismatch Error)

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Standby              Disconnected
<emphasis>    Disconnection Reason    Master/Standy RE DRAM Size Mismatch</emphasis
>
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

Release Information

Command introduced in Junos OS Release 11.1.

show user-plane class-of-service classifier

IN THIS SECTION

Syntax

```
show user-plane class-of-service classifier
<classifier-name classifier-name>
<name name>
<type>
```

Description

Show mapping of code point to forwarding class/loss priority at user-plane logical interface.

Options

none	Display all classifiers.
`classifier-name` *`classifier-name`*	(Optional) Name of classifier
`name` *`name`*	(Optional) Display named classifier.
`type` *`type`*	(Optional) Displays the classifier of the specified type. The classifier type can be one of the following: • dscp—Differentiated Services code point (DSCP) • dscp-ipv6—Differentiated Services code point (DSCP) for IPv6 • exp—MPLS experimental code point • ieee-802.1—IEEE-802.1 code point • ieee-802.1ad—IEEE-802.1ad (DEI) code point • inet6-precedence—IPv6 precedence code point • inet-precedence—IPv4 precedence code point

Required Privilege Level

view

Output Fields

Table 45 on page 451 describes the output fields for the `show user-plane class-of-service classifier` command. Output fields are listed in the approximate order in which they appear.

Table 45: show user-plane class-of-service classifier Output Fields

Field Name	Field Description
`Classifier`	Name of the classifier at user plane.
`Code point type`	Type of the classifier: `exp` (not on EX Series switch), `dscp`, `dscp-ipv6` (not on EX Series switch), `ieee-802.1`, or `inet-precedence`.
`Index`	Internal index of the classifier.
`Code point`	Code point value used for classification
`Forwarding class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
`Loss priority`	Loss priority value used for classification. For most platforms, the value is `high` or `low`. For some platforms, the value is `high`, `medium-high`, `medium-low`, or `low`.

Sample Output

show user-plane class-of-service classifier

```
user@host>show user-plane class-of-service classifier
Classifier: dscp-default, Code point type: dscp, Index: 7
  Code point         Forwarding class                    Loss priority
  000001             q0                                  low
  000010             q0                                  low
  000011             q0                                  low
  000100             q0                                  low
  000101             q0                                  low
  000110             q0                                  low
  000111             q0                                  low
  001000             q0                                  low
  001001             q0                                  low
  001010             q2                                  low
  001011             q0                                  low
  001100             q2                                  high
  001101             q0                                  low
  001110             q2                                  high
  001111             q0                                  low
  010000             q0                                  low
  010001             q0                                  low
  010010             q0                                  low
  010011             q0                                  low
  010100             q0                                  low
  010101             q0                                  low
  010110             q0                                  low
  010111             q0                                  low
  011000             q0                                  low
  011001             q0                                  low
  011010             q0                                  low
  011011             q0                                  low
  011100             q0                                  low
  011101             q0                                  low
  011110             q0                                  low
  011111             q0                                  low
  100000             q0                                  low
  100001             q0                                  low
  100010             q0                                  low
```

```
  100011          q0                                  low
  100100          q0                                  low
  100101          q0                                  low
  100110          q0                                  low
  100111          q0                                  low
  101000          q0                                  low
  101001          q0                                  low
  101010          q0                                  low
  101011          q0                                  low
  101100          q0                                  low
  101101          q0                                  low
  101110          q1                                  low
  101111          q0                                  low
  110000          q3                                  low
  110001          q0                                  low
  110010          q0                                  low
  110011          q0                                  low
  110100          q0                                  low
  110101          q0                                  low
  110110          q0                                  low
  110111          q0                                  low
  111000          q3                                  low
  111001          q0                                  low
  111010          q0                                  low
  111011          q0                                  low
  111100          q0                                  low
  111101          q0                                  low
  111110          q0                                  low
  111111          q0                                  low

Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 8
  Code point      Forwarding class                    Loss priority
  000001          q0                                  low
  000010          q0                                  low
  000011          q0                                  low
  000100          q0                                  low
  000101          q0                                  low
  000110          q0                                  low
  000111          q0                                  low
  001000          q0                                  low
  001001          q0                                  low
  001010          q2                                  low
  001011          q0                                  low
```

001100	q2	high
001101	q0	low
001110	q2	high
001111	q0	low
010000	q0	low
010001	q0	low
010010	q0	low
010011	q0	low
010100	q0	low
010101	q0	low
010110	q0	low
010111	q0	low
011000	q0	low
011001	q0	low
011010	q0	low
011011	q0	low
011100	q0	low
011101	q0	low
011110	q0	low
011111	q0	low
100000	q0	low
100001	q0	low
100010	q0	low
100011	q0	low
100100	q0	low
100101	q0	low
100110	q0	low
100111	q0	low
101000	q0	low
101001	q0	low
101010	q0	low
101011	q0	low
101100	q0	low
101101	q0	low
101110	q1	low
101111	q0	low
110000	q3	low
110001	q0	low
110010	q0	low
110011	q0	low
110100	q0	low
110101	q0	low
110110	q0	low

```
  110111          q0                              low
  111000          q3                              low
  111001          q0                              low
  111010          q0                              low
  111011          q0                              low
  111100          q0                              low
  111101          q0                              low
  111110          q0                              low
  111111          q0                              low

Classifier: dscp-ipv6-compatibility, Code point type: dscp-ipv6, Index: 9
  Code point      Forwarding class                Loss priority
  000001          q0                              low
  000010          q0                              low
  000011          q0                              low
  000100          q0                              low
  000101          q0                              low
  000110          q0                              low
  000111          q0                              low
  001000          q0                              low
  001001          q0                              low
  001010          q0                              low
  001011          q0                              low
  001100          q0                              low
  001101          q0                              low
  001110          q0                              low
  001111          q0                              low
  010000          q0                              low
  010001          q0                              low
  010010          q0                              low
  010011          q0                              low
  010100          q0                              low
  010101          q0                              low
  010110          q0                              low
  010111          q0                              low
  011000          q0                              low
  011001          q0                              low
  011010          q0                              low
  011011          q0                              low
  011100          q0                              low
  011101          q0                              low
  011110          q0                              low
  011111          q0                              low
```

```
  100000          q0                    low
  100001          q0                    low
  100010          q0                    low
  100011          q0                    low
  100100          q0                    low
  100101          q0                    low
  100110          q0                    low
  100111          q0                    low
  101000          q0                    low
  101001          q0                    low
  101010          q0                    low
  101011          q0                    low
  101100          q0                    low
  101101          q0                    low
  101110          q0                    low
  101111          q0                    low
  110000          q3                    low
  110001          q0                    low
  110010          q0                    low
  110011          q0                    low
  110100          q0                    low
  110101          q0                    low
  110110          q0                    low
  110111          q0                    low
  111000          q3                    low
  111001          q0                    low
  111010          q0                    low
  111011          q0                    low
  111100          q0                    low
  111101          q0                    low
  111110          q0                    low
  111111          q0                    low

Classifier: exp-default, Code point type: exp, Index: 10
  Code point      Forwarding class      Loss priority
  001             q0                    high
  010             q1                    low
  011             q1                    high
  100             q2                    low
  101             q2                    high
  110             q3                    low
  111             q3                    high
```

```
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point         Forwarding class                    Loss priority
  001                q0                                  high
  010                q1                                  low
  011                q1                                  high
  100                q2                                  low
  101                q2                                  high
  110                q3                                  low
  111                q3                                  high

Classifier: ipprec-default, Code point type: inet-precedence, Index: 12
  Code point         Forwarding class                    Loss priority
  001                q2                                  low
  010                q0                                  low
  011                q0                                  low
  100                q0                                  low
  101                q1                                  low
  110                q3                                  low
  111                q3                                  high

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                    Loss priority
  001                q0                                  high
  010                q0                                  low
  011                q0                                  high
  100                q0                                  low
  101                q0                                  high
  110                q3                                  low
  111                q3                                  high
```

show user-plane class-of-service classifier name

```
user@host> show user-plane class-of-service classifier name ipprec-default type inet-
precedenc
Classifier: ipprec-default, Code point type: inet-precedence, Index: 12
  Code point         Forwarding class                    Loss priority
  001                q2                                  low
  010                q0                                  low
  011                q0                                  low
  100                q0                                  low
  101                q1                                  low
```

```
110                 q3                                  low
111                 q3                                  high
```

show user-plane class-of-service interface

IN THIS SECTION

Syntax

```
show user-plane class-of-service interface
<interface-name interface-name>
<detail>
```

Description

Displays BNG User Plane specific class-of-service (CoS) information for subscriber logical interfaces. This command displays the subscriber logical interface association for the classifier, rewrite rules, and scheduler map objects.

