Engineering
Simplicity

Juniper BNG CUPS User Guide

Published
2024-08-15

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

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

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

4

Juniper BNG CUPS CLI Operational Statements

About This Guide

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

1

CHAPTER

Overview

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 clusters 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.

Figure 1 on page 2 shows the Juniper BNG CUPS architecture.

Figure 1: Juniper BNG CUPS Solution

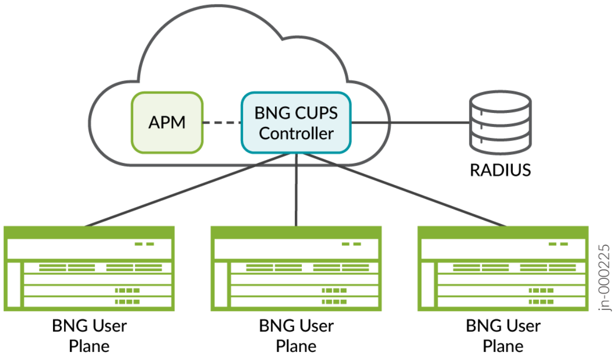

Benefits of Juniper BNG CUPS

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.

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.

Juniper BNG CUPS use cases:

- 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 two 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's practically undetectable to the users.

  - Warm 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 is typically used to support N:1 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.

- 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
- Auto-sensed 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-related commands run on the BNG CUPS Controller. To help with troubleshooting, some of the operational commands run on the BNG User Planes.

The Juniper BNG CUPS CLI uses a slightly different layout from the traditional commands used for integrated BNG Junos OS CLI. The goal of the Juniper BNG CUPS CLI is to reduce the need for you to understand where information is kept. At a high level, the subscriber management show commands are either subsystem based or object based. Figure 2 on page 7 shows the Juniper BNG CUPS CLI hierarchy for the show commands.

Figure 2: Juniper BNG CUPS CLI Hierarchy

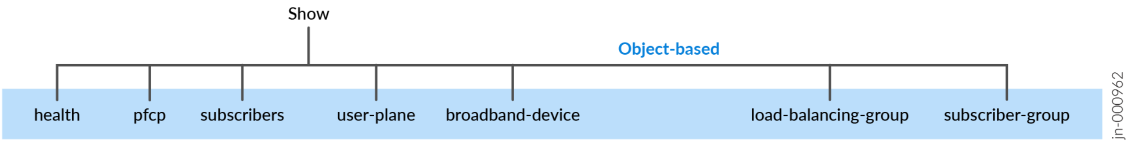

Below is a summary list of the functional components and where you run their operations.

> **NOTE**: For a complete set of commands, see the *Juniper BNG CUPS User Guide*.

You run operational commands for the following functional components on the BNG CUPS Controller:

- Accounting

- Agent
- Broadband device
- Subscriber groups
- Load balancing groups
- Health
- User Plane
- Subscriber session state
- Node management
- Routing instance
- Services
- DHCP/DHCPv6 local server and relay
- PPPoE and PPP
- L2TP
- Dynamic auto-sensed VLANs
- AAA

You run operational commands for the following functional components on the BNG User Plane:

- Node management
- Subscriber management

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 16 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

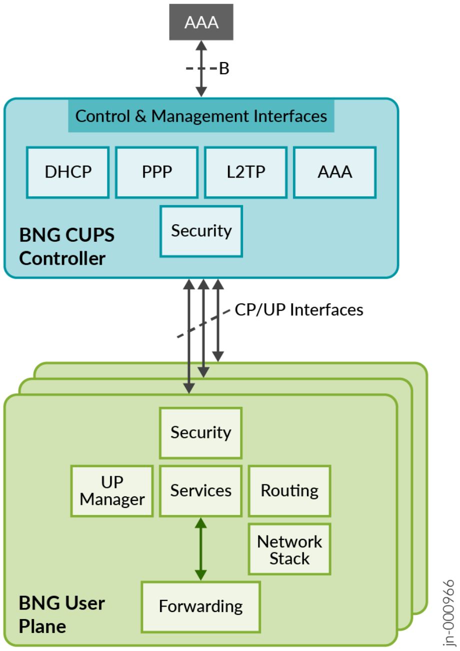

Smart Session Load Balancing

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 13 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

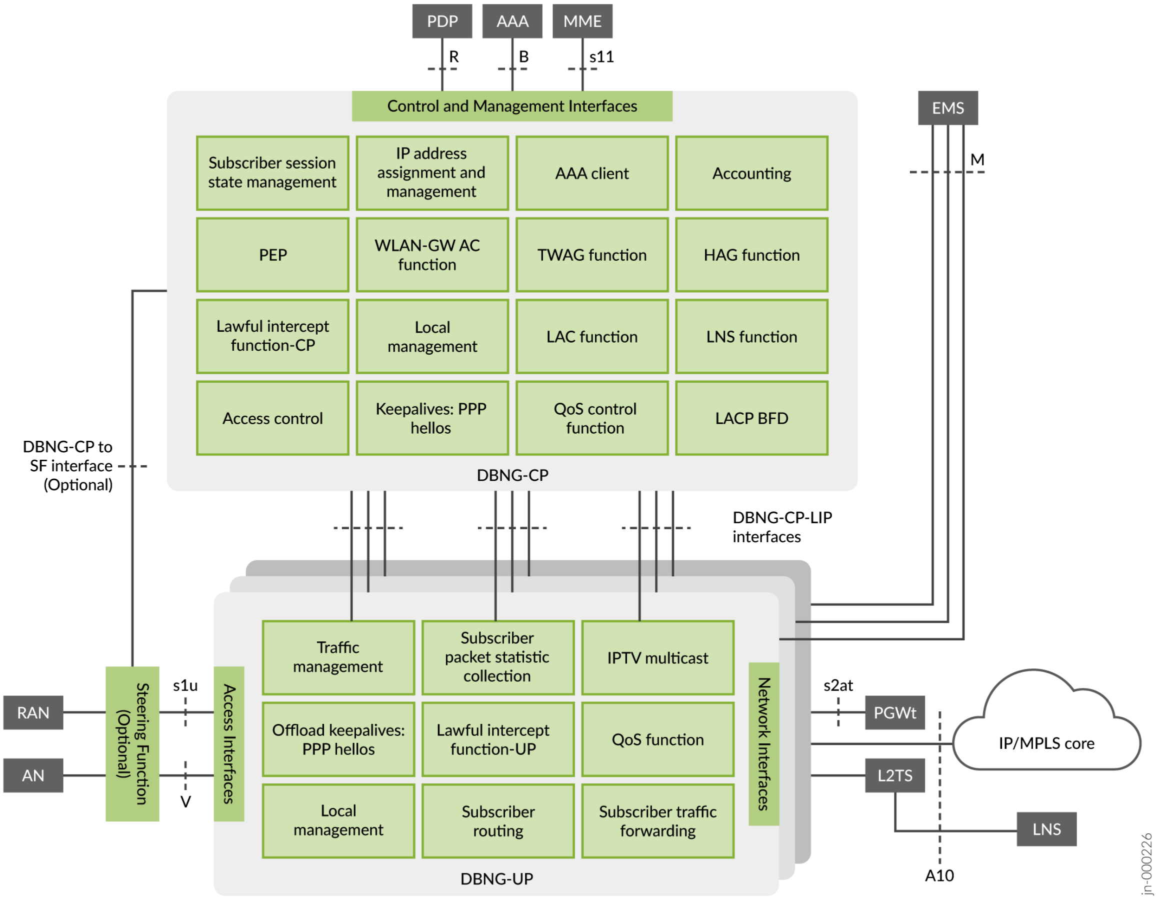

The combination of the control plane functions is referred to as a control plane of the disaggregated BNG. Similarly, a combination of the user plane specific functions is referred to as a user plane of the disaggregated BNG.

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 16 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 No Link Title.

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-common-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 17 shows the configuration group hierarchy.

Figure 5: Configuration Group Hierarchy

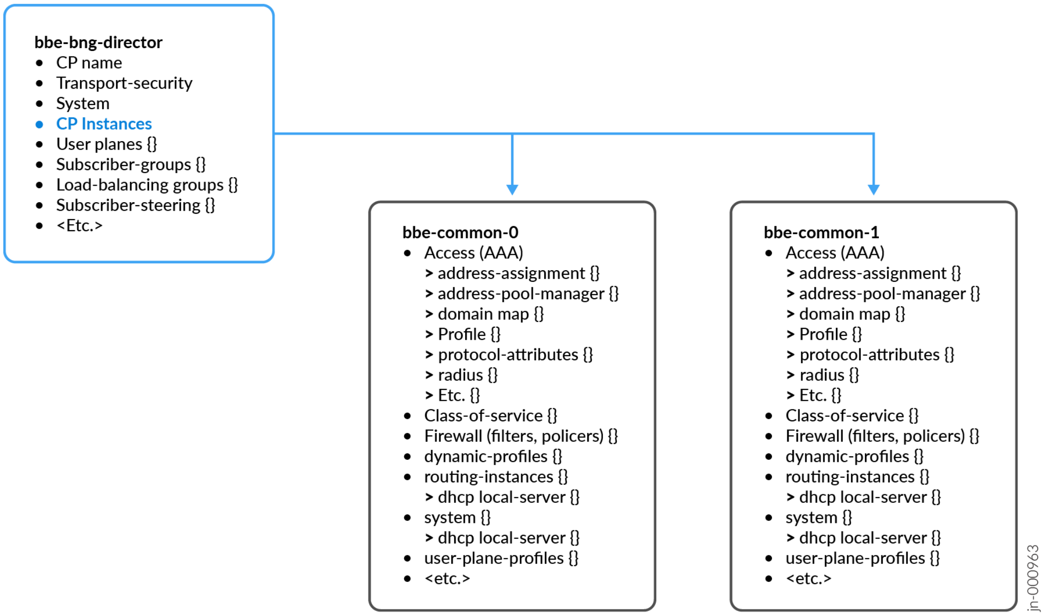

Configure the bbe-bng-director Group

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-common-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-common-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` for the control plane instance.

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-common-0`) to use for subscriber management configuration.

- The list of BNG User Planes assigned to the control plane instance.

Configure the bbe-common-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-common-0, bbe-common-1, bbe-common-2, bbe-common-3, or bbe-common-4.

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

```
groups bbe-common-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.

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 {
                    control-plane-name cpi-boston;
                }
            }
        }
    }
```

NOTE: Also, you will need to perform a similar configuration for 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, 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 the 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 <cluster- context>
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 <cluster-context> --version 23.4R2 cpi-test-1
```

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 <cluster-context>
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 <cluster-context>
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

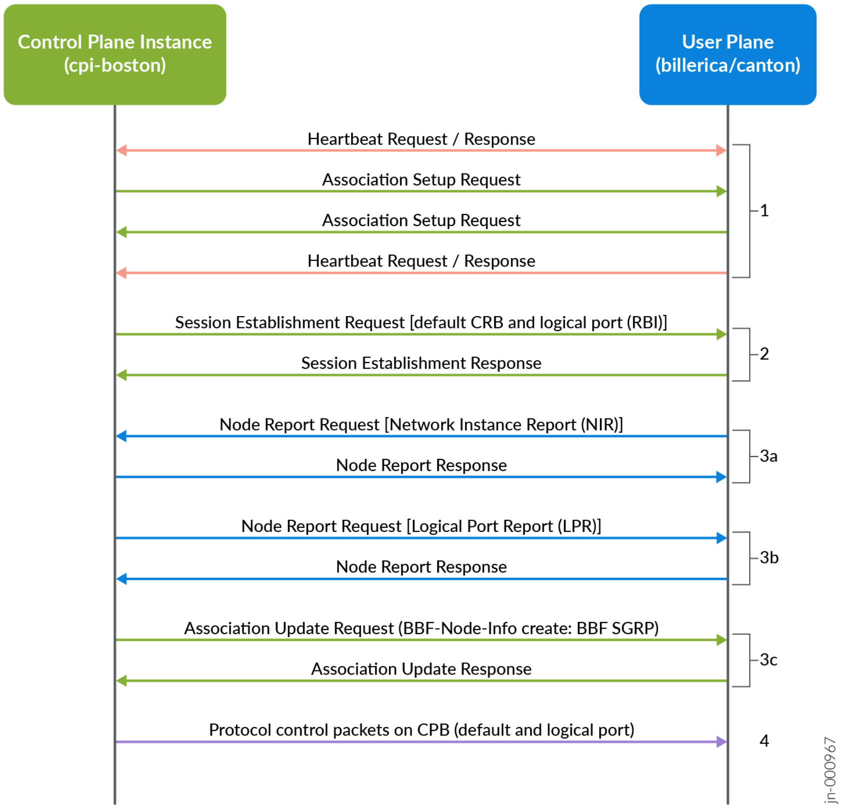
Control Plane Instance
(cpi-boston)
User Plane
(billerica/canton)
Heartbeat Request / Response
Association Setup Request
Association Setup Request
Heartbeat Request / Response
1
Session Establishment Request [default CRB and logical port (RBI)]
Session Establishment Response
2
Node Report Request [Network Instance Report (NIR)]
Node Report Response
3a
Node Report Request [Logical Port Report (LPR)]
Node Report Response
3b
Association Update Request (BBF-Node-Info create: BBF SGRP)
Association Update Response
3c
Protocol control packets on CPB (default and logical port)
4
jn-000967

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

   ```
   root@cpi-boston> show 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

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-common-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-common-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 7 on page 34 shows the BNG CUPS Controller container and the State Cache container.