NOTE: This topic lists all possible options and output fields for the `show user-plane class-of-service interface` command. Options and output fields can vary depending on the platform, software release, and operating system (Junos OS or Junos OS Evolved).

NOTE: On routing platforms with dual Routing Engines, running this command on the backup Routing Engine, with or without any of the available options, is not supported and produces the following error message:

```
error: the class-of-service subsystem is not running
```

Options

none	Display CoS associations for all physical and logical interfaces.
`interface-name` *`interface-name`*	(Optional) Display class-of-service (CoS) associations for the specified interface.
detail	(Optional) Display QoS and CoS information based on the interface.

If the `interface` *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the `interface` *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

Required Privilege Level

view

Output Fields

Table 46 on page 460 describes the output fields for the `show user-plane class-of-service interface` command. Output fields are listed in the approximate order in which they appear.

Table 46: show user-plane class-of-service interface Output Fields

Field Name	Field Description
`Physical interface`	Name of a physical interface.
`Index`	Index of this interface or the internal index of this object. (Enhanced subscriber management) Index values for dynamic CoS traffic control profiles and dynamic scheduler maps are larger for enhanced subscriber management than they are for legacy subscriber management.
`Dedicated Queues`	Status of dedicated queues configured on an interface. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.
`Maximum usable queues`	Number of queues you can configure on the interface.
`Maximum usable queues`	Maximum number of queues you can use.
`Total non-default queues created`	Number of queues created in addition to the default queues. (Enhanced subscriber management) This field is not displayed for enhanced subscriber management.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Rewrite Input IEEE Code-point`	IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
`Shaping rate`	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the `Shaping rate` field is displayed for either the physical interface or the logical interface.
`Scheduler map`	Name of the output scheduler map associated with this interface. (Enhanced subscriber management) The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Scheduler map forwarding class sets`	Name of the output fabric scheduler map associated with a QFabric system Interconnect device interface.
`Input shaping rate`	Maximum transmission rate on the input interface.
`Input scheduler map`	Name of the input scheduler map associated with this interface.
`Chassis scheduler map`	Name of the scheduler map associated with the packet forwarding component queues.
`Rewrite`	Name and type of the rewrite rules associated with this interface.
`Traffic-control-profile`	Name of the associated traffic control profile. (Enhanced subscriber management) The name of the dynamic traffic control profile object is associated with a generated UID (for example, `TC_PROF_100_199_SERIES_UID1006`) instead of with a subscriber interface.
`Classifier`	Name and type of classifiers associated with this interface.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Forwarding-class-map`	Name of the forwarding map associated with this interface.
`Congestion-notification`	Congestion notification state, `enabled` or `disabled`.
`Dedicated Buffer Profile`	Name of the dedicated buffer profile associated with the interface.
`Monitoring Profile Name`	Name of the monitoring profile defined to monitor the peak queue length for virtual output queues (VOQs) for the interface.
`Logical interface`	Name of a logical interface.
`Object`	Category of an object: `Classifier`, `Fragmentation-map` (for LSQ interfaces only), `Scheduler-map`, `Rewrite`, `Translation Table`, or `traffic-class-map`.
`Name`	Name of an object.
`Type`	Type of an object: `dscp`, `dscp-ipv6`, `exp`, `ieee-802.1`, `ip`, `inet-precedence`, or `ieee-802.1ad`.
`Link-level type`	Encapsulation on the physical interface.
`MTU`	MTU size on the physical interface.
`Speed`	Speed at which the interface is running.
`Loopback`	Whether loopback is enabled and the type of loopback.
`Source filtering`	Whether source filtering is enabled or disabled.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Flow control`	Whether flow control is enabled or disabled.
`Auto-negotiation`	Whether autonegotiation is enabled or disabled.
`Remote-fault`	Remote fault status. • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.
`Device flags`	The `Device flags` field provides information about the physical device and displays one or more of the following values: • `Down`—Device has been administratively disabled. • `Hear-Own-Xmit`—Device receives its own transmissions. • `Link-Layer-Down`—The link-layer protocol has failed to connect with the remote endpoint. • `Loopback`—Device is in physical loopback. • `Loop-Detected`—The link layer has received frames that it sent, thereby detecting a physical loopback. • `No-Carrier`—On media that support carrier recognition, no carrier is currently detected. • `No-Multicast`—Device does not support multicast traffic. • `Present`—Device is physically present and recognized. • `Promiscuous`—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • `Quench`—Transmission on the device is quenched because the output buffer is overflowing. • `Recv-All-Multicasts`—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • `Running`—Device is active and enabled.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Interface flags`	The `Interface flags` field provides information about the physical interface and displays one or more of the following values:
	• `Admin-Test`—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
	• `Disabled`—Interface is administratively disabled.
	• `Down`—A hardware failure has occurred.
	• `Hardware-Down`—Interface is nonfunctional or incorrectly connected.
	• `Link-Layer-Down`—Interface keepalives have indicated that the link is incomplete.
	• `No-Multicast`—Interface does not support multicast traffic.
	• `No-receive No-transmit`—Passive monitor mode is configured on the interface.
	• `Point-To-Point`—Interface is point-to-point.
	• `Pop all MPLS labels from packets of depth`—MPLS labels are removed as packets arrive on an interface that has the `pop-all-labels` statement configured. The depth value can be one of the following:
	  ◦ `1`—Takes effect for incoming packets with one label only.
	  ◦ `2`—Takes effect for incoming packets with two labels only.
	  ◦ `[ 1 2 ]`—Takes effect for incoming packets with either one or two labels.
	• `Promiscuous`—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
	• `Recv-All-Multicasts`—Interface is in multicast promiscuous mode and provides no multicast filtering.
	• `SNMP-Traps`—SNMP trap notifications are enabled.
	• `Up`—Interface is enabled and operational.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Flags	The Logical interface flags field provides information about the logical interface and displays one or more of the following values: • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Link flags`	Provides information about the physical link and displays one or more of the following values: • `ACFC`—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • `Give-Up`—Link protocol does not continue connection attempts after repeated failures. • `Loose-LCP`—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • `Loose-LMI`—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • `Loose-NCP`—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • `Keepalives`—Link protocol keepalives are enabled. • `No-Keepalives`—Link protocol keepalives are disabled. • `PFC`—Protocol field compression is configured. The PPP session negotiates the PFC option.
`Hold-times`	Current interface hold-time up and hold-time down, in milliseconds.
`CoS queues`	Number of CoS queues configured.
`Last flapped`	Date, time, and how long ago the interface went from down to up. The format is `Last flapped:` *year-month-day hour*:*minute*:*second*:*timezone* (*hour*:*minute*:*second* `ago`). For example, `Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)`.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Statistics last cleared`	Number and rate of bytes and packets received and transmitted on the physical interface. • `Input bytes`—Number of bytes received on the interface. • `Output bytes`—Number of bytes transmitted on the interface. • `Input packets`—Number of packets received on the interface. • `Output packets`—Number of packets transmitted on the interface.
Exclude Overhead Bytes	Exclude the counting of overhead bytes from aggregate queue statistics. • `Disabled`—Default configuration. Includes the counting of overhead bytes in aggregate queue statistics. • `Enabled`—Excludes the counting of overhead bytes from aggregate queue statistics for just the physical interface. • `Enabled for hierarchy`—Excludes the counting of overhead bytes from aggregate queue statistics for the physical interface as well as all child interfaces, including logical interfaces and interface sets.
`IPv6 transit statistics`	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Input errors`	Input errors on the interface. The labels are explained in the following list: • `Errors`—Sum of the incoming frame terminations and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Giants`—Number of frames received that are larger than the giant threshold. • `Bucket Drops`—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `HS link CRC errors`—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • `HS link FIFO overflows`—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Output errors`	Output errors on the interface. The labels are explained in the following list: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • `Errors`—Sum of the outgoing frame terminations and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Aged packets`—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • `HS link FIFO underflows`—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • `MTU errors`—Number of packets whose size exceeds the MTU of the interface.
`Egress queues`	Total number of egress Maximum usable queues on the specified interface.
`Queue counters`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism.
`SONET alarms` `SONET defects`	SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: `SONET PHY`, `SONET section`, `SONET line`, and `SONET path`.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`SONET PHY`	Counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. A state other than `OK` indicates a problem. The `SONET PHY` field has the following subfields: • `PLL Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal
`SONET section`	Counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. A state other than `OK` indicates a problem. The `SONET section` field has the following subfields: • `BIP-B1`—Bit interleaved parity for SONET section overhead • `SEF`—Severely errored framing • `LOS`—Loss of signal • `LOF`—Loss of frame • `ES-S`—Errored seconds (section) • `SES-S`—Severely errored seconds (section) • `SEFS-S`—Severely errored framing seconds (section)

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`SONET line`	Active alarms and defects, plus counts of specific SONET errors with detailed information. • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. A state other than `OK` indicates a problem. The `SONET line` field has the following subfields: • `BIP-B2`—Bit interleaved parity for SONET line overhead • `REI-L`—Remote error indication (near-end line) • `RDI-L`—Remote defect indication (near-end line) • `AIS-L`—Alarm indication signal (near-end line) • `BERR-SF`—Bit error rate fault (signal failure) • `BERR-SD`—Bit error rate defect (signal degradation) • `ES-L`—Errored seconds (near-end line) • `SES-L`—Severely errored seconds (near-end line) • `UAS-L`—Unavailable seconds (near-end line) • `ES-LFE`—Errored seconds (far-end line) • `SES-LFE`—Severely errored seconds (far-end line) • `UAS-LFE`—Unavailable seconds (far-end line)

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. The SONET path field has the following subfields: • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Received SONET overhead Transmitted SONET overhead	Values of the received and transmitted SONET overhead: • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-*N* signal. • Z3 and Z4—Allocated for future use.
Received path trace Transmitted path trace	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.
HDLC configuration	Information about the HDLC configuration. • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
CoS information	Information about the CoS queue for the physical interface. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress Maximum usable queues on the specified interface.
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Transmitted Packets`	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values.
`Transmitted Bytes`	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
`Tail-dropped packets`	Number of packets dropped because of tail drop.

`RED-dropped packets`	Number of packets dropped because of random early detection (RED).
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP packets dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP packets dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP packets dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP packets dropped because of RED.
	• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories:
	  • `Low`—Number of low-loss priority packets dropped because of RED.
	  • `Medium-low`—Number of medium-low loss priority packets dropped because of RED.
	  • `Medium-high`—Number of medium-high loss priority packets dropped because of RED.
	  • `High`—Number of high-loss priority packets dropped because of RED.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`RED-dropped bytes`	Number of bytes dropped because of RED. The byte counts vary by PIC type.
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP bytes dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP bytes dropped because of RED.

`Transmit rate`	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
`Rate Limit`	Rate limiting configuration of the queue. Possible values are : • `None`—No rate limit. • `exact`—Queue transmits at the configured rate.
`Buffer size`	Delay buffer size in the queue.
`Priority`	Scheduling priority configured as `low` or `high`.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler: `low`, `medium-low`, `medium-high`, `high`, or `none`.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.
`Excess Priority`	Priority of the excess bandwidth traffic on a scheduler.
`Drop profiles`	Display the assignment of drop profiles. • `Loss priority`—Packet loss priority for drop profile assignment. • `Protocol`—Transport protocol for drop profile assignment. • `Index`—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • `Name`—Name of the drop profile. • `Type`—Type of the drop profile: `discrete` or `interpolated`. • `Fill Level`—Percentage fullness of a queue. • `Drop probability`—Drop probability at this fill level.

Table 46: show user-plane class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`Adjustment information`	Display the assignment of shaping-rate adjustments on a scheduler node or queue.
	• `Adjusting application`—Application that is performing the shaping-rate adjustment.
	  • The adjusting application can appear as `ancp LS-0`, which is the Junos OS Access Node Control Profile process (`ancpd`) that performs shaping-rate adjustments on schedule nodes.
	  • The adjusting application can appear as `DHCP`, which adjusts the shaping-rate and overhead-accounting class-of-service attributes based on DSL Forum VSA conveyed in DHCP option 82, suboption 9 (Vendor Specific Information). The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	  • The adjusting application can also appear as `pppoe`, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
	• `Adjustment type`—Type of adjustment: `absolute` or `delta`.
	• `Configured shaping rate`—Shaping rate configured for the scheduler node or queue.
	• `Adjustment value`—Value of adjusted shaping rate.
	• `Adjustment target`—Level of shaping-rate adjustment performed: `node` or `queue`.
	• `Adjustment overhead-accounting mode`—Configured shaping mode: `frame` or `cell`.
	• `Adjustment overhead bytes`—Number of bytes that the ANCP agent adds to or subtracts from the actual downstream frame overhead before reporting the adjusted values to CoS.
	• `Adjustment target`—Level of shaping-rate adjustment performed: `node` or `queue`.
	• `Adjustment multicast index`—

Sample Output

show user-plane class-of-service interface

```
user@host> show user-plane class-of-service interface
  Logical interface: demux0.3222274073, Index: 3222274073
Object                 Name                   Type                    Index
Classifier             dscp-ipv6-compatibility dscp-ipv6                  9
Classifier             ipprec-compatibility   ip                         13

  Logical interface: pp0.3222274074, Index: 3222274074
Object                 Name                   Type                    Index
Traffic-control-profile cp::tcp1_UID1129      Output              4295491731
Scheduler-map          cp::smap_UID1128       Output              4294967355
Classifier             ipprec-compatibility   ip                         13
```

show user-plane class-of-service interface *interface-name*

```
user@host> show user-plane class-of-service interface pp0.3222274074
  Logical interface: pp0.3222274074, Index: 3222274074
Object                 Name                   Type                    Index
Traffic-control-profile cp::tcp1_UID1129      Output              4295491731
Scheduler-map          cp::smap_UID1128       Output              4294967355
Classifier             ipprec-compatibility   ip                         13
```

show class-of-service interface detail

```
user@host> show user-plane class-of-service interface pp0.3222274074 detail
  Logical interface: pp0.3222274074, Index: 3222274074
Object                 Name                   Type                    Index
Traffic-control-profile cp::tcp1_UID1129      Output              4295491731
Scheduler-map          cp::smap_UID1128       Output              4294967355
Classifier             ipprec-compatibility   ip                         13
```

show user-plane class-of-service interface-set

IN THIS SECTION

Syntax

```
show user-plane class-of-service interface-set
<interface-set-name>
```

Description

Show mapping of CoS objects to an interface set.

Options

interface-set *interface-set-name*	(Optional) Display CoS associations for the specified interface set.

Required Privilege Level

view

Output Fields

Table 46 on page 460 describes the output fields for the `show user-plane class-of-service interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 47: show user-plane class-of-service interface-set Output Fields

Field Name	Field Description
`Interface-set`	Name of a logical interface set composed of one or more logical interfaces for which hierarchical scheduling is enabled.
`Index`	Index number of this interface set or the internal index number of this object.
`Physical interface`	Name of a physical interface.
`Output traffic control profile`	Name of the output traffic control profile attached to the logical interface set.
`Output traffic control profile remaining`	Name of the output traffic control profile for remaining traffic attached to the logical interface set.