Figure 7: BNG CUPS Controller High Availability

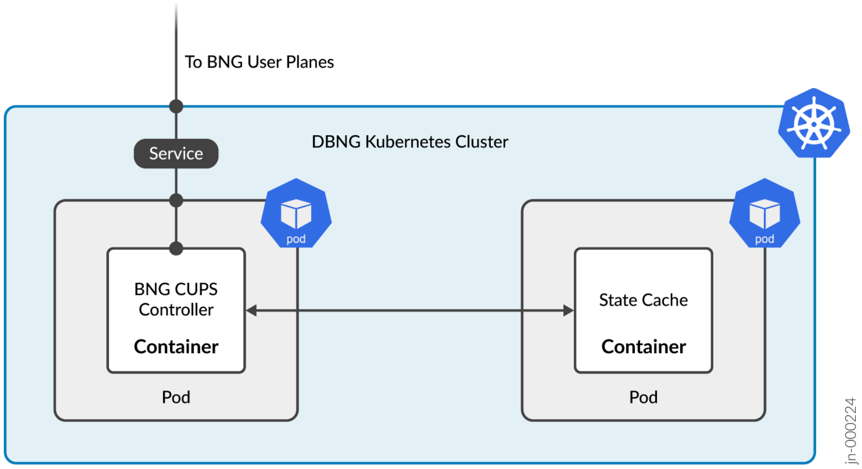

BNG User Planes High Availability

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 35.

Figure 8 on page 35 shows the BNG User Plane high availability andwith GRES support between RE0 and RE1.

Figure 8: BNG User Plane High Availability and GRES

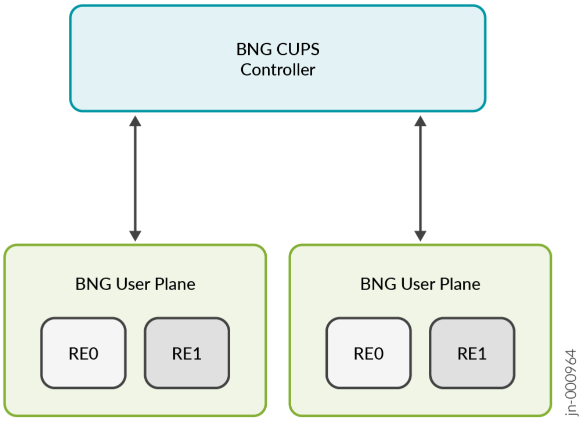

Use Juniper BNG CUPS Subscriber Groups

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 9 on page 36 (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 9: Subscriber Session Resiliency

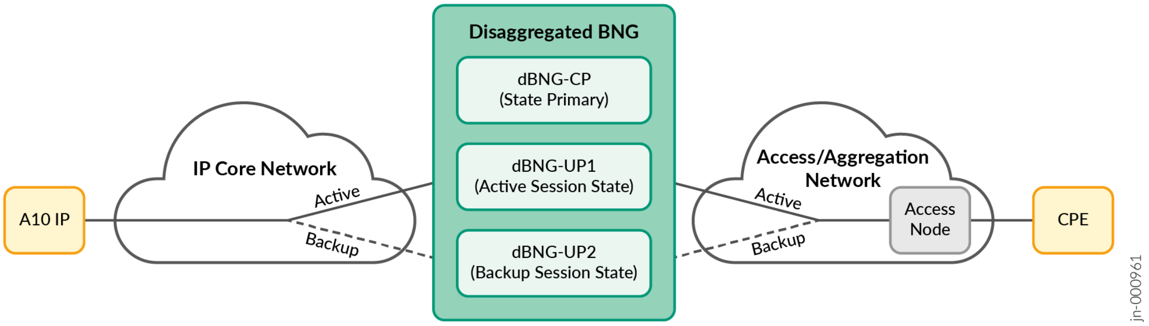

Subscriber sessions that are subject to the same restoration capability are placed into the same subscriber group. Grouping subscribers together helps to increase core routing efficiency.

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:

```
[edit groups bbe-bng-director bng-controller]
subscriber-groups {
   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;
     }
  }
}
```

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-groups {
   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 user 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.

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 10 on page 43 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 10: Hitless BNG User Plane Maintenance

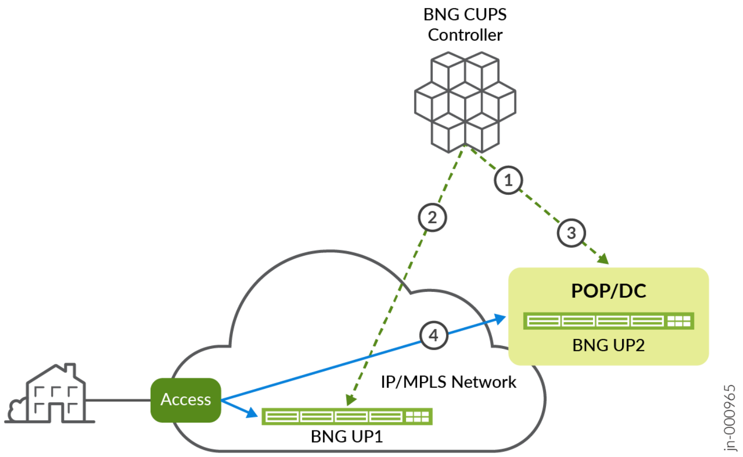

Following are the steps that occur when performing maintenance on a BNG User Plane (see Figure 10 on page 43.

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. BNG User Plane *BNG-UP2* is now preprogrammed with the subscriber state information from BNG User Plane *BNG-UP1* as a hot standby BNG User Plane. 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 10 on page 43.

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 11 on page 47 shows how BNG CUPS Controller implements subscriber load balancing.

Figure 11: Juniper BNG CUPS Subscriber Load Balancing

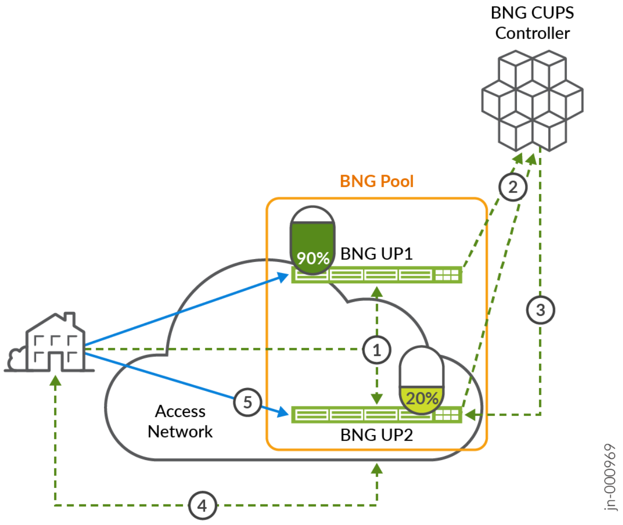

Following is the work flow that Juniper BNG CUPS uses for subscriber load balancing (see Figure 11 on page 47).

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 12 on page 50, 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 12: Load Balancing Combined with Subscriber Session Steering

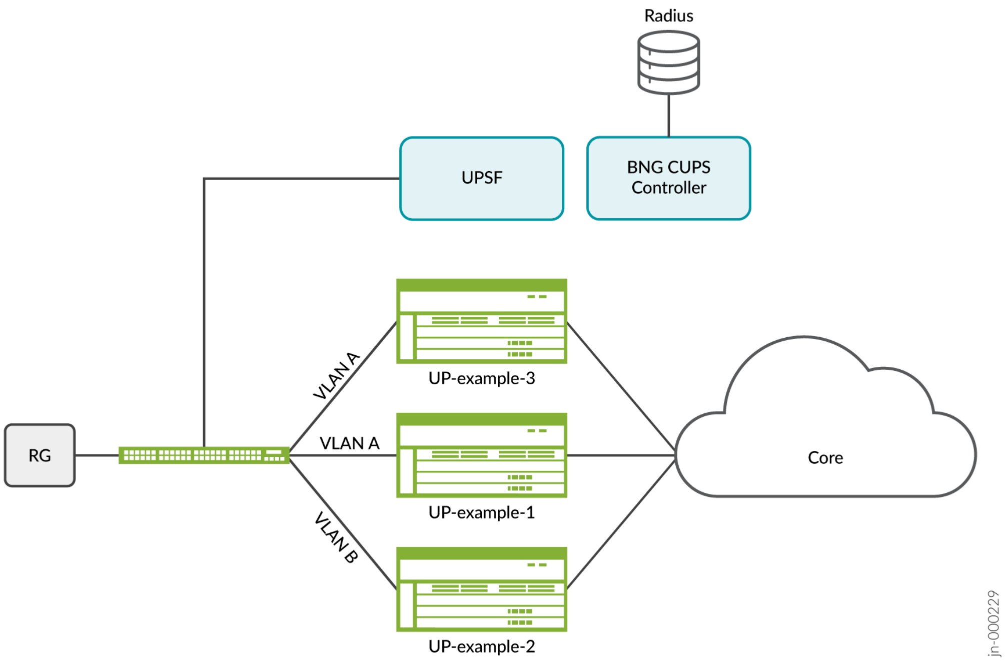

For weight-based load balancing, you use the dynamic profile configuration to specify subscriber weight or logical interface set weight.

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, we would 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 13 on page 56 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 13: WT-474 Subscriber Session Steering Architecture

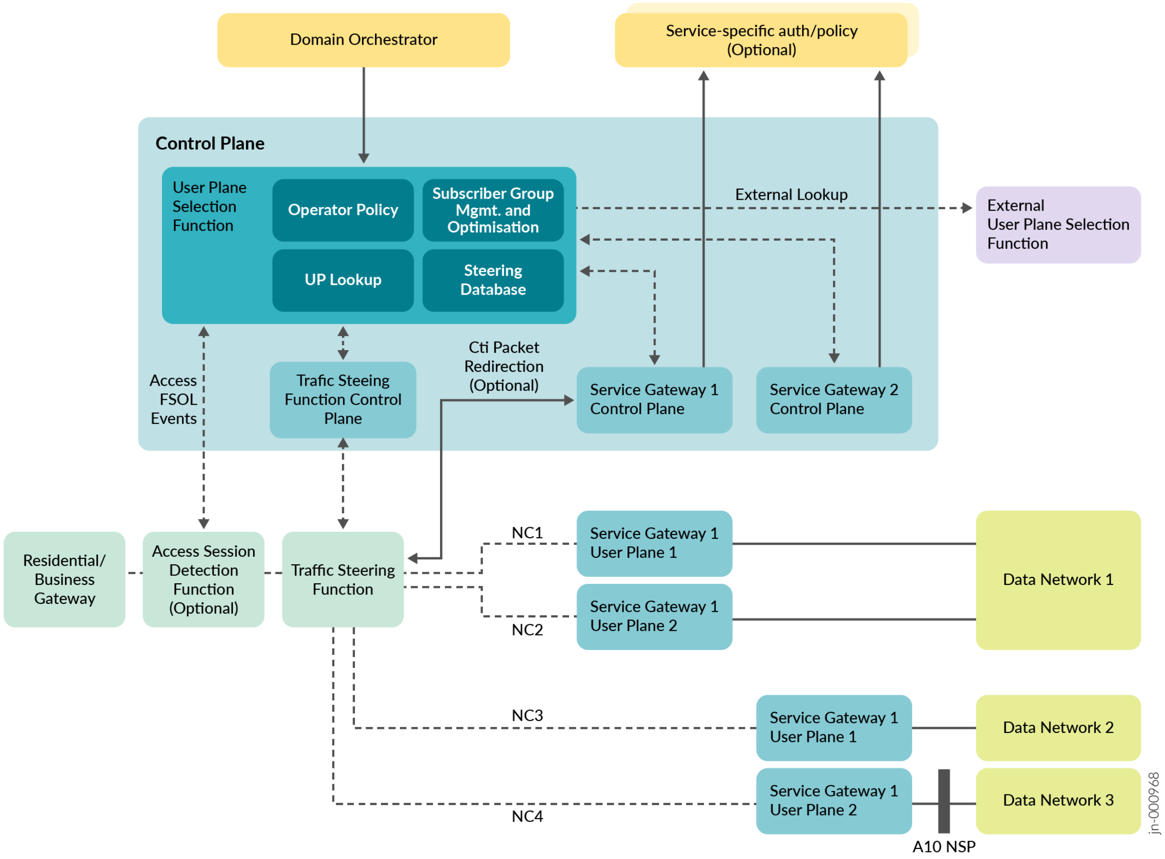

Following are new functions that are listed in the WT474 architecture:

- 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 14 on page 59.)

Figure 14: Subscriber Session Steering

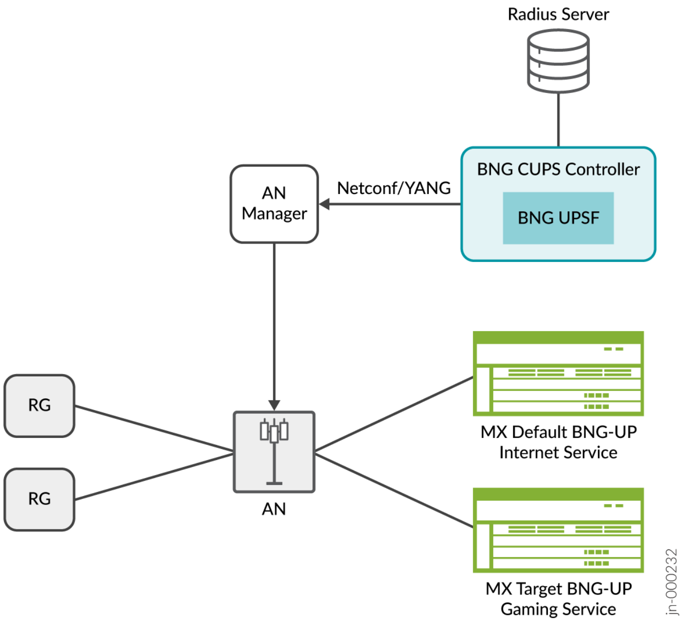

The default BNG User Plane is the ingress BNG User Plane for the subscriber login control packets. The target BNG User Plane is where user plane selection function places the subscriber. Depending on the network architecture, the default BNG User Plane and the target BNG User Plane might be the same 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 in to *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 12 on page 50, 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*.

Configure Juniper BNG CUPS

SUMMARY

This document presents sample configurations that you can use to set up Juniper BNG CUPS and configure subscriber access and subscriber management.

IN THIS SECTION

Configure Multicast

You can set up multicast in your Juniper BNG CUPS environment. The following sections contain example configurations to help you set up multicast in Juniper BNG CUPS.

- "Configure Global Multicast Settings" on page 63
- "Configure Centralized Multicast" on page 65
- "Configure Distributed Multicast" on page 66

Configure Global Multicast Settings

You can use the following example configuration to help you configure global multicast settings.

Configure multicast on the BNG User Planes, as follows:

```
[edit]
 protocols {
    igmp {
        query-interval 125;
        query-response-interval 10;
        query-last-member-interval 1;
        robust-count 2;
    }
    mld {
        query-interval 125;
        query-response-interval 10;
        query-last-member-interval 1;
        robust-count 2;
    }
}
policy-options {
    policy-statement OIF-MAP-V4 {
        term A {
            from {
                route-filter 230.10.10.1/24 orlonger;
                route-filter 230.20.20.1/32 exact;
            }
            then {
                map-to-interface ge-1/0/1.33;
                accept;
            }
        }
        then reject;
    }
    policy-statement OIF-MAP-V6 {
        term A {
            from {
                route-filter ff3e:0:0:0:0:0:0:101/64 orlonger;
                route-filter ff05:230::1/128 exact;
            }
            then {
                map-to-interface ge-1/0/1.33;
                accept;
            }
        }
        then reject;
    }
    policy-statement igmp-group-policy {
        term A1 {
            from {
                route-filter 230.0.0.1/24 orlonger;
            }
            then accept;
        }
        then reject;
    }
    policy-statement mld-group-policy {
        term A1 {
            from {
                route-filter ff05::/64 orlonger;
            }
            then accept;
        }
        then reject;
    }
    policy-statement ssm-map-v4 {
        term A1 {
            from {
                route-filter 230.0.0.1/24 orlonger;
            }
            then {
                ssm-source 194.0.0.22;
                accept;
            }
        }
    }
    policy-statement ssm-map-v6 {
        term A1 {
            from {
                route-filter ff05::/64 orlonger;
            }
            then {
                ssm-source 3000::1;
                accept;
            }
        }
    }
}
routing-options {
    multicast {
        ssm-groups 233.0.0.0/8;
        cont-stats-collection-interval 600;
    }
}
```

Configure Centralized Multicast

You can use the following example configuration to help you configure a centralized multicast setup.

Configure centralized multicast on the BNG CUPS Controller, as follows:

```
[edit dynamic-profiles profile-name]
protocols {
    igmp {
        interface "$junos-interface-name" {
            version 3;
            immediate-leave;
            promiscuous-mode;
            ssm-map-policy ssm-map-v4;
            group-policy igmp-group-policy;
            oif-map OIF-MAP-V4;
        }
    }
}
```

Configure Distributed Multicast

You can use the following example configuration to help you configure a distributed multicast setup.

Configure distributed multicast on the BNG CUPS Controller, as follows:

```
[edit dynamic-profiles profile-name]
protocols {
    mld {
        interface "$junos-interface-name" {
            version 2;
            immediate-leave;
            promiscuous-mode;
            distributed;
            ssm-map-policy ssm-map-v6;
            group-policy mld-group-policy;
        }
    }
}
```

Configure Lawful Intercept

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

1. Configure lawful intercept on the BNG CUPS Controller, as follows:

```

[edit]
services radius-flow-tap {
       policy <name> {
             …
       }
       snmp notify-targets <IP address list>;
}
```

2. Configure services on the BNG User Planes that are associated to the BNG CUPS Controller, as follows:

   ```
   [edit]
   services radius-flow-tap {
          forwarding-class <fc>;
          {
                 …
          }
          routing-instance <ri>;
          source-ipv4-address <address>;
   }
   ```

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>;
```

```
            }
        }
    }
}
```

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 69 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: • `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 effect the running application.
`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`* `--version` *`software-release 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. • `version` *`software-release`*—The software release for the new CPi pod. Enter a release. • *`label`*—Specify a label that is used for CPi commands.
`sudo -E dbng cpi rm --context` *`context-name 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. • *`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. • `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
`sudo -E dbng link --version` *`software-release`* `--context` *`context-name`*	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.
`sudo -E dbng rename-context --context` *`context-name`* `--new-name` *`new-name`*	Renames a context. Does not effect the currently running BNG CUPS Controller 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.
`sudo -E dbng restart --context` *`context-name`* `[--force] [--wait]` *`service-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.

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. • `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.
`sudo -E dbng setup --context` *context-name* `[--bbecloudsetup] [--update] [--ssh` *ip-address:port-number*`]`	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. • `bbecloudsetup`—Use the values for a BBE Cloudsetup deployed cluster. This option is only valid if you are running a BBE Cloudsetup deployed cluster. • `update`—You will only be prompted for missing values during setup. • `ssh` *ip-address:port-number*—Enables SSH. Enter the SSH IP address and port number. The IP address must be the IP address that you SSH to. The IP address can also be a DNS name.
`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. • `pipe`—Allows you to pipe input into the command.

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

Command Name	Action
`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.
`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. • `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]`	Stop 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. • `now`—Stops the BNG CUPS Controller immediately, instead of waiting for the two minute delay.
`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. • `output`—Allows you to request the output in JSON format. • `terse`—Displays a summarized output of the storage health.

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

Command Name	Action
`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.
`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. • `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 inTable 2 on page 76. 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.

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 CLI Configuration Statements.

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 Operational Commands.

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).

Figure 15 on page 81 shows an example of the BBE Event Collection and Visualization Dashboard.

Figure 15: BBE Event Collection and Visualization Dashboard

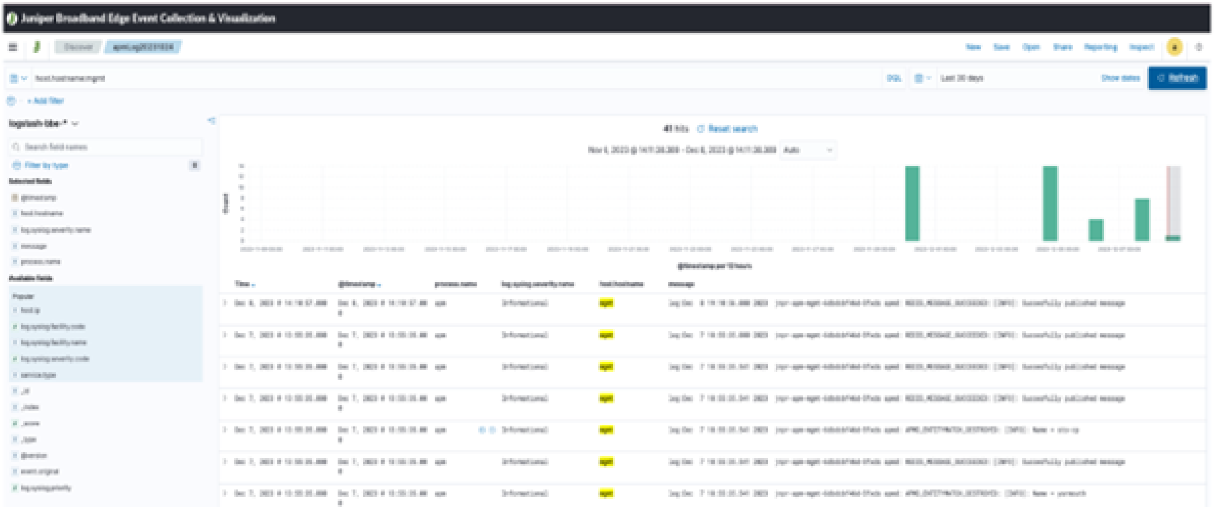

BBE Event Collection and Visualization can be installed in the same Kubernetes cluster as the Juniper BBE applications (BNG CUPS Controller and APM). The installation follows the same model (installed from the jump host). For BBE Event Collection and Visualization installation instructions, see Broadband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

3

CHAPTER

Juniper BNG CUPS CLI Configuration Statements

Juniper BNG CUPS 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.

address-pool-manager

IN THIS SECTION

Syntax

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

Hierarchy Level

```
[edit 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.

port *port-number* The port that APM is listening on for incoming address pool manager connections.

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-director

IN THIS SECTION

Syntax

```
bng-director {
    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-profiles 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-groups  {
         }
         load-balancing-groups {
         }
       }
    }
  }
}
```

Hierarchy Level

```
[edit groups]
```

Description

Configures the BNG Director on the BNG CUPS Controller. The BNG Director manages all the control plane instances (CPi).

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-profiles`

Secify a security profile. See "security-profiles" on page 126.

`user-plane`

Specify the BNG User Planes to be associated with the BNG CUPS Controller. See "user-planes (bng-controller)" on page 138.

`control-plane-instances`

See "control-plane-instances" on page 90.

`subscriber-groups`	See subscriber-groups.
`load-balancing-groups`	See "load-balancing-groups" on page 99.

captive-portal-content-delivery-profile (Services)

IN THIS SECTION

Syntax

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

Hierarchy Level

```
[edit system services subscriber-management mode control-plane user-plane bng-user-plane-name
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 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 <value>;
           backup-tag <value>;
     }
     source-partition-qualifier string;
     excluded-address last-octet number;
     dhcp-gateway-address-last-octet number;
     protocol-attributes dhcp-attribute;
    }
 }
```

Hierarchy Level

```
[edit 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 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 {
accounting;
    interface interface-name {
        (accounting | no-accounting);
        disable;
        distributed;
        group-limit limit;
        group-policy [ policy-names ];
        group-threshold
        immediate-leave;
        log-interval
        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 dynamic-profiles profile-name]
```

Description

Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets. This command runs on the BNG CUPS Controller.

The remaining statements are explained separately. See CLI Explorer.

Default

IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

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.

interface (Protocols IGMP)

IN THIS SECTION

Syntax

```
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 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.

interface (Protocols MLD)

IN THIS SECTION

Syntax

```
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 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.