Sample Output

show user-plane class-of-service interface-set

```
user@host> show user-plane class-of-service interface-set bb0.1-1
Interface-set: bb0.1-1, Index: 65543
Physical interface: et-0/0/11
  Output traffic control profile: cp::tcp2_UID1152, Index: 4295491847
  Output traffic control profile remaining: cp::tcp3-rem_UID1153, Index: 4295491848
```

show user-plane class-of-service rewrite-rule

IN THIS SECTION

Syntax

```
show user-plane class-of-service rewrite-rule
<rewrite-rule-name rewrite-rule-name>
<type type>
```

Description

Displays BNG User plane specific class-of-service (CoS) rewrite rules information. This command displays the mapping of forwarding classes and loss priority to code point values.

Options

none	Display all rewrite rules.
`rewrite-rule-name` *`rewrite-rule-name`*	(Optional) Display the specified rewrite rule.

`type type` (Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:

- **dscp**—Differentiated Services code point (DSCP) for IPv4
- **dscp-ipv6**—Differentiated Services code point (DSCP) for IPv6
- **exp**—MPLS experimental code point
- **frame-relay-de**—Frame relay discard eligible bit code point
- **ieee-802.1**—IEEE-802.1 code point
- **ieee-802.1ad**—IEEE-802.1ad (DEI) code point
- **inet6-precedence**—IPv6 precedence code point
- **inet-precedence**—IPv4 precedence code point.

Required Privilege Level

view

Output Fields

Table 48 on page 483 describes the output fields for the `show user-plane class-of-service rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 48: show user-plane class-of-service rewrite-rule Output Fields

Field Name	Field Description
Rewrite rule	Name of the rewrite rule at user-plane.
Code point type	Type of rewrite rule: **dscp**, **dscp-ipv6**, **exp**, **frame-relay-de**, or **inet-precedence**.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch.

Table 48: show user-plane class-of-service rewrite-rule Output Fields *(Continued)*

Field Name	Field Description
Index	Internal index for this particular rewrite rule.
Loss priority	Loss priority for rewriting.
Code point	Code point value to rewrite.

Sample Output

show user-plane class-of-service rewrite-rule

```
user@host> show user-plane class-of-service rewrite-rule
Rewrite rule: cp::dscp-default, Code point type: dscp, Index: 138
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000000
  q0                                  high                000000
  q1                                  low                 101110
  q1                                  high                101110
  q2                                  low                 000000
  q2                                  high                000000
  q3                                  low                 000000
  q3                                  high                000000

Rewrite rule: cp::dscp-ipv6-default, Code point type: dscp-ipv6, Index: 139
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000000
  q0                                  high                000000
  q1                                  low                 101110
  q1                                  high                101110
  q2                                  low                 000000
  q2                                  high                000000
  q3                                  low                 000000
  q3                                  high                000000
```

```
Rewrite rule: cp::exp-default, Code point type: exp, Index: 140
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000
  q0                                  high                000
  q1                                  low                 011
  q1                                  high                011
  q2                                  low                 000
  q2                                  high                000
  q3                                  low                 000
  q3                                  high                000

Rewrite rule: cp::ieee8021p-default, Code point type: ieee-802.1, Index: 141
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000
  q0                                  high                000
  q1                                  low                 011
  q1                                  high                011
  q2                                  low                 000
  q2                                  high                000
  q3                                  low                 000
  q3                                  high                000

Rewrite rule: cp::ipprec-default, Code point type: inet-precedence, Index: 142
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000
  q0                                  high                000
  q1                                  low                 101
  q1                                  high                101
  q2                                  low                 000
  q2                                  high                000
  q3                                  low                 000
  q3                                  high                000
```

show user-plane class-of-service rewrite-rule name

```
user@host> show user-plane class-of-service rewrite-rule name cp::ipprec-default type inet-
precedence
Rewrite rule: cp::ipprec-default, Code point type: inet-precedence, Index: 142
  Forwarding class                    Loss priority       Code point
  q0                                  low                 000
  q0                                  high                000
```

```
q1                  low         101
q1                  high        101
q2                  low         000
q2                  high        000
q3                  low         000
q3                  high        000
```

show user-plane class-of-service scheduler-map

IN THIS SECTION

Syntax

```
show user-plane class-of-service scheduler-map
<scheduler-map-name>
```

Description

Displays BNG User plane specific class-of-service (CoS) scheduler map information. This command displays the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options

none	Display all schedular maps.
scheduler-map-name	Name of the scheduler for which you want to view information. Displays a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level

view

Output Fields

Table 49 on page 487 lists all possible output fields for the `show user-plane class-of-service scheduler-map` command. The fields that appear may vary depending on the platform and software release.

Table 49: show user-plane class-of-service scheduler-map Output Fields

Field Name	Field Description
`Scheduler map`	Name of the scheduler map. The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Index`	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles. Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
`Scheduler`	Name of the scheduler.
`Forwarding class`	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Table 49: show user-plane class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
Transmit rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder, which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are none, meaning no rate limiting, and exact, meaning the queue only transmits at the configured rate.
Buffer Rate	The buffer rate is the target rate of a VOQ, which is the intended egress queue rate during typical congestion. The target rate is an import parameter for VOQs, as the ratio of the drain rate to the target rate determines how much buffer adjustment happens at ingress.
Maximum buffer delay	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword remainder to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
Buffer size	The size of the buffer as a percent of the total buffer size for the port. The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent. On QFX10000 switches, the buffer size is the amount of time in milliseconds of port bandwidth that a queue can use to continue to transmit packets during periods of congestion, before the buffer runs out and packets begin to drop.
Buffer dynamic threshold	The alpha value of the shared-buffer pool at the egress buffer partition for each queue.
Priority	Scheduling priority: low or high.
Excess priority	Priority of excess bandwidth: low, medium-low, medium-high, high, or none.

Table 49: show user-plane class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
`Explicit Congestion Notification`	Explicit congestion notification (ECN) state: • Disable—ECN is disabled on the specified scheduler • Enable—ECN is enabled on the specified scheduler ECN is disabled by default.
`Adjust minimum`	Minimum shaping rate for an adjusted queue, in bps.
`Adjust percent`	Bandwidth adjustment applied to a queue, in percent.
`Drop profiles`	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
`Loss priority`	Packet loss priority for drop profile assignment.
`Protocol`	Transport protocol for drop profile assignment.
`Name`	Name of the drop profile.

Sample Output

show user-plane class-of-service scheduler-map (Logical Interface)

```
user@host> show user-plane class-of-service scheduler-map
Scheduler map: cp::smap_UID1128, Index: 4294967355

  Scheduler: cp::s1_UID1120, Forwarding class: q0, Index: 4295491712
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
```

```
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s2_UID1121, Forwarding class: q1, Index: 4295491714
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: strict-high
    Excess Priority: unspecified
    Shaping rate: 10000000 bps
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s3_UID1122, Forwarding class: q2, Index: 4295491715
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Shaping rate: 10000000 bps
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s4_UID1123, Forwarding class: q3, Index: 4295491716
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: high
    Excess Priority: unspecified
    Shaping rate: 10000000 bps
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128
```

```
  Scheduler: cp::s5_UID1124, Forwarding class: q4, Index: 4295491717
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s6_UID1125, Forwarding class: q5, Index: 4295491718
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s7_UID1126, Forwarding class: q6, Index: 4295491719
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
      High            TCP           128

  Scheduler: cp::s8_UID1127, Forwarding class: q7, Index: 4295491720
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP       128
      Low             TCP           128
      High            non-TCP       128
```

```
        High            TCP           128
```

show user-plane class-of-service traffic-control-profile

IN THIS SECTION

Syntax

```
show user-plane class-of-service traffic-control-profile
<profile-name>
```

Description

Displays BNG User plane or logical interface specific class-of-service (CoS) traffic control profile information. This command displays information on traffic shaping and scheduling profiles.

Options

none	Displays all traffic control profiles.
profile-name	Name of the traffic control profile for which you want to view information.