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 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 {
    accounting;
    interface interface-name {
        (accounting | no-accounting);
        disable;
        distributed;
        group-limit  limit;
        group-policy [ policy-names ];
        immediate-leave;
        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 dynamic-profiles profile-name]
```

Description

Enable MLD on the router. MLD must be enabled for the router to receive multicast packets. This command runs on the BNG CUPS Controller.

Default

MLD is disabled on the router. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options

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.

multicast

IN THIS SECTION

Syntax

```
multicast {
    asm-override-ssm;
    backup-pe-group group-name {
        backups [ addresses ];
        local-address address;
    }
    cont-stats-collection-interval interval;
    flow-map flow-map-name {
        bandwidth (bps | adaptive);
        forwarding-cache {
            timeout (never non-discard-entry-only | minutes);
        }
        policy [ policy-names ];
        redundant-sources [ addresses ];
    }
    forwarding-cache {
        threshold suppress value <reuse value>;
        timeout minutes;
    }
    interface interface-name {
        enable;
        maximum-bandwidth bps;
        no-qos-adjust;
        reverse-oif-mapping {
            no-qos-adjust;
        }
        subscriber-leave-timer seconds;
    }
    local-address address
    omit-wildcard-address
    pim-to-igmp-proxy {
        upstream-interface [ interface-names ];
    }
    pim-to-mld-proxy {
        upstream-interface [ interface-names ];
    }
    rpf-check-policy [ policy-names ];
    scope scope-name {
        interface [ interface-names ];
        prefix destination-prefix;
    }
    scope-policy [ policy-names ];
    ssm-groups [ addresses ];
    ssm-map ssm-map-name {
        policy [ policy-names ];
        source [ addresses ];
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <disable>;
    }
}
```

Hierarchy Level

```
[edit routing-options]
```

Description

Configure multicast routing options properties. Note that you cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the scope statement does apply individually to a specific routing instance.

NOTE: The `multicast` command runs on the BNG CUPS Controller.

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.

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.

overrides

IN THIS SECTION

Syntax

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

Hierarchy Level

```
[edit groups group-name 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 bng-director bng-controller user-planes bng-user-plane-name dynamic-address-poolsl-
plane]
```

Description

Defines the BNG User Plane partition attribute. The partition attribute defines the geographical region or area to which the BNG User Plane belongs.

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 bng-director bng-controller]
```

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

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 (Protocols IGMP)

IN THIS SECTION

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit protocols igmp]
```

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-interval (Protocols MLD)

IN THIS SECTION

Syntax

```
query-interval seconds;
```

Hierarchy Level

```
[edit protocols mld]
```

Description

Specify how often the querier router 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-engine-services

IN THIS SECTION

Syntax

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

Hierarchy Level

```
[edit system services subscriber-management mode control-plane user-plane bng-user-plane-name
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. The services interfaces on the Routing Engine are identified with an si- prefix (for example, si-1/1/0). The si- interface contains all redirect and rewrite traffic and services for the Routing Engine. 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.

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.

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 bng-director bng-controller user-planes transport]
```

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-interface (Services Interfaces)

IN THIS SECTION

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit system services subscriber-management mode control-plane user-plane bng-user-plane-name
service-set service-set-name interface-service]
```

Description

Specify the name for the services interface associated with an interface-wide service set. This command runs on the BNG CUPS Controller.

Options

`interface-name` Identifier of the service interface.

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.

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.

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;
      }
    }
  }
```

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.

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;
   }
 }
```

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.

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 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 bng-dirctor bng-controller user-planes]
```

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 (see "security-profiles" on page 126).

user-plane-profile

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-common-0
```

Description

A user-plane-profile is a template for configuring a BNG User Plane in terms of interfaces, pfcp behavior, and subscriber management override behavior. A user-plane-profile 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. So, 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-common-0).

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 108).

interfaces *interface-name*
Specify interfaces configuration to be used for the BNG User Plane (see "interfaces (Static and Dynamic Subscribers)" on page 499.

user-planes (bng-controller)

IN THIS SECTION

Syntax

```
user-planes {
    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 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 135.

`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-profile" on page 136.

weight

IN THIS SECTION

Syntax

```
weight weight-number;
```

Hierarchy Level

```
[edit dynamic-profiles dynamic-profiles-name interfaces $junos-interface-ifd-name unit $junos-
interface-unit load-balance]
[edit dynamic-profiles dynamic-profiles- name interfaces interface-set $junos-phy-ifd-interface-
set-name load-balance]
```

Description

Set 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 CLI Operational Statements

Juniper BNG CUPS 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.

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 switchover

IN THIS SECTION

Syntax

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

Description

Used to 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.

request user-plane manager restart up-name

IN THIS SECTION

Syntax

```
request user-plane manager restart up-name bng-user-plane-name
```

Description

Restart the managing SMD service instance on the BNG CUPS Controller associated to the specified BNG User Plane.

Options

up-plane ***bng-user-plane-name***	The BNG User Plane for which you want to restart the managing SMD service instance on its associated BNG CUPS Controller.

Required Privilege Level

view

Output Fields

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

Sample Output

request user-plane manager restart up-name

```
user@host> request user-plane manager restart up-name exampl-up-1
```

restart bbe-cpm-daemon

IN THIS SECTION

Syntax

```
restart bbe-cpm-daemon
```

Description

Restarts the Control 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. This command runs on the BNG CUPS Controller.

Sample Output

restart bbe-cpm-daemon

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

restart bbe-stats-daemon

IN THIS SECTION

Syntax

```
restart bbe-stats-daemon
```

Description

Restarts the Enhanced Session Management Statistics 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. This command runs on the BNG CUPS Controller.

Sample Output

restart bbe-stats-daemon

```
user@host> restart bbe-stats-daemon
Control Plane Manager for dBNG started, pid <process-id>
```

restart bbe-stats-svcsd

IN THIS SECTION

Syntax

```
restart bbe-stats-svcsd
```

Description

Restarts the Statistics Services 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. This command runs on the BNG User Planes.

Sample Output

restart bbe-stats-svcsd

```
user@host> restart bbe-stats-svcsd
Control Plane Manager for dBNG started, pid <process-id>
```

restart bbe-upm-daemon

IN THIS SECTION

Syntax

```
restart bbe-upm-daemon
```

Description

Restarts the User Plane Manager daemon. This command runs on the BNG CUPS Controller.

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 bbe-upsf-daemon

IN THIS SECTION

Syntax

```
restart bbe-upsf-daemon
```

Description

Restarts the User Plane Selection Function daemon. This command runs on the BNG CUPS Controller.

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-upsf-daemon

```
user@host> restart bbe-upsf-daemon
Control Plane Manager for dBNG started, pid <process-id>
```

restart cp-smg-server

IN THIS SECTION

Syntax

```
restart cp-smg-server
```

Description

Restarts the Enhanced Session Management BNG CUPS Controller process. This command runs on the BNG CUPS Controller.

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 cp-smg-server

```
user@host> restart cp-smg-server
Control Plane Manager for dBNG started, pid <process-id>
```

restart replication-client-process

IN THIS SECTION

Syntax

```
restart replication-client-process
```

Description

Restarts the Replication Client Process. A BNG User Plane hosts the Replication Client Process daemon and the Replication Server Process daemon. These daemons replicate the state between the BNG CUPS Controller and the BNG User Plane and the routing engines. This command runs on the BNG CUPS Controller.

Avoid using this command unless Juniper Networks Technical Assistance Center (JTAC) directs you to use it.

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 replication-client-process

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

restart replication-server-process

IN THIS SECTION

Syntax

```
restart replication-server-process
```

Description

Restarts the Replication Server Process. A BNG User Plane hosts the Replication Client Process daemon and the Replication Server Process daemon. These daemons replicate the state between the BNG CUPS Controller and the BNG User Plane and the routing engines. This command runs on the BNG CUPS Controller.

Avoid using this command unless Juniper Networks Technical Assistance Center (JTAC) directs you to use it.

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 replication-server-process

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

restart up-helper-service

IN THIS SECTION

Syntax

```
restart up-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 up-helper-service

```
user@host> restart up-helper-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 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 171 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.

Table 3: show broadband-device Output Fields *(Continued)*

Field Name	Field Description
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 174 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 micro-service-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` *`micro-service-name`*	Displays health information for the specified micro service.
`service` *`service-name`*	Displays health information for the associated endpoints for the specified service that is part of the micro service. 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 176 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
`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.
`Subsystem`	The subsystem for which the information is being displayed.
`Shared-memory`	The health of the shared memory.
`Initial State Recovery`	Displays whether the initial state recovery succeeded.

Table 5: show health *(Continued)*

Field Name	Field Description
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.
HealthBNG User Plane	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.

Table 5: show health *(Continued)*

Field Name	Field Description
Address	BNG User Plane IP address.
Active/Backup-sess	The number of active and backup subscriber sessions served by the BNG User Plane.
State (user-plane)	The state of the BNG User Plane. The state can be one of the following: • initializing • ready • connecting • connected • disconnecting • disconnected • security-updating • warm-init • deconfiguring • misconfigured
User-plane	The BNG User Plane for which the information is being reported.
Id	ID of the BNG User Plane.
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                 Health          Unhealthy-services    Uptime
host                 healthy         0                     06:08:28
```

show health subsystem

```
user@host> show health subsystem cpihardening
Subsystem: cpihardening
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
  ScachePublish                 reconciled        06:08:20        0
  Apm                           connected         06:08:20        0
```

show health subsystem <micro-service-name> service

```
user@host> show health subsystem cpi-boston service ppp-service
Subsystem: cpi-boston
Service: ppp-service
  Status: up
  State: ready
  Up-time: 2d 12:43:16
  Restarts: 0
  Memory Usage  : 1059540KB
  CPU% (threads): 0.1% (1)
  Endpoint-health: healthy
  Endpoints:Id          Flapped   State            Up-time
  PppSmdIpc:4           0         connected        0d 10:43:22
  PppSmdIpc:3           0         connected        0d 10:43:22
  PppSmdIpc:5           0         connected        0d 10:43:22
  PppSmdIpc:1           0         connected        0d 10:43:22
  PppSmdIpc:2           0         connected        0d 10:43:22
```

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 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 6 on page 184 describes the output fields for the `show igmp group` command. Output fields are listed in the approximate order in which they appear.

Table 6: 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.
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
```

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

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 7 on page 188 describes the output fields for the `show igmp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 7: 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.
`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.

Table 7: show igmp statistics Output Fields *(Continued)*

Field Name	Field Description
`Received`	Number of messages received.
`Sent`	Number of messages sent.
`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
```

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

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 8 on page 192 lists the output fields for the `show load-balancing-group` command.

Table 8: 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.

Sample Output

show load-balancing group

```
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 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 9 on page 194 describes the output fields for the `show mld group` command. Output fields are listed in the approximate order in which they appear.

Table 9: 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.
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
```

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

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 10 on page 198 describes the output fields for the `show mld statistics` command. Output fields are listed in the approximate order in which they appear.

Table 10: 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 10: 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
```

Release Information

Statement introduced in Juniper BNG CUPS Release 22.4R1.

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 11 on page 201 lists the output fields for the `show network-access address-assignment address-pool-manager status` command. Output fields are listed in alphabetical order.

Table 11: 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

Table 11: show network-access address-assignment address-pool-manager status Output Fields *(Continued)*

Field Name	Field Description
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 12 on page 204 lists the output fields for the `show network-access address-assignment domain` command. Output fields are listed in alphabetical order.

Table 12: 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.
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.

Table 12: show network-access address-assignment domain Output Fields *(Continued)*

Field Name	Field Description
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 13 on page 207 lists the output fields for the `show network-access address-assignment domain-state` command. Output fields are listed in alphabetical order.

Table 13: show network-access address-assignment domain-state Output Fields

Field Name	Field Description
Domain Name	Name of the pool domain.
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 35 on page 296 lists the output fields for the `show routing-instances` command. Output fields are listed in the approximate order in which they appear.

Table 14: show routing-instances Output Fields

Field Name	Field Description
`User Plane Name`	Name of the BNG User Plane.

Table 14: show routing-instances Output Fields *(Continued)*

Field Name	Field Description
`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 15 on page 211 lists the output fields for the `show subscriber-group` command.

Table 15: 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.

Table 15: show subscriber-group Output Fields *(Continued)*

Field Name	Field Description
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.
User-Plane	The BNG User Plane that the specified subscriber group belongs to.
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.

Sample Output

show subscriber-group

```
user@host> show subscriber-group
Name         ID SGRP Mode      SGRP State    User Plane    User Plane   Active UP
swwf-mx204-d 2  Control Plane  healthy       swwf-mx204-d  ---          swwf-mx204-d
Italy        3  Control Plane  healthy       swwf-mx204-d  swwf-mx204-e swwf-mx204-d
Greece       4  Control Plane  active-active swwf-mx204-d  swwf-mx204-e ---
Spain        5  Control Plane  backup-backup swwf-mx204-d  swwf-mx204-e ---
```

show subscriber-group

```
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
<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

`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 16 on page 216 lists the output fields for the `show subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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 16: 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
<device header for cpi-boston>
SID       CP-instance    SGRP             Active-UP
234096    cpi-example1     test-vest003      alkaid
230077    cpi-example1     test-west001      alkaid
.
.
<device header for cpi-test1>
SID       CP-instance    SGRP             Active-UP
28603     cpi-test1      north-frame001   northboro
```

show subscribers user-name

```
user@host> show subscribers user-name user@host.com display detail
<device header for cpi-example1>
Type: DHCP
User Name: user@host.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
Active User Plane: alkaid
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 17 on page 236 lists the output fields for the `show subscribers subscriber-group` command. Output fields are listed in the approximate order in which they appear.

Table 17: 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.

Table 17: show subscribers subscriber-group Output Fields *(Continued)*

Field Name	Field Description
RI	Routing instance associated with the subscriber.
SID	ID number for a subscriber session.
Resiliency-state	Lists weather 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>
```

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.

Required Privilege Level

view

Output Fields

Table 18 on page 239 lists the output fields for the `show subscribers summary` command. Output fields are listed in the approximate order in which they appear.

Table 18: show subscribers summary Output Fields

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 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 19 on page 242 lists the output fields for the `show subscribers user-plane` command. Output fields are listed in the approximate order in which they appear.

Table 19: 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.

Table 19: show subscribers user-plane Output Fields *(Continued)*

Field Name	Field Description
`Resiliency-state`	Lists weather 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

IN THIS SECTION

Syntax

```
show user-plane
```

Description

Display a summary of information for all configured BNG User Planes, which includes their associations, health state, and CPi binding.

Options

Required Privilege Level

view

Output Fields

Table 20 on page 244 lists the output fields for the `show user-plane` command. Output fields are listed in the approximate order in which they appear.

Table 20: show user-plane

Field Name	Field Description
`Name`	The BNG User Plane name.
`Uptime`	The amount of time the service has been up.
`Health`	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.
`Address`	BNG User Plane IP address.

Table 20: show user-plane *(Continued)*

Field Name	Field Description
`Active/Backup-sess`	The number of active and backup subscriber sessions served by the BNG User Plane.
`State`	The state of the BNG User Plane. The state can be one of the following: • `initializing` • `ready` • `connecting` • `connected` • `disconnecting` • `disconnected` • `security-updating` • `warm-init` • `deconfiguring` • `misconfigured`
`CPi`	The control plane instance that is handling the BNG User Plane.

Sample Output

show user-plane

```
user@host> show 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 user-plane class-of-service scheduler-map

IN THIS SECTION

Syntax

```
show user-plane class-of-service scheduler-map
<static-only>
<scheduler-name>
<up-name user-plane-name>
```

Description

Displays BNG User plane specific information for the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options

static-only Displays only statically configured scheduler maps.

scheduler-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.

user-plane-name Name of the user-plane for which you want to view information.

Required Privilege Level

view

Output Fields

Table 21 on page 247 describes the output fields for the `show user-plane class-of-service scheduler-map` command. Output fields are listed in the approximate order in which they appear.

Table 21: show user-plane class-of-service scheduler-map Output Fields

Field Name	Field Description
`Scheduler map`	Name of the 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.
`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.
`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.

Table 21: show user-plane class-of-service scheduler-map Output Fields *(Continued)*

Field Name	Field Description
`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.
`Priority`	Scheduling priority: `low` or `high`.
`Excess priority`	Priority of excess bandwidth: `low`, `medium-low`, `medium-high`, `high`, or `none`.
`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

```
user@host> show user-plane class-of-service scheduler-map static-only smap-mixed up-name up-
test
Scheduler map: smap-mixed, Index: 15931

  Scheduler: sched-be, Forwarding class: best-effort, Index: 44487
    Transmit rate: unspecified, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
    Priority: low
    Excess Priority: unspecified, Excess rate: 12 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP         39865    dp-static
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: sched-ef, Forwarding class: expedited-forwarding, Index: 44324
    Transmit rate: unspecified, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
    Priority: low
    Excess Priority: unspecified, Excess rate: 10 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP             1    <default-drop-profile>
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: sched-af, Forwarding class: assured-forwarding, Index: 44452
    Transmit rate: unspecified, Rate Limit: none, Buffer size: remainder, Buffer Limit: none,
    Priority: low
    Excess Priority: unspecified, Excess rate: 10 percent,
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP         39865    dp-static
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```

show user-plane class-of-service traffic-control-profile

IN THIS SECTION

Syntax

```
show user-plane class-of-service traffic-control-profile
<static-only>
<profile-name>
<up-name user-plane-name>
```

Description

Display information for traffic shaping and scheduling profiles, for the specified BNG User Plane.

Options

`static-only`	Displays only statically configured traffic control profiles.
`profile-name`	Name of the traffic control profile for which you want to view information.
`user-plane-name`	The name of the BNG User Plane for which you want to view information.

Required Privilege Level

view

Output Fields

Table 22 on page 251 describes the output fields for the `show class-of-service traffic-control-profile` command. Output fields are listed in the approximate order in which they appear.

Table 22: show user-plane class-of-service traffic-control-profile Output Fields

Field Name	Field Description
`Traffic control profile`	Name of the traffic control profile. You can configure objects of the same type with the same name on the user plane and the control plane. The display designates between the control plane and the user plane by adding cp or up to the name. Also, the dynamically generated UID is displayed.
`Index`	Index number of the traffic control profile.
`Shaping rate`	Configured shaping rate, in bps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
`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.
`User-plane`	The associated BNG User Plane.

Sample Output

show user-plane class-of-service traffic-control-profile static-only

```
user@host> show user-plane class-of-service traffic-control-profile static-only tcp-mixed up-
name test-123
Traffic control profile: tcp-mixed, Index: 18213
  Scheduler map: smap-mixed
```

show user-plane <user-plane-name> class-of-service traffic-control-profile

```
user@host> show user-plane test-123 class-of-service traffic-control-profile
Traffic control profile: tcp-delete-me, Index: 17350
  Shaping rate: 10000000
  Scheduler map: <default>

Traffic control profile: tcp-mixed, Index: 18213
  Scheduler map: smap-mixed

Traffic control profile: tcp-andover_UID1074, Index: 4299161706
  User-plane: test-123
  Scheduler map: smap-mixed

Traffic control profile: tcp-iflset_UID1075, Index: 4299161705
  User-plane: test-123
  Shaping rate: 155000000
  Scheduler map: <default>
```

show user-plane class-of-service traffic-control-profile static-only

```
user@host> show user-plane class-of-service traffic-control-profile static-only up-name test-123
Traffic control profile: tcp-delete-me, Index: 17350
  Shaping rate: 10000000
  Scheduler map: <default>

Traffic control profile: tcp-mixed, Index: 18213
  Scheduler map: smap-mixed
```

show user-plane firewall filter

IN THIS SECTION

Syntax

```
show firewall
<filter filter-name>
<up-name user-plane-name>
```

Description

Display firewall instance information for the specified BNG User Plane.

Options

filter *filter-name*	Name of a configured filter.
up-name *user-plane-name*	The BNG User Plane for which you want to view firewall filter information.

Required Privilege Level

view

Output Fields

Table 4 on page 174 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 23: show user-plane firewall filter 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 user-plane firewall filter

```
user@host> show user-plane firewall filter dynamic-filter_UID1080-demux0.3221225473-out up-name
up-example-1
Filter: dynamic-filter_UID1080-demux0.3221225473-out          User Plane: up-example-1
```

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

root

Output Fields

Table 24 on page 255 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 24: 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 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 25 on page 257 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 25: 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.

Table 25: show user-plane igmp interface Output Fields *(Continued)*

Field Name	Field Description
`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..
`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.

Table 25: show user-plane igmp interface Output Fields *(Continued)*

Field Name	Field Description
`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.
`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 26 on page 261 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 26: 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.

Table 26: show user-plane ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
`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.
`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.

Table 26: show user-plane ipv6 router-advertisement Output Fields *(Continued)*

Field Name	Field Description
`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 27 on page 265 lists the output fields for the `show user-plane maintenance up-name` command.

Table 27: 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.
`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.

Table 27: show user-plane maintenance up-name Output Fields *(Continued)*

Field Name	Field Description
`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 28 on page 268 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 28: 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.
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

Table 28: show user-plane mld interface Output Fields *(Continued)*

Field Name	Field Description
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.
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 29 on page 271 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 29: 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

Table 29: show user-plane pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
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
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 29: 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 30 on page 275 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 30: 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 30: 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 idenfier (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 31 on page 279 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 31: 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 31: 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 32 on page 283 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 32: 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 33 on page 285 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 33: 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 34 on page 288 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 34: 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 34: 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 34: 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 34: 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 34: 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 35 on page 296 lists the output fields for the `show user-plane route` command. Output fields are listed in the approximate order in which they appear.

Table 35: 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 35: 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 35: 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 35: 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 35 on page 296 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 36: show user-plane routing-instances Output Fields

Field Name	Field Description
`User Plane`	Name of the BNG User Plane.

Table 36: 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 37 on page 308 lists the output fields for the `show user-plan statistics` command. Output fields are listed in the approximate order in which they appear.

Table 37: 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 37: 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
```

request network-access aaa address-assignment domain-profile

IN THIS SECTION

Syntax

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

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-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 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 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 bbe-cpm-daemon

IN THIS SECTION

Syntax

```
restart bbe-cpm-daemon
```

Description

Restarts the Control 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. This command runs on the BNG CUPS Controller.

Sample Output

restart bbe-cpm-daemon

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

restart bbe-stats-daemon

IN THIS SECTION

Syntax

```
restart bbe-stats-daemon
```

Description

Restarts the Enhanced Session Management Statistics 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. This command runs on the BNG CUPS Controller.

Sample Output

restart bbe-stats-daemon

```
user@host> restart bbe-stats-daemon
Control Plane Manager for dBNG started, pid <process-id>
```

restart bbe-stats-svcsd

IN THIS SECTION

Syntax

```
restart bbe-stats-svcsd
```

Description

Restarts the Statistics Services 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. This command runs on the BNG User Planes.

Sample Output

restart bbe-stats-svcsd

```
user@host> restart bbe-stats-svcsd
Control Plane Manager for dBNG started, pid <process-id>
```

restart bbe-upm-daemon

IN THIS SECTION

Syntax

```
restart bbe-upm-daemon
```

Description

Restarts the User Plane Manager daemon. This command runs on the BNG CUPS Controller.

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 bbe-upsf-daemon

IN THIS SECTION

Syntax

```
restart bbe-upsf-daemon
```

Description

Restarts the User Plane Selection Function daemon. This command runs on the BNG CUPS Controller.

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-upsf-daemon

```
user@host> restart bbe-upsf-daemon
Control Plane Manager for dBNG started, pid <process-id>
```

restart cp-smg-server

IN THIS SECTION

Syntax

```
restart cp-smg-server
```

Description

Restarts the Enhanced Session Management BNG CUPS Controller process. This command runs on the BNG CUPS Controller.

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 cp-smg-server

```
user@host> restart cp-smg-server
Control Plane Manager for dBNG started, pid <process-id>
```

restart replication-client-process

IN THIS SECTION

Syntax

```
restart replication-client-process
```

Description

Restarts the Replication Client Process. A BNG User Plane hosts the Replication Client Process daemon and the Replication Server Process daemon. These daemons replicate the state between the BNG CUPS Controller and the BNG User Plane and the routing engines. This command runs on the BNG CUPS Controller.

Avoid using this command unless Juniper Networks Technical Assistance Center (JTAC) directs you to use it.

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 replication-client-process

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

restart replication-server-process

IN THIS SECTION

Syntax

```
restart replication-server-process
```

Description

Restarts the Replication Server Process. A BNG User Plane hosts the Replication Client Process daemon and the Replication Server Process daemon. These daemons replicate the state between the BNG CUPS Controller and the BNG User Plane and the routing engines. This command runs on the BNG CUPS Controller.

Avoid using this command unless Juniper Networks Technical Assistance Center (JTAC) directs you to use it.