Required Privilege Level

view

Output Fields

Table 50 on page 493 describes the possible output fields for the `show class-of-service traffic-control-profile` command. The output fields that display depend on the platform and software device of your device. Output fields are listed in the approximate order in which they appear.

Table 50: show user-plane class-of-service traffic-control-profile Output Fields

Field Name	Field Description
`Traffic control profile`	Name of the traffic control profile. (Enhanced subscriber management for MX Series routers) The name of the dynamic traffic control profile object is associated with a generated UID (for example, `TC_PROF_100_199_SERIES_UID1000`) instead of with a subscriber interface.
`Index`	Index number of the traffic control profile. (Enhanced subscriber management for MX Series routers) Index values for dynamic CoS traffic control profiles are larger for enhanced subscriber management than they are for legacy subscriber management.
`Shaping rate`	Configured shaping rate, in bps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.

Table 50: show user-plane class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Shaping rate burst`	Configured burst size for the shaping rate, in bytes. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
`Scheduler map`	Name of the associated scheduler map. (Enhanced subscriber management for MX Series routers) The name of the dynamic scheduler map object is associated with a generated UID (for example, `SMAP-1_UID1002`) instead of with a subscriber interface.
`Guaranteed rate`	Configured guaranteed rate, in bps or cps. **NOTE**: This value depends on the service category chosen. Possible values: • `cbr`—Guaranteed rate is equal to the configured peak rate in cps. • `rtvbr`—Guaranteed rate is equal to the configured sustained rate in cps. • `nrtvbr`—Guaranteed rate is equal to the configured sustained rate in cps.
`User-plane`	The associated BNG User Plane.

Sample Output

show user-plane class-of-service traffic-control-profile

```
user@host> show user-plane class-of-service traffic-control-profile
Traffic control profile: cp::tcp1_UID1129, Index: 4295491731
  Shaping rate: 50000000
  Scheduler map: cp::smap_UID1128
```

show user-plane <user-plane-name> class-of-service traffic-control-profile *profile-name*

```
user@host> show user-plane test-123 class-of-service traffic-control-profile cp::tcp1_UID1129
Traffic control profile: cp::tcp1_UID1129, Index: 4295491731
  Shaping rate: 50000000
  Scheduler map: cp::smap_UID1128
```

show user-plane interfaces

IN THIS SECTION

Syntax

```
show user-plane interfaces
<interface-name>
```

Description

Shows BNG User Plane interface information.

Options

none	Displays information for all interfaces on the BNG User Plane.
interface-name	(Optional) Displays information for the specified interface.

Required Privilege Level

view

Output Fields

See Table 51 on page 496 for the output fields for the `show user-plane interfaces` command.

Table 51: show user-plane interfaces Output Fields

Field Name	Field Description
`Interface`	Interface name.
`Admin`	Administrative state of the interface (`Up` or `Down`).
`Link`	Status of the physical link (`Up` or `Down`).
`Proto`	Protocol family configured on the interface.
`Local`	IP address of the logical interface.
`Remote`	IP address of the remote interface.

Sample Output

show user-plane interfaces

```
user@host> show user-plane interfaces
Interface               Admin Link Proto    Local                Remote
demux0.2147516420       up    up   inet     0.0.0.0            --> 0/0
demux0.3221225474       up    up   inet     20.0.0.1           --> 0/0
demux0.3221225475       up    up   inet     20.0.0.1           --> 0/0
pp0.3221225478          up    up   inet     16.0.0.1           --> 0/0
```

show user-plane interfaces demux0

```
user@host> show user-plane interfaces demux0
Interface               Admin Link Proto    Local                Remote
demux0.2147516420       up    up   inet     0.0.0.0            --> 0/0
demux0.3221225474       up    up   inet     20.0.0.1           --> 0/0
demux0.3221225475       up    up   inet     20.0.0.1           --> 0/0
```

show user-plane interfaces interface-set queue

IN THIS SECTION

Syntax

```
show user-plane interfaces interface-set queue interface-set-name
```

Description

Displays information about the interface set queue for BNG User Planes.

Options

Following are all possible options for the `show user-plane interfaces interface-set queue` command. The options that appear may vary depending on the platform and software release.

interface-set-name (Optional) Display information about the specified interface set.

Required Privilege Level

view

Output Fields

Table 52 on page 498 lists all possible fields for the `show user-plane interfaces interface-set queue` command. The fields that appear may vary depending on the platform and software release.

Table 52: Ethernet show user-plane interfaces interface-set queue Output Fields

Field Name	Field Description
Egress supported	Total number of egress queues supported on the specified interface set.

Table 52: Ethernet show user-plane interfaces interface-set queue Output Fields *(Continued)*

Field Name	Field Description
Egress in use	Total number of egress queues used on the specified interface set.
Queue	Egress or ingress queue number for the statistics being displayed.
Forwarding classes	Forwarding class name for the statistics being displayed.
Queued	**Packet** and **Byte** statistics for the specified queue. • **Packets**—Number of packets queued and input rate in packets per second. • **Bytes**—Number of bytes queued and input rate in bytes per second.
Transmitted	**Packet** and **Byte** statistics for the specified forwarding class. • **Packets**—Number of packets transmitted and transmit rate in packets per second. • **Bytes**—Number of bytes transmitted and transmit rate in bytes per second. • **Tail-dropped packets**—Number of packets tail dropped. • **RED-dropped packets**—Number of RED-dropped packets for the **low**, **medium-low**, **medium-high**, and **high** loss priorities. • **RED-dropped bytes**—Number of RED-dropped bytes for the **low**, **medium-low**, **medium-high**, and **high** loss priorities.

Sample Output

show user-plane interfaces interface-set queue

```
user@host> show user-plane interfaces interface-set queue bb0.1-1
Egress: 8 supported, 8 in use
Queue: 0, Forwarding classes: q0
  Queued:
    Packets              :               1720368                  20101 pps
```

```
    Bytes                :            2064441600             192982528 bps
  Transmitted:
    Packets              :                 61143                   537 pps
    Bytes                :              73371600               5161072 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1659225                 19564 pps
    RED-dropped bytes    :            1991070000             187821456 bps
Queue: 1, Forwarding classes: q1
  Queued:
    Packets              :               1656425                 21295 pps
    Bytes                :            1987710000             204451656 bps
  Transmitted:
    Packets              :                 98284                  1077 pps
    Bytes                :             117940800              10351808 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1558141                 20218 pps
    RED-dropped bytes    :            1869769200             194099848 bps
Queue: 2, Forwarding classes: q2
  Queued:
    Packets              :               1684704                 20284 pps
    Bytes                :            2021644800             194733280 bps
  Transmitted:
    Packets              :                 58703                   537 pps
    Bytes                :              70443600               5161280 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1626001                 19747 pps
    RED-dropped bytes    :            1951201200             189572000 bps
Queue: 3, Forwarding classes: q3
  Queued:
    Packets              :               1621076                 21464 pps
    Bytes                :            1945291200             206063128 bps
  Transmitted:
    Packets              :                 96426                  1075 pps
    Bytes                :             115711200              10322152 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1524650                 20389 pps
    RED-dropped bytes    :            1829580000             195740976 bps
Queue: 4, Forwarding classes: q4
  Queued:
    Packets              :               1600565                 19928 pps
    Bytes                :            1920678000             191320352 bps
  Transmitted:
    Packets              :                 52963                   539 pps
```

```
    Bytes                :          63555600           5182112 bps
    Tail-dropped packets :                 0                 0 pps
    RED-dropped packets  :           1547602             19389 pps
    RED-dropped bytes    :        1857122400         186138240 bps
Queue: 5, Forwarding classes: q5
  Queued:
    Packets              :           1580649             20462 pps
    Bytes                :        1896778800         196441744 bps
  Transmitted:
    Packets              :             51874               537 pps
    Bytes                :          62248800           5161280 bps
    Tail-dropped packets :                 0                 0 pps
    RED-dropped packets  :           1528775             19925 pps
    RED-dropped bytes    :        1834530000         191280464 bps
Queue: 6, Forwarding classes: q6
  Queued:
    Packets              :           1560936             20952 pps
    Bytes                :        1873123200         201146120 bps
  Transmitted:
    Packets              :             51111               537 pps
    Bytes                :          61333200           5161072 bps
    Tail-dropped packets :                 0                 0 pps
    RED-dropped packets  :           1509825             20415 pps
    RED-dropped bytes    :        1811790000         195985048 bps
Queue: 7, Forwarding classes: q7
  Queued:
    Packets              :           1540270             20769 pps
    Bytes                :        1848324000         199391368 bps
  Transmitted:
    Packets              :             50478               539 pps
    Bytes                :          60573600           5182112 bps
    Tail-dropped packets :                 0                 0 pps
    RED-dropped packets  :           1489792             20230 pps
    RED-dropped bytes    :        1787750400         194209256 bps
```

show user-plane interfaces logical

IN THIS SECTION

Syntax

```
show user-plane interfaces logical physical-interface-name
<breif | detail | terse>
<forwarding class forwarding-class-name>
<queue physical-interface-name | forwarding class forwarding-class-name>
```

Description

Lists interface statistics for the specified BNG User Plane subscriber interfaces.