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 replication-server-process

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

restart up-helper-service

IN THIS SECTION

Syntax

```
restart up-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 up-helper-service

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

restart up-smg-server

IN THIS SECTION

Syntax

```
restart up-smg-server
```

Description

Restarts the Enhanced Session Management BNG User Plane process. This command runs on the BNG CUPS Controller.

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-server

```
user@host> restart up-smg-server
Control Plane Manager for dBNG started, pid <process-id>
```

5

CHAPTER

Junos OS CLI Configuration Statements

Firewall Filter Match Conditions for IPv4 Traffic

You can configure a firewall filter with match conditions for Internet Protocol version 4 (IPv4) traffic (`family inet`).

NOTE: For MX Series routers with MPCs, you need to initialize the filter counter for Trio-only match filters in the MIB by walking the corresponding SNMP MIB, for example, `show snmp mib walk` *`name`* `ascii`. This forces Junos to learn the filter counters, and ensures that the filter statistics are displayed (this is because the first poll to filter statistics may not show all counters). This guidance applies to all enhanced mode firewall filters, filters with flexible conditions, and filters with certain terminating actions. See those topics, listed under Related Documentation, for details.

Table 38 on page 334 describes the *`match-conditions`* you can configure at the `[edit firewall family inet filter` *`filter-name`* `term` *`term-name`* `from]` hierarchy level.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic

Match Condition	Description
`address` *`address`* `[ except ]`	Match the IPv4 source or destination address field unless the `except` option is included. If the option is included, do not match the IPv4 source or destination address field.
`destination-address` *`address`* `[ except ]`	Match the IPv4 destination address field unless the `except` option is included. If the option is included, do not match the IPv4 destination address field. You cannot specify both the `address` and `destination-address` match conditions in the same term.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
destination-port *number*	Match the UDP or TCP destination port field. You cannot specify both the port and destination-port match conditions in the same term. If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).
destination-port-except *number*	Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.
destination-prefix-list *name* [except]	Match destination prefixes in the specified list unless the except option is included. If the option is included, do not match the destination prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list *prefix-list-name*] hierarchy level.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`dscp` *number*	Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. Support was added for filtering on Differentiated Services Code Point (DSCP) and forwarding class for Routing Engine sourced packets, including IS-IS packets encapsulated in generic routing encapsulation (GRE). Subsequently, when upgrading from a previous version of Junos OS where you have both a class of service (CoS) and firewall filter, and both include DSCP or forwarding class filter actions, the criteria in the firewall filter automatically takes precedence over the CoS settings. The same is true when creating new configurations; that is, where the same settings exist, the firewall filter takes precedence over the CoS, regardless of which was created first. You can specify a numeric value from `0` through `63`. To specify the value in hexadecimal form, include `0x` as a prefix. To specify the value in binary form, include `b` as a prefix. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*, defines one code point: `ef` (46). • RFC 2597, *Assured Forwarding PHB Group*, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: • `af11` (10), `af12` (12), `af13` (14) • `af21` (18), `af22` (20), `af23` (22) • `af31` (26), `af32` (28), `af33` (30) • `af41` (34), `af42` (36), `af43` (38)
`dscp-except` *number*	Do not match on the DSCP number. For more information, see the `dscp` match condition.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

<table>
<tr><th>Match Condition</th><th colspan="2">Description</th></tr>
<tr><td>esp-spi <i>spi-value</i></td><td colspan="2">Match the IPsec encapsulating security payload (ESP) SPI value. Match on this specific SPI value. You can specify the ESP SPI value in hexadecimal, binary, or decimal form.
<b>NOTE</b>: This match condition is not supported on PTX series routers.</td></tr>
<tr><td>esp-spi-except <i>spi-value</i></td><td colspan="2">Match the IPsec ESP SPI value. Do not match on this specific SPI value.
<b>NOTE</b>: This match condition is not supported on PTX series routers.</td></tr>
<tr><td>first-fragment</td><td colspan="2">Match if the packet is the first fragment of a fragmented packet. Do not match if the packet is a trailing fragment of a fragmented packet. The first fragment of a fragmented packet has a fragment offset value of 0.
This match condition is an alias for the bit-field match condition fragment-offset 0 match condition.
To match both first and trailing fragments, you can use two terms that specify different match conditions: first-fragment and is-fragment.</td></tr>
<tr><td rowspan="6">flexible-match-mask <i>value</i></td><td>bit-length</td><td>Length of the data to be matched in bits, not needed for string input (0..128)</td></tr>
<tr><td>bit-offset</td><td>Bit offset after the (match-start + byte) offset (0..7)</td></tr>
<tr><td>byte-offset</td><td>Byte offset after the match start point</td></tr>
<tr><td>flexible-mask-name</td><td>Select a flexible match from predefined template field</td></tr>
<tr><td>mask-in-hex</td><td>Mask out bits in the packet data to be matched</td></tr>
<tr><td>match-start</td><td>Start point to match in packet</td></tr>
</table>

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description	
	`prefix`	Value data/string to be matched
`flexible-match-range` *value*	`bit-length`	Length of the data to be matched in bits (0..32)
	`bit-offset`	Bit offset after the (match-start + byte) offset (0..7)
	`byte-offset`	Byte offset after the match start point
	`flexible-range-name`	Select a flexible match from predefined template field
	`match-start`	Start point to match in packet
	`range`	Range of values to be matched
	`range-except`	Do not match this range of values
`forwarding-class` *class*	Match the forwarding class of the packet. Specify `assured-forwarding`, `best-effort`, `expedited-forwarding`, or `network-control`.	
`forwarding-class-except` *class*	Do not match the forwarding class of the packet. For details, see the `forwarding-class` match condition.	
`fragment-flags` *number*	(Ingress only) Match the three-bit IP fragmentation flags field in the IP header. In place of the numeric field value, you can specify one of the following keywords (the field values are also listed): `dont-fragment` (0x4), `more-fragments` (0x2), or `reserved` (0x8).	

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`fragment-offset` *value*	Match the 13-bit fragment offset field in the IP header. The value is the offset, in 8-byte units, in the overall datagram message to the data fragment. Specify a numeric value, a range of values, or a set of values. An offset value of `0` indicates the first fragment of a fragmented packet. The `first-fragment` match condition is an alias for the `fragment-offset 0` match condition. To match both first and trailing fragments, you can use two terms that specify different match conditions (`first-fragment` and `is-fragment`).
`fragment-offset-except` *number*	Do not match the 13-bit fragment offset field.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description

`icmp-code` *`number`*	Match the ICMP message code field.
	NOTE: When using this match condition, you should also use the `protocol icmp` match condition in the same term (as shown below) to ensure that **icmp** packets are being evaluated.
	`term Allow _ICMP {`
	`        from protocol icmp {`
	`            icmp-code ip-header-bad;`
	`            icmp-type echo-reply;`
	`        }`
	`        then {`
	`            policer ICMP_Policier;`
	`            count Allow_ICMP;`
	You must also configure the `icmp-type` *`message-type`* match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.
	In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:
	• parameter-problem: `ip-header-bad` (0), `required-option-missing` (1)
	• redirect: `redirect-for-host` (1), `redirect-for-network` (0), `redirect-for-tos-and-host` (3), `redirect-for-tos-and-net` (2)
	• time-exceeded: `ttl-eq-zero-during-reassembly` (1), `ttl-eq-zero-during-transit` (0)
	• unreachable: `communication-prohibited-by-filtering` (13), `destination-host-prohibited` (10), `destination-host-unknown` (7), `destination-network-prohibited` (9), `destination-network-unknown` (6), `fragmentation-needed` (4), `host-precedence-violation` (14), `host-unreachable` (1), `host-unreachable-for-TOS` (12), `network-unreachable` (0), `network-unreachable-for-TOS` (11), `port-unreachable` (3), `precedence-cutoff-in-effect` (15), `protocol-unreachable` (2), `source-host-isolated` (8), `source-route-failed` (5)

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

<table>
<tr><th>Match Condition</th><th>Description</th></tr>
<tr><td>icmp-code-except <em>message-code</em></td><td>Do not match the ICMP message code field. For details, see the icmp-code match condition.</td></tr>
<tr><td>icmp-type <em>number</em></td><td>

Match the ICMP message type field.

NOTE: When using this match condition, you should also use the protocol icmp match condition in the same term (as shown below) to ensure that **icmp** packets are being evaluated.

```
term Allow _ICMP {
                from protocol icmp {
                    icmp-code ip-header-bad;
                    icmp-type echo-reply;
                }
                then {
                    policer ICMP_Policier;
                    count Allow_ICMP;
```

You must also configure the icmp-type *message-type* match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.

NOTE: For Junos OS Evolved, you must configure the protocol match statement in the same term.

In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).

</td></tr>
<tr><td>icmp-type-except <em>message-type</em></td><td>Do not match the ICMP message type field. For details, see the icmp-type match condition.</td></tr>
</table>

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`ip-options` *`values`*	Match the 8-bit IP option field, if present, to the specified value or list of values. In place of a numeric value, you can specify one of the following text synonyms (the option values are also listed): `loose-source-route` (131), `record-route` (7), `router-alert` (148), `security` (130), `stream-id` (136),`strict-source-route` (137), or `timestamp` (68). To match *any* value for the IP option, use the text synonym `any`. To match on *multiple* values, specify the list of values within square brackets ('`[`' and '`]`'). To match a *range* of values, use the value specification `[` *`value1-value2`* `]`. For example, the match condition `ip-options [ 0-147 ]` matches on an IP options field that contains the `loose-source-route`, `record-route`, or `security` values, or any other value from 0 through 147. However, this match condition does not match on an IP options field that contains only the `router-alert` value (148). For most interfaces, a filter term that specifies an `ip-option` match on one or more *specific* IP option values (a value other than `any`) causes packets to be sent to the Routing Engine so that the kernel can parse the IP option field in the packet header. • For a firewall filter term that specifies an `ip-option` match on one or more specific IP option values, you cannot specify the `count`, `log`, or `syslog` nonterminating actions *unless* you also specify the `discard` terminating action in the same term. This behavior prevents double-counting of packets for a filter applied to a transit interface on the router. • Packets processed on the kernel might be dropped in case of a system bottleneck. To ensure that matched packets are instead sent to the Packet Forwarding Engine (where packet processing is implemented in hardware), use the `ip-options any` match condition. The 10-Gigabit Ethernet Modular Port Concentrator (MPC), 100-Gigabit Ethernet MPC, 60-Gigabit Ethernet MPC, 60-Gigabit Queuing Ethernet MPC, and 60-Gigabit Ethernet Enhanced Queuing MPC on MX Series routers are capable of parsing the IP option field of the IPv4 packet header. For interfaces configured on those MPCs, *all* packets that are matched using the `ip-options` match condition are sent to the Packet Forwarding Engine for processing. **NOTE**: • On MX series routers, filter matches using `ip-options` cannot be used with egress (output) filters.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`ip-options-except` *`values`*	Do not match the IP option field to the specified value or list of values. For details about specifying the *`values`*, see the `ip-options` match condition.
`is-fragment`	Using this condition causes a match if the More Fragments flag is enabled in the IP header or if the fragment offset is not zero. **NOTE**: To match both first and trailing fragments, you can use two terms that specify different match conditions (`first-fragment` and `is-fragment`).
`loss-priority` *`level`*	Match the packet loss priority (PLP) level. Specify a single level or multiple levels: `low`, `medium-low`, `medium-high`, or `high`. For IP traffic on M320, MX Series, and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), you must include the `tri-color` statement at the `[edit class-of-service]` hierarchy level to commit a PLP configuration with any of the four levels specified. If the `tri-color` statement is not enabled, you can only configure the `high` and `low` levels. This applies to all protocol families.
`loss-priority-except` *`level`*	Do not match the PLP level. For details, see the `loss-priority` match condition.
`packet-length` *`bytes`*	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead. You can also specify a range of values to be matched.
`packet-length-except` *`bytes`*	Do not match the length of the received packet, in bytes. For details, see the `packet-length` match type.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`port` *`number`*	Match the UDP or TCP source or destination port field. If you configure this match condition, you cannot configure the `destination-port` match condition or the `source-port` match condition in the same term. If you configure this match condition, we recommend that you also configure the `protocol udp` or `protocol tcp` match statement in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the text synonyms listed under `destination-port`.
`port-except` *`number`*	Do not match either the source or destination UDP or TCP port field. For details, see the `port` match condition.
`precedence` *`ip-precedence-value`*	Match the IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): `critical-ecp` (0xa0), `flash` (0x60), `flash-override` (0x80), `immediate` (0x40), `internet-control` (0xc0), `net-control` (0xe0), `priority` (0x20), or `routine` (0x00). You can specify precedence in hexadecimal, binary, or decimal form.
`precedence-except` *`ip-precedence-value`*	Do not match the IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): `critical-ecp` (0xa0), `flash` (0x60), `flash-override` (0x80), `immediate` (0x40), `internet-control` (0xc0), `net-control` (0xe0), `priority` (0x20), or `routine` (0x00). You can specify precedence in hexadecimal, binary, or decimal form.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`prefix-list` *`name`* `[ except ]`	Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the `except` option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list. The prefix list is defined at the `[edit policy-options prefix-list` *`prefix-list-name`*`]` hierarchy level.
`protocol` *`number`*	Match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `dstopts` (60), `egp` (8), `esp` (50), `fragment` (44), `gre` (47), `hop-by-hop` (0), `icmp` (1), `icmp6` (58), `icmpv6` (58), `igmp` (2), `ipip` (4), `ipv6` (41), `ospf` (89), `pim` (103), `rsvp` (46), `sctp` (132), `tcp` (6), `udp` (17), or `vrrp` (112).
`protocol-except` *`number`*	Do not match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `dstopts` (60), `egp` (8), `esp` (50), `fragment` (44), `gre` (47), `hop-by-hop` (0), `icmp` (1), `icmp6` (58), `icmpv6` (58), `igmp` (2), `ipip` (4), `ipv6` (41), `ospf` (89), `pim` (103), `rsvp` (46), `sctp` (132), `tcp` (6), `udp` (17), or `vrrp` (112).
`service-filter-hit`	Match a packet received from a filter where a `service-filter-hit` action was applied.
`source-address` *`address`* `[ except ]`	Match the IPv4 address of the source node sending the packet unless the `except` option is included. If the option is included, do not match the IPv4 address of the source node sending the packet. You cannot specify both the `address` and `source-address` match conditions in the same term.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`source-port` *`number`*	Match the UDP or TCP source port field. You cannot specify the `port` and `source-port` match conditions in the same term. If you configure this match condition for IPv4 traffic, we recommend that you also configure the `protocol udp` or `protocol tcp` match statement in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the text synonyms listed with the `destination-port` *`number`* match condition.
`source-port-except` *`number`*	Do not match the UDP or TCP source port field. For details, see the `source-port` match condition.
`source-prefix-list` *`name`* `[ except ]`	Match source prefixes in the specified list unless the `except` option is included. If the option is included, do not match the source prefixes in the specified list. Specify the name of a prefix list defined at the `[edit policy-options prefix-list` *`prefix-list-name`*`]` hierarchy level.
`tcp-established`	Match TCP packets of an established TCP session (packets other than the first packet of a connection). This is an alias for `tcp-flags "(ack | rst)"`. This match condition does not implicitly check that the protocol is TCP. To check this, specify the `protocol tcp` match condition.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`tcp-flags` *`value`*	Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header. To specify individual bit fields, you can specify the following text synonyms or hexadecimal values: • `fin` (0x01) • `syn` (0x02) • `rst` (0x04) • `push` (0x08) • `ack` (0x10) • `urgent` (0x20) In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet. You can string together multiple flags using the bit-field logical operators. For combined bit-field match conditions, see the `tcp-established` and `tcp-initial` match conditions. If you configure this match condition, we recommend that you also configure the `protocol tcp` match statement in the same term to specify that the TCP protocol is being used on the port. For IPv4 traffic only, this match condition does not implicitly check whether the datagram contains the first fragment of a fragmented packet. To check for this condition for IPv4 traffic only, use the `first-fragment` match condition.
`tcp-initial`	Match the initial packet of a TCP connection. This is an alias for `tcp-flags "(!ack & syn)"`. This condition does not implicitly check that the protocol is TCP. If you configure this match condition, we recommend that you also configure the `protocol tcp` match condition in the same term.

Table 38: Firewall Filter Match Conditions for IPv4 Traffic *(Continued)*

Match Condition	Description
`ttl` *`number`*	Match the IPv4 time-to-live number. Specify a TTL value or a range of TTL values. For *`number`*, you can specify one or more values from `0` through `255`. This match condition is supported only on M120, M320, MX Series, and T Series routers.
`ttl-except` *`number`*	Do not match on the IPv4 TTL number. For details, see the `ttl` match condition.

RELATED DOCUMENTATION

- Firewall Filter Match Conditions for IPv6 Traffic | **348**
- Firewall Filter Nonterminating Actions | **361**
- Firewall Filter Terminating Actions | **368**

Firewall Filter Match Conditions for IPv6 Traffic

You can configure a firewall filter with match conditions for Internet Protocol version 6 (IPv6) traffic (`family inet6`).

> **NOTE**: For MX Series routers with MPCs, you need to initialize the filter counter for Trio-only match filters by walking the corresponding SNMP MIB, for example, `show snmp mib walk` *`name`* `ascii`. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all enhanced mode firewall filters, filters with flexible conditions, and filters with the certain terminating actions. See those topics, listed under Related Documentation, for details.

Table 39 on page 349 describes the match conditions you can configure at the `[edit firewall family inet6 filter` *`filter-name`* `term` *`term-name`* `from]` hierarchy level.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic

Match Condition	Description
address *address* [except]	Match the IPv6 source or destination address field unless the except option is included. If the option is included, do not match the IPv6 source or destination address field.
destination-address *address* [except]	Match the IPv6 destination address field unless the except option is included. If the option is included, do not match the IPv6 destination address field. You cannot specify both the address and destination-address match conditions in the same term.
destination-port *number*	Match the UDP or TCP destination port field. You cannot specify both the port and destination-port match conditions in the same term. If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).
destination-port-except *number*	Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description	
`destination-prefix-list` *prefix-list-name* `[ except ]`	Match the IPv6 destination prefix to the specified list unless the `except` option is included. If the option is included, do not match the IPv6 destination prefix to the specified list. The prefix list is defined at the [`edit policy-options prefix-list` *prefix-list-name*] hierarchy level.	
`extension-header` *header-type*	Match an extension header type that is contained in the packet by identifying a Next Header value. **NOTE**: This match condition is only supported on MPCs in MX Series routers. In the first fragment of a packet, the filter searches for a match in any of the extension header types. When a packet with a fragment header is found (a subsequent fragment), the filter only searches for a match of the next extension header type because the location of other extension headers is unpredictable. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `destination` (60), `esp` (50), `fragment` (44), `hop-by-hop` (0), `mobility` (135), or `routing` (43). To match *any* value for the extension header option, use the text synonym `any`. For MX Series routers with MPCs, initialize new firewall filters that include this condition by walking the corresponding SNMP MIB.	
`extension-headers-except` *header-type*	Do not match an extension header type that is contained in the packet. For details, see the `extension-headers` match condition. **NOTE**: This match condition is only supported on MPCs in MX Series routers.	
`first-fragment`	Match if the packet is the first fragment.	
`flexible-match-mask` *value*	`bit-length`	Length of integer input (1..32 bits); (Optional) Length of string input (1..128 bits)

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description	
	`bit-offset`	Bit offset after the (match-start + byte) offset (0..7)
	`byte-offset`	Byte offset after the match start point
	`flexible-mask-name`	Select a flexible match from predefined template field
	`mask-in-hex`	Mask out bits in the packet data to be matched
	`match-start`	Start point to match in packet
	`prefix`	Value data/string to be matched
`flexible-match-range` *value* Ranges should use the following format: *Integer-Integer*	`bit-length`	Length of the data to be matched in bits (0..32)
	`bit-offset`	Bit offset after the (match-start + byte) offset (0..7)
	`byte-offset`	Byte offset after the match start point
	`flexible-range-name`	Select a flexible match from predefined template field
	`match-start`	Start point to match in packet
	`range`	Range of values to be matched
	`range-except`	Do not match this range of values

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`forwarding-class` *`class`*	Match the forwarding class of the packet. Specify `assured-forwarding`, `best-effort`, `expedited-forwarding`, or `network-control`.
`forwarding-class-except` *`class`*	Do not match the forwarding class of the packet. For details, see the `forwarding-class` match condition.
`hop-limit` *`hop-limit`*	Match the hop limit to the specified hop limit or set of hop limits. For *`hop-limit`*, specify a single value or a range of values from 0 through 255. Supported on interfaces hosted on MICs or MPCs in MX Series routers only. **NOTE**: This match condition is supported on PTX series routers when `enhanced-mode` is configured on the router.
`hop-limit-except` *`hop-limit`*	Do not match the hop limit to the specified hop limit or set of hop limits. For details, see the `hop-limit` match condition. Supported on interfaces hosted on MICs or MPCs in MX Series routers only. **NOTE**: This match condition is supported on PTX series routers when `enhanced-mode` is configured on the router.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`icmp-code` *`message-code`*	Match the ICMP message code field. If you configure this match condition, we recommend that you also configure the `next-header icmp` or `next-header icmp6` match condition in the same term. If you configure this match condition, you must also configure the `icmp-type` *`message-type`* match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated: • parameter-problem: `ip6-header-bad` (0), `unrecognized-next-header` (1), `unrecognized-option` (2) • time-exceeded: `ttl-eq-zero-during-reassembly` (1), `ttl-eq-zero-during-transit` (0) • destination-unreachable: `administratively-prohibited` (1), `address-unreachable` (3), `no-route-to-destination` (0), `port-unreachable` (4)
`icmp-code-except` *`message-code`*	Do not match the ICMP message code field. For details, see the `icmp-code` match condition.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description	
`icmp-type` *`message-type`*	Match the ICMP message type field. If you configure this match condition, we recommend that you also configure the `next-header icmp` or `next-header icmp6` match condition in the same term. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `certificate-path-advertisement` (149), `certificate-path-solicitation` (148), `destination-unreachable` (1), `echo-reply` (129), `echo-request` (128), `home-agent-address-discovery-reply` (145), `home-agent-address-discovery-request` (144), `inverse-neighbor-discovery-advertisement` (142), `inverse-neighbor-discovery-solicitation` (141), `membership-query` (130), `membership-report` (131), `membership-termination` (132), `mobile-prefix-advertisement-reply` (147), `mobile-prefix-solicitation` (146), `neighbor-advertisement` (136), `neighbor-solicit` (135), `node-information-reply` (140), `node-information-request` (139), `packet-too-big` (2), `parameter-problem` (4), `private-experimentation-100` (100), `private-experimentation-101` (101), `private-experimentation-200` (200), `private-experimentation-201` (201), `redirect` (137), `router-advertisement` (134), `router-renumbering` (138), `router-solicit` (133), or `time-exceeded` (3). For `private-experimentation-201` (201), you can also specify a range of values within square brackets.	