Options

`detail | extensive | terse`	(Optional) Displays the specified level of output.
`logical` *`physical-interface-name`*	Displays interface statistics for the specified BNG User Plane subscriber's interface name.
`forwarding class` *`forwarding-class-name`*	(Optional) Forwarding class name for this queue. Shows detailed statistics for the queue associated with the specified forwarding class.

Required Privilege Level

view

Output Fields

Table 53 on page 503 lists the output fields for the `show user-plane interfaces logical` command. Output fields are listed in the approximate order in which they appear.

Table 53: show user-plane interfaces logical Output Fields

Field Name	Field Description
`Logical interface`	Name of the logical interface.
`Flags`	Information about the interface. Possible values are described in the “Interface Flags” section under Common Output Fields Description
`State`	State of the logical interface (`up` or `down`).
`Session ID`	Session ID.
`Session AC Name`	Name of the access concentrator.
`Underlying interface`	Interface on which PPPoE is running.
`Ignor End-Of List tag`	Disables the `End-of-List` tag to continue processing of other tags after the `End-of-List` tag in a PPPoE Active Discovery Offer (PADO) packet.

Sample Output

show user-plane interfaces logical *physical-interface-name* brief

```
user@host> show user-plane interfaces logical pp0.3222274049 brief
Logical Interface pp0.3222274049
 Flags: Up Point-To-Point PPPOE
 PPPoE:
   State: SessionUp, Session ID: 1,
   Session AC name: cpi-ssen, Remote MAC address: 00:12:01:00:00:01,
   Underlying interface: demux0.3222274048 (Index 537395211)
   Ignore End-Of-List tag: Disable
       100.0.0.1 --> 0.0.0.0
```

show user-plane interfaces logical *physical-interface-name* terse

```
user@host> show user-plane interfaces logical pp0.3222274049 terse
Logical Interface pp0.3222274049
    Flags: Up Point-To-Point PPPOE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: cpi-ssen, Remote MAC address: 00:12:01:00:00:01,
      Underlying interface: demux0.3222274048 (Index 537395211)
      Ignore End-Of-List tag: Disable
           100.0.0.1 --> 0.0.0.0
```

show user-plane interfaces logical *physical-interface-name* detail

```
user@host> show user-plane interfaces logical pp0.3222274049 detail
Logical Interface pp0.3222274049 (Index 537395213) (SNMP ifIndex 200524301) (Generation 3)
    Flags: Up Point-To-Point PPPOE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: cpi-ssen, Remote MAC address: 00:12:01:00:00:01,
      Underlying interface: demux0.3222274048 (Index 537395211)
```

```
    Ignore End-Of-List tag: Disable
Egress: 8 supported, 8 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                                0                    0                    0
  1                                0                    0                    0
  2                                0                    0                    0
  3                                0                    0                    0
  4                                0                    0                    0
  5                                0                    0                    0
  6                                0                    0                    0
  7                                0                    0                    0
  Traffic statistics:
   Input  bytes  :                  40
   Output bytes  :                  36
   Input  packets:                   1
   Output packets:                   2
  Local traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit traffic statistics:
   Input  bytes  :                  40                    0 bps
   Output bytes  :                  36                    0 bps
   Input  packets:                   1                    0 bps
   Output packets:                   2                    0 bps
  Protocol inet, MTU: 1470
  Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
  Route table: 0
    Flags: Unnumbered
    Donor interface: lo0.0 (Index 322)
    Addresses, Flags: Is-Primary
    Destination: Unspecified, Local: 100.0.0.1, Broadcast: Unspecified, Generation: 3
```

show user-plane interfaces queue

IN THIS SECTION

Syntax

```
show user-plane interfaces queue physical-interface-name
<forwarding class forwarding-class-name>
```

Description

Lists queue information for the specified BNG User Plane subscriber's interface name.

Options

none	Show detailed queue statistics for all physical interfaces.
`forwarding class` *`forwarding-class-name`*	(Optional) Forwarding class name for this queue. Shows detailed statistics for the queue associated with the specified forwarding class.

Required Privilege Level

root

Output Fields

Table 54 on page 507 lists the output fields for the `show user-plane interfaces queue` command. Output fields are listed in the approximate order in which they appear.

Table 54: show user-plane interfaces queue Output Fields

Field Name	Field Description
`Queue`	Queue number.
`Forwarding classes`	Forwarding class name.
`Queued Packets`	Number of packets queued to this queue. For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see *Additional Information*.
`Queued Bytes`	Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see . For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see *Additional Information*.
`Transmitted Packets`	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the `Packet Forwarding Engine Chassis Queues` field) shows the prefragmentation values.

Table 54: show user-plane interfaces queue Output Fields *(Continued)*

Field Name	Field Description
`Transmitted Bytes`	Number of bytes transmitted by this queue. The byte counts vary by interface hardware. . **NOTE**: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation. **NOTE**: For Layer 2 statistics, see *Additional Information*
`Tail-dropped packets`	Number of packets dropped because of tail drop.
`RL-dropped packets`	Number of packets dropped due to rate limiting. For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see *Additional Information*.
`RL-dropped bytes`	Number of bytes dropped due to rate limiting. For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see *Additional Information*.

`RED-dropped packets`	Number of packets dropped because of random early detection (RED).
	
	• The output classifies dropped packets into the following categories:
	
	  • `Low`—Number of low-loss priority packets dropped because of RED.
	
	  • `Medium-low`—Number of medium-low loss priority packets dropped because of RED.
	
	  • `Medium-high`—Number of medium-high loss priority packets dropped because of RED.
	
	  • `High`—Number of high-loss priority packets dropped because of RED.

Table 54: show user-plane interfaces queue Output Fields *(Continued)*

Field Name	Field Description

`RED-dropped bytes`	Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see .
	• The output classifies dropped bytes into the following categories:
	  • `Low, non-TCP`—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • `Low, TCP`—Number of low-loss priority TCP bytes dropped because of RED.
	  • `High, non-TCP`—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • `High, TCP`—Number of high-loss priority TCP bytes dropped because of RED.

`ECN-CE packets`	The number of packets that experienced congestion, either on a specific queue or on another device prior to being received.
`ECN-CE bytes`	The number of bytes that experienced congestion, either on a specific queue or on another device prior to being received.
`Queue Buffer Usage:`	• `Reserved buffer`—The size of the memory buffer that is allocated for storing packets • `Current`—The amount of buffer memory that is currently in use on this queue.
`Queue-depth bytes`	Displays the amount of queue buffer that is in occupation at this instance. This is an indicator of the amount of data that is present in a queue at that point in time. The amount of data present is in the units of bytes.
`Peak`	Diplays the peak buffer occupancy for the queue while `buffer-monitor-enable` is enabled at the `[edit chassis fpc` *`slot-number`* `traffic-manager]` hierarchy level.
`Last-packet enqueued`	If `packet-timestamp` is enabled for an FPC, shows the day, date, time, and year in the format *day-of-the-week month day-date hh:mm:ss yyyy* when a packet was enqueued in the CoS queue. When the timestamp is aggregated across all active Packet Forwarding Engines, the latest timestamp for each CoS queue is reported.

Sample Output

show user-plane interfaces queue *physical-interface-name*

```
user@host> show user-plane interfaces queue pp0.3222274049
Egress: 8 supported, 8 in use
Queue: 0, Forwarding classes: q0
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                   0                    0 pps
    RL-dropped packets  :                    0                    0 pps
    RL-dropped bytes    :                    0                    0 bps
    RED-dropped packets :                    0                    0 pps
     Low                :                    0                    0 pps
     Medium-low         :                    0                    0 pps
     Medium-high        :                    0                    0 pps
     High               :                    0                    0 pps
    RED-dropped bytes   :                    0                    0 bps
     Low                :                    0                    0 bps
     Medium-low         :                    0                    0 bps
     Medium-high        :                    0                    0 bps
     High               :                    0                    0 bps
Queue: 1, Forwarding classes: q1
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                   0                    0 pps
    RL-dropped packets  :                    0                    0 pps
    RL-dropped bytes    :                    0                    0 bps
    RED-dropped packets :                    0                    0 pps
     Low                :                    0                    0 pps
     Medium-low         :                    0                    0 pps
     Medium-high        :                    0                    0 pps
```

```
      High              :                    0                    0 pps
    RED-dropped bytes   :                    0                    0 bps
      Low               :                    0                    0 bps
      Medium-low        :                    0                    0 bps
      Medium-high       :                    0                    0 bps
      High              :                    0                    0 bps
Queue: 2, Forwarding classes: q2
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                   0                    0 pps
    RL-dropped packets  :                    0                    0 pps
    RL-dropped bytes    :                    0                    0 bps
    RED-dropped packets :                    0                    0 pps
      Low               :                    0                    0 pps
      Medium-low        :                    0                    0 pps
      Medium-high       :                    0                    0 pps
      High              :                    0                    0 pps
    RED-dropped bytes   :                    0                    0 bps
      Low               :                    0                    0 bps
      Medium-low        :                    0                    0 bps
      Medium-high       :                    0                    0 bps
      High              :                    0                    0 bps
Queue: 3, Forwarding classes: q3
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets :                   0                    0 pps
    RL-dropped packets  :                    0                    0 pps
    RL-dropped bytes    :                    0                    0 bps
    RED-dropped packets :                    0                    0 pps
      Low               :                    0                    0 pps
      Medium-low        :                    0                    0 pps
      Medium-high       :                    0                    0 pps
      High              :                    0                    0 pps
    RED-dropped bytes   :                    0                    0 bps
      Low               :                    0                    0 bps
```

```
     Medium-low           :                    0                     0 bps
     Medium-high          :                    0                     0 bps
     High                 :                    0                     0 bps
Queue: 4, Forwarding classes: q4
  Queued:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
  Transmitted:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets    :                    0                     0 pps
    RL-dropped bytes      :                    0                     0 bps
    RED-dropped packets   :                    0                     0 pps
     Low                  :                    0                     0 pps
     Medium-low           :                    0                     0 pps
     Medium-high          :                    0                     0 pps
     High                 :                    0                     0 pps
    RED-dropped bytes     :                    0                     0 bps
     Low                  :                    0                     0 bps
     Medium-low           :                    0                     0 bps
     Medium-high          :                    0                     0 bps
     High                 :                    0                     0 bps
Queue: 5, Forwarding classes: q5
  Queued:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
  Transmitted:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets    :                    0                     0 pps
    RL-dropped bytes      :                    0                     0 bps
    RED-dropped packets   :                    0                     0 pps
     Low                  :                    0                     0 pps
     Medium-low           :                    0                     0 pps
     Medium-high          :                    0                     0 pps
     High                 :                    0                     0 pps
    RED-dropped bytes     :                    0                     0 bps
     Low                  :                    0                     0 bps
     Medium-low           :                    0                     0 bps
     Medium-high          :                    0                     0 bps
     High                 :                    0                     0 bps
```

```
Queue: 6, Forwarding classes: q6
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RL-dropped packets   :                    0                     0 pps
    RL-dropped bytes     :                    0                     0 bps
    RED-dropped packets  :                    0                     0 pps
     Low                 :                    0                     0 pps
     Medium-low          :                    0                     0 pps
     Medium-high         :                    0                     0 pps
     High                :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
     Low                 :                    0                     0 bps
     Medium-low          :                    0                     0 bps
     Medium-high         :                    0                     0 bps
     High                :                    0                     0 bps
Queue: 7, Forwarding classes: q7
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RL-dropped packets   :                    0                     0 pps
    RL-dropped bytes     :                    0                     0 bps
    RED-dropped packets  :                    0                     0 pps
     Low                 :                    0                     0 pps
     Medium-low          :                    0                     0 pps
     Medium-high         :                    0                     0 pps
     High                :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
     Low                 :                    0                     0 bps
     Medium-low          :                    0                     0 bps
     Medium-high         :                    0                     0 bps
     High                :                    0                     0 bps
```

show user-plane subscribers

IN THIS SECTION

Syntax

```
show user-plane subscribers
<detail | extensive | terse>
<accounting-statistics>
<address ip-address>
<client-type client-type>
<count>
<id subscriber-id>
<interface interface-name>
<license-usage>
<mac-address mac-address>
<physical-interface physical-interface-name>
<routing-instance routing-instance-name>
<stacked-vlan-id stacked-vlan-id>
<subscriber-state subscriber-state>
<Summary>
<vlan-id vlan-id>
```

Description

Displays information for the BNG User Plane subscribers.

Options

None	Displays information by interface, IP address or VLAN ID, and LS:RI.
`detail | extensive | terse`	(Optional) Displays the specified level of output.
`accounting-statistics`	Displays subscriber accounting statistics. If you do not configure the statement, the CLI displays a value of 0 for the accounting statistics.
`address` *ip-address*	(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid and rejects the command.
`client-type` *client-type*	(Optional) Display subscribers whose client type matches one of the following client types: • `dhcp`—Dynamic Host Configuration Protocol (DHCP) clients only. • `dot1x`—802.1X clients only. • `essm`—Extensible Subscribers Services Manager (ESSM) clients only. • `fixed-wireless-access`—Fixed wireless access clients only. • `fwauth`—FwAuth (authenticated across a firewall) clients only. • `l2tp`—Layer 2 Tunneling Protocol (L2TP) clients only. • `mlppp`—Multilink Point-to-Point Protocol (MLPPP) clients only. • `ppp`—Point-to-Point Protocol (PPP) clients only. • `pppoe`—Point-to-Point Protocol over Ethernet (PPPoE) clients only. • `static`—Static clients only.

- `vlan`—VLAN clients only.
- `vlan-oob`—VLAN out-of-band (triggered by by Access Node Control Protocol or ANCP) clients only.
- `vpls-pw`—Virtual private LAN service (VPLS) pseudowire clients only.
- `xauth`—Extended Authentication (XAuth) clients only.

`count`	(Optional) Displays the count of total subscribers and active subscribers for any specified option.
`id` ***`subscriber-id`***	(Optional) Displays a subscriber of a session with an ID matching the specific subscriber ID.
`interface` ***`interface-name`***	(Optional) Displays subscribers with an interface that matches the specified interface.
`license-usage`	(Optional) Displays the license usage for the last 48 hours.
`mac-address` ***`mac-address`***	(Optional) Displays subscribers with a MAC address that matches the specified MAC address.
`physical-interface` ***`physical-interface-name`***	(Optional) Displays subscribers whose physical interface matches the specified physical interface.
`routing-instance` ***`routing-instance-name`***	(Optional) Displays subscribers with a routing instance that matches the specified routing instance.
`stacked-vlan-id` ***`stacked-vlan-id`***	(Optional) Displays subscribers with a stacked VLAN ID that matches the specified stacked VLAN ID.
`subscriber-state` ***`subscriber-state`***	(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).
`summary`	(Optional) Displays summary information for subscribers.
`vlan-id` ***`vlan-id`***	(Optional) Displays subscribers that have a VLAN ID that matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the `stacked-vlan-id` *`stacked-vlan-id`* option to match the outer VLAN tag.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Output Fields

Table 55 on page 517 lists the output fields for the `show user-plane subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 55: show user-plane subscribers Output Fields

Field Name	Field Description
`Interface`	Interface associated with the subscriber. The router or switch displays subscribers with interfaces that match or begin with the specified interface. The asterisk (*) character indicates a continuation of addresses for the same session.
`IP Address/VLAN ID`	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*. No IP address or VLAN ID is assigned to a Layer 2 Tunneling Protocol (L2TP) tunnel-switched session. For these subscriber sessions the value is `Tunnel-switched`.
`LS:RI`	Logical system and routing instance associated with the subscriber.
`Type`	Subscriber client type (DHCP, L2TP, PPP, PPPoE, and VLAN).
`IP Address`	Subscriber IPv4 address.
`Logical System`	Logical system associated with the subscriber.

Table 55: show user-plane subscribers Output Fields *(Continued)*

Field Name	Field Description
`Routing Instance`	Routing instance associated with the subscriber.
`Interface`	Name of the enhanced subscriber management logical interface, in the form `demux0.`*nnnn* (for example, `demux0.3221225472`), to which access-internal and framed subscriber routes are mapped.
`Interface Type`	Whether the subscriber interface is `Static` or `Dynamic`.
`Underlying Interface`	Name of the underlying interface for the subscriber session.
`MAC Address`	MAC address associated with the subscriber.
`State`	Current state of the subscriber session (`Init`, `Configured`, `Active`, `Terminating`, `Tunneled`).
`PFE Flow ID`	Forwarding flow identifier.
`VLAN ID`	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
`Session ID`	ID number for a subscriber session.
`Service Sessions`	Number of service sessions (that is, a service that is activated using RADIUS CoA) associated with the subscribers.
`Service Session ID`	ID number for a subscriber service session.
`Service Session Name`	Service session profile name.
`IPv4 Input Filter Name`	Name assigned to the IPv4 input filter (client or service session).

Table 55: show user-plane subscribers Output Fields *(Continued)*

Field Name	Field Description
`IPv4 Output Filter Name`	Name assigned to the IPv4 output filter (client or service session).
`Service Activation Time`	Time at which the first family in this service became active.
`Accounting Interval Service`	How often statistics are collected for the accounting profile.
`Dynamic Configuration`	Values for variables that are passed into the dynamic profile from RADIUS.
`Subscribers by State`	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscribers currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers in the process of being terminated. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.
`Subscribers by Client Type`	Number of subscribers, summarized by client type. Client types can include DHCP, L2TP, PPP, PPPoE, and VLAN. Also displays the total number of subscribers for all client types.
`Subscribers by LS:RI`	Number of subscribers, summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations.
`Redundancy Service Id`	The redundancy service ID number.

Sample Output

show user-plane subscribers

```
user@host> show user-plane subscribers
Interface             IP Address/VLAN ID                                 LS:RI
demux0.3222276148      0x8100.210 0x8100.209                        default:default
demux0.3222276149      0x8100.210 0x8100.210                        default:default
demux0.3222276150      0x8100.210 0x8100.211                        default:default
demux0.3222276151      0x8100.210 0x8100.212                        default:default
demux0.3222276152      0x8100.210 0x8100.213                        default:default
demux0.3222276153      0x8100.210 0x8100.214                        default:default
demux0.3222276154      0x8100.210 0x8100.215                        default:default
demux0.3222276155      0x8100.210 0x8100.216                        default:default
demux0.3222276156      0x8100.210 0x8100.217                        default:default
demux0.3222276157      0x8100.210 0x8100.218                        default:default
demux0.3222276158     192.168.10.254                                default:default
demux0.3222276159     192.168.10.252                                default:default
demux0.3222276160     192.168.10.253                                default:default
demux0.3222276161     192.168.10.248                                default:default
demux0.3222276162     192.168.10.249                                default:default
demux0.3222276163     192.168.10.250                                default:default
```

show user-plane subscribers count

```
user@host> show user-plane subscribers count
Total Subscribers: 1, Active Subscribers: 1
```

show user-plane subscribers detail

```
user@host> show user-plane subscribers detail
Oct 15 11:03:30
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221749760
Interface type: Dynamic
```

```
Underlying Interface: ps1.32767
State: Active
Session ID: 43
PFE Flow ID: 537395225
VLAN Id: 1
Stacked VLAN Id: 100

Type: DHCP
IP Address: 5.0.0.102
Logical System: default
Routing Instance: default
Interface: demux0.3221749761
Interface type: Dynamic
Underlying Interface: demux0.3221749760
MAC Address: 00:21:01:00:00:01
State: Active
Session ID: 44
PFE Flow ID: 537395227
VLAN Id: 1
Stacked VLAN Id: 100
SGRP ID: 3

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221749762
Interface type: Dynamic
Underlying Interface: ps1.32767
State: Active
Session ID: 45
PFE Flow ID: 537395228
VLAN Id: 100
Stacked VLAN Id: 100

Type: PPPoE
IP Address: 5.0.0.103
Logical System: default
Routing Instance: default
Interface: pp0.3221749763
Interface type: Dynamic
Underlying Interface: demux0.3221749762
MAC Address: 00:23:01:00:00:01
State: Active
```

```
Session ID: 46
PFE Flow ID: 537395229
VLAN Id: 100
Stacked VLAN Id: 100
SGRP ID: 3
```

show user-plane subscribers client-type (PPPoE)

```
user@host> show user-plane subscribers client-type pppoe
Oct 15 11:03:36
Interface             IP Address/VLAN ID                              LS:RI
pp0.3221749763        5.0.0.103                                  default:default
```

show user-plane subscribers client-type (DHCP)

```
user@host> show user-plane subscribers client-type dhcp
Oct 15 11:03:39
Interface             IP Address/VLAN ID                              LS:RI
demux0.3221749761     5.0.0.102                                  default:default
```

show user-plane subscribers extensive

```
user@host> show user-plane subscribers extensive
Oct 15 11:03:42
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221749760
Interface type: Dynamic
Underlying Interface: ps1.32767
State: Active
Session ID: 43
PFE Flow ID: 537395225
VLAN Id: 1
Stacked VLAN Id: 100

Type: DHCP
IP Address: 5.0.0.102
```

```
Logical System: default
Routing Instance: default
Interface: demux0.3221749761
Interface type: Dynamic
Underlying Interface: demux0.3221749760
MAC Address: 00:21:01:00:00:01
State: Active
Session ID: 44
PFE Flow ID: 537395227
VLAN Id: 1
Stacked VLAN Id: 100
SGRP ID: 3

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221749762
Interface type: Dynamic
Underlying Interface: ps1.32767
State: Active
Session ID: 45
PFE Flow ID: 537395228
VLAN Id: 100
Stacked VLAN Id: 100

Type: PPPoE
IP Address: 5.0.0.103
Logical System: default
Routing Instance: default
Interface: pp0.3221749763
Interface type: Dynamic
Underlying Interface: demux0.3221749762
MAC Address: 00:23:01:00:00:01
State: Active
Session ID: 46
PFE Flow ID: 537395229
VLAN Id: 100
Stacked VLAN Id: 100
SGRP ID: 3
```

show user-plane subscribers interface *interface-name* extensive

```
user@host> show user-plane subscribers interface interface-name extensive
Type: PPPoE
IP Address: 192.168.1.199
Logical System: default
Routing Instance: default
Interface: pp0.3221749860
Interface type: Dynamic
Underlying Interface: demux0.3221749760
MAC Address: 00:00:64:08:01:65
State: Active
Session ID: 111
PFE Flow ID: 537395323
VLAN Id: 1000
Stacked VLAN Id: 1000
SGRP ID: 2048
Redundancy Service Id:
```