`icmp-type-except` *`message-type`*	Do not match the ICMP message type field. For details, see the `icmp-type` match condition.	
`is-fragment`	Match if the packet is a fragment.	
`last-fragment`	Match if the packet is the last fragment.	

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`loss-priority` *`level`*	Match the packet loss priority (PLP) level. Specify a single level or multiple levels: `low`, `medium-low`, `medium-high`, or `high`. For IP traffic on MX Series routers with Enhanced II Flexible PIC Concentrators (FPCs), you must include the `tri-color` statement at the `[edit class-of-service]` hierarchy level to commit a PLP configuration with any of the four levels specified. If the `tri-color` statement is not enabled, you can only configure the `high` and `low` levels. This applies to all protocol families.
`loss-priority-except` *`level`*	Do not match the PLP level. For details, see the `loss-priority` match condition.
`next-header` *`header-type`*	Match the first 8-bit Next Header field in the packet. Support for the `next-header` firewall match condition is available in Junos OS Release 13.3R6 and later. For IPv6, we recommend that you use the `payload-protocol` term rather than the `next-header` term when configuring a firewall filter with match conditions. Although either can be used, `payload-protocol` provides the more reliable match condition because it uses the actual payload protocol to find a match, whereas `next-header` simply takes whatever appears in the first header following the IPv6 header, which may or may not be the actual protocol. In addition, if `next-header` is used with IPv6, the accelerated filter block lookup process is bypassed and the standard filter used instead. Match the first 8-bit Next Header field in the packet. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): `ah` (51), `dstops` (60), `egp` (8), `esp` (50), `fragment` (44), `gre` (47), `hop-by-hop` (0), `icmp` (1), `icmp6` (58), `icmpv6` (58), `igmp` (2), `ipip` (4), `ipv6` (41), `mobility` (135), `no-next-header` (59), `ospf` (89), `pim` (103), `routing` (43), `rsvp` (46), `sctp` (132), `tcp` (6), `udp` (17), or `vrrp` (112). **NOTE**: `next-header icmp6` and `next-header icmpv6` match conditions perform the same function. `next-header icmp6` is the preferred option. `next-header icmpv6` is hidden in the Junos OS CLI.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`next-header-except` *`header-type`*	Do not match the 8-bit Next Header field that identifies the type of header between the IPv6 header and payload. For details, see the `next-header` match type.
`packet-length` *`bytes`*	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
`packet-length-except` *`bytes`*	Do not match the length of the received packet, in bytes. For details, see the `packet-length` match type.
`payload-protocol` *`protocol-type`*	Match the payload protocol type. In place of the *`protocol-type`* numeric value, you can specify one of the following text synonyms (the field values are also listed): specify one or a set of of the following: `ah` (51), `dstopts` (60), `egp` (8), `esp` (50), `fragment` (44), `gre` (47), `hop-by-hop` (0), `icmp` (1), `icmp6` (58, `igmp` (2), `ipip` (4), `ipv6` (41), `no-next-header`, `ospf` (89), `pim` (103), `routing`, `rsvp` (46), `sctp` (132), `tcp` (6), `udp` (17), or `vrrp` (112) (dstopts (60), fragment (44), hop-by-hop 0), and routing are not available in Junos OS Release 16.1 and later). You can also use the `payload-protocol` condition to match an extension header type that the Juniper Networks firmware cannot interpret. You can specify a range of extension header values within square brackets. When the firmware finds the first extension header type that it cannot interpret in a packet, the `payload-protocol` value is set to that extension header type. The firewall filter only examines the first extension header type that the firmware cannot interpret in the packet. **NOTE**: This match condition is only supported on MPCs on MX Series Routers. Initialize new firewall filters that include this condition by walking the corresponding SNMP MIB.
`payload-protocol-except` *`protocol-type`*	Do not match the payload protocol type. For details, see the `payload-protocol` match type. **NOTE**: This match condition is only supported on MPCs on MX Series Routers

Table 39: Firewall Filter Match Conditions for IPv6 Traffic ***(Continued)***

Match Condition	Description
port *number*	Match the UDP or TCP source or destination port field. If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term. If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the text synonyms listed under destination-port.
port-except *number*	Do not match the UDP or TCP source or destination port field. For details, see the port match condition.
prefix-list *prefix-list-name* [except]	Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the except option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list. The prefix list is defined at the [edit policy-options prefix-list *prefix-list-name*] hierarchy level.
service-filter-hit	Match a packet received from a filter where a service-filter-hit action was applied.
source-address *address* [except]	Match the IPv6 address of the source node sending the packet unless the except option is included. If the option is included, do not match the IPv6 address of the source node sending the packet. You cannot specify both the address and source-address match conditions in the same term.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`source-port` *`number`*	Match the UDP or TCP source port field. You cannot specify the `port` and `source-port` match conditions in the same term. If you configure this match condition, we recommend that you also configure the `next-header udp` or `next-header tcp` match condition in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the text synonyms listed with the `destination-port` *`number`* match condition.
`source-port-except` *`number`*	Do not match the UDP or TCP source port field. For details, see the `source-port` match condition.
`source-prefix-list` *`name`* `[ except ]`	Match the IPv6 address prefix of the packet source field unless the `except` option is included. If the option is included, do not match the IPv6 address prefix of the packet source field. Specify a prefix list name defined at the `[edit policy-options prefix-list` *`prefix-list-name`*`]` hierarchy level.
`tcp-established`	Match TCP packets other than the first packet of a connection. This is a text synonym for `tcp-flags "(ack | rst)"` (`0x14`). **NOTE**: This condition does not implicitly check that the protocol is TCP. To check this, specify the `protocol tcp` match condition. If you configure this match condition, we recommend that you also configure the `next-header tcp` match condition in the same term.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`tcp-flags` *`flags`*	Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header. To specify individual bit fields, you can specify the following text synonyms or hexadecimal values: • `fin` (0x01) • `syn` (0x02) • `rst` (0x04) • `push` (0x08) • `ack` (0x10) • `urgent` (0x20) In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet. You can string together multiple flags using the bit-field logical operators. For combined bit-field match conditions, see the `tcp-established` and `tcp-initial` match conditions. If you configure this match condition, we recommend that you also configure the `next-header tcp` match condition in the same term to specify that the TCP protocol is being used on the port.
`tcp-initial`	Match the initial packet of a TCP connection. This is a text synonym for `tcp-flags "(! ack & syn)"`. This condition does not implicitly check that the protocol is TCP. If you configure this match condition, we recommend that you also configure the `next-header tcp` match condition in the same term.

Table 39: Firewall Filter Match Conditions for IPv6 Traffic *(Continued)*

Match Condition	Description
`traffic-class` *number*	Match the 8-bit field that specifies the class-of-service (CoS) priority of the packet. This field was previously used as the type-of-service (ToS) field in IPv4. You can specify a numeric value from `0` through `63`. To specify the value in hexadecimal form, include `0x` as a prefix. To specify the value in binary form, include `b` as a prefix. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*, defines one code point: `ef` (46). • RFC 2597, *Assured Forwarding PHB Group*, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: • `af11` (10), `af12` (12), `af13` (14) • `af21` (18), `af22` (20), `af23` (22) • `af31` (26), `af32` (28), `af33` (30) • `af41` (34), `af42` (36), `af43` (38)
`traffic-class-except` *number*	Do not match the 8-bit field that specifies the CoS priority of the packet. For details, see the `traffic-class` match description.

NOTE: If you specify an IPv6 address in a match condition (the `address`, `destination-address`, or `source-address` match conditions), use the syntax for text representations described in RFC 4291, *IP Version 6 Addressing Architecture*.

RELATED DOCUMENTATION

Firewall Filter Match Conditions for IPv4 Traffic | **334**

Firewall Filter Nonterminating Actions | **361**

Firewall Filter Nonterminating Actions

Firewall filters support different sets of nonterminating actions for each protocol family, which include an implicit accept action. In this context, *nonterminating* means that other actions can follow these actions whereas no other actions can follow a *terminating* action. As such, you cannot configure the `next term` action with a *terminating* action in the same filter term. You can, however, configure the `next term` action with another *nonterminating* action in the same filter term.

Table 40 on page 361 describes the nonterminating actions you can configure for a firewall filter term.

Table 40: Nonterminating Actions for Firewall Filters

Nonterminating Action	Description	Protocol Families
`count` *`counter-name`*	Count the packet in the named counter.	• `family any` • `family inet` • `family inet6`
`dont-fragment (set | clear)`	Configure the value of the Don’t Fragment bit (flag) in the IPv4 header to specify whether the datagram can be fragmented: • `set`—Change the flag value to one, preventing fragmentation. • `clear`—Change the flag value to zero, allowing fragmentation. **NOTE**: The `dont-fragment (set | clear)` actions are supported only on MPCs.	`family inet`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
`dscp` *`value`*	Set the IPv4 Differentiated Services code point (DSCP) bit. You can specify a numerical value from `0` through `63`. To specify the value in hexadecimal form, include `0x` as a prefix. To specify the value in binary form, include `b` as a prefix. The default DSCP value is `be` (best effort), or `0`. You can also specify one of the following text synonyms: • `af11`—Assured forwarding class 1, low drop precedence (1) • `af12`—Assured forwarding class 1, medium drop precedence (2) • `af13`—Assured forwarding class 1, high drop precedence (3); and so on through `af43`, Assured forwarding class 4, high drop precedence • `be`—Best effort • `cs0`—Class selector 0; and so on through `cs7`, Class selector 0 • `ef`—Expedited forwarding **NOTE**: MPC line cards running on MX series routers support any value (from 0 to 63) in conjunction with the `set dscp` firewall filter action.	`family inet`
`force-premium`	By default, a hierarchical policer processes the traffic it receives according to the traffic's forwarding class. Premium, expedited-forwarding traffic, has priority for bandwidth over aggregate, best-effort traffic. The `force-premium` filter ensures that traffic matching the term is treated as premium traffic by a subsequent hierarchical policer, regardless of its forwarding class. This traffic is given preference over any aggregate traffic received by that policer. **NOTE**: The `force-premium` filter option is supported only on MPCs.	• `family any` • `family inet` • `family inet6`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
`forwarding-class` *`class-name`*	Classify the packet to the named forwarding class: • *forwarding-class-name* • `assured-forwarding` • `best-effort` • `expedited-forwarding` • `network-control`	• `family any` • `family inet` • `family inet6`
`hierarchical-policer`	Police the packet using the specified hierarchical policer	• `family any` • `family inet` • `family inet6`
`log`	Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the `show firewall log` command at the command-line interface (CLI). **NOTE**: The Layer 2 (L2) families log action is available only for MX Series routers with MPCs (MPC mode if the router has only MPCs, or mix mode if it has MPCs and DCPs). For MX Series routers with DPCs, the log action for L2 families is ignored if configured.	• `family inet` • `family inet6`
`loss-priority (high | medium-high | medium-low | low)`	Set the packet loss priority (PLP) level. You cannot also configure the `three-color-policer` nonterminating action for the same firewall filter term. These two nonterminating actions are mutually exclusive. For IP traffic on MX Series routers with Enhanced II Flexible PIC Concentrators (FPCs), you must include the `tri-color` statement at the `[edit class-of-service]` hierarchy level to commit a PLP configuration with any of the four levels specified. If the `tri-color` statement is not enabled, you can only configure the `high` and `low` levels. This applies to all protocol families.	• `family any` • `family inet` • `family inet6`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
`next`	Continue to the next term in a filter.	• `family inet` • `family inet6`
`next-ip` *`ip-address`*	Direct packets to the specified destination IPv4 address.	`family inet`
`next-ip6` *`ipv6-address`*	Direct packets to the specified destination IPv6 address.	`family inet6`
`policer` *`policer-name`*	Name of policer to use to rate-limit traffic.	• `family any` • `family inet` • `family inet6`
`port-mirror` *`instance-name`*	Port-mirror the packet based on the specified family. This action is supported on M120 routers, M320 routers configured with Enhanced III FPCs, MX Series routers, and PTX Series Packet Transport Routers only. We recommend that you do not use both the `next-hop-group` and the `port-mirror` actions in the same firewall filter.	• `family any` • `family inet` • `family inet6`
`routing-instance` *`routing-instance-name`*	Direct packets to the specified routing instance.	• `family inet` • `family inet6` •
`sample`	Sample the packet. **NOTE**: Junos OS does not sample packets originating from the router. If you configure a filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.	• `family inet` • `family inet6`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
service-accounting	Use the inline counting mechanism when capturing subscriber per-service statistics. Count the packet for service accounting. The count is applied to a specific named counter (__junos-dyn-service-counter) that RADIUS can obtain. The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.	• family any • family inet • family inet6
service-accounting-deferred	Use the deferred counting mechanism when capturing subscriber per-service statistics. The count is applied to a specific named counter (__junos-dyn-service-counter) that RADIUS can obtain. The service-accounting and service-accounting-deferred keywords are mutually exclusive, both per-term and per-filter.	• family any • family inet • family inet6
service-filter-hit	(Only if the service-filter-hit flag is marked by a previous filter in the current type of chained filters) Direct the packet to the next type of filters. Indicate to subsequent filters in the chain that the packet was already processed. This action, coupled with the service-filter-hit match condition in receiving filters, helps to streamline filter processing.	• family any • family inet • family inet6
syslog	Log the packet to the system log file. The syslog firewall action for existing inet and inet6 families, and the syslog action in L2 family filters includes the following L2 information: Input interface, action, VLAN ID1, VLAN ID2, Ethernet type, source and destination MAC addresses, protocol, source and destination IP addresses, source and destination ports, and the number of packets. **NOTE**: The L2 families syslog action is available only for MX Series routers with MPCs (MPC mode if the router has only MPCs, or mix mode if it has MPCs and DCPs). For MX Series routers with DPCs, the syslog action for L2 families is ignored if configured.	• family inet • family inet6

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
`three-color-policer (single-rate | two-rate)` *`policer-name`*	Police the packet using the specified single-rate or two-rate three-color-policer. **NOTE**: You cannot also configure the `loss-priority` action for the same firewall filter term. These two actions are mutually exclusive.	• `family inet` • `family inet6`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
`traffic-class` *`value`*	Specify the traffic-class code point. You can specify a numerical value from `0` through `63`. To specify the value in hexadecimal form, include `0x` as a prefix. To specify the value in binary form, include `b` as a prefix. The default traffic-class value is best effort, that is, `be` or `0`. In place of the numeric value, you can specify one of the following text synonyms: • `af11`—Assured forwarding class 1, low drop precedence • `af12`—Assured forwarding class 1, medium drop precedence • `af13`—Assured forwarding class 1, high drop precedence • `af21`—Assured forwarding class 2, low drop precedence • `af22`—Assured forwarding class 2, medium drop precedence • `af23`—Assured forwarding class 2, high drop precedence • `af31`—Assured forwarding class 3, low drop precedence • `af32`—Assured forwarding class 3, medium drop precedence • `af33`—Assured forwarding class 3, high drop precedence • `af41`—Assured forwarding class 4, low drop precedence • `af42`—Assured forwarding class 4, medium drop precedence • `af43`—Assured forwarding class 4, high drop precedence • `be`—Best effort • `cs0`—Class selector 0 • `cs1`—Class selector 1 • `cs2`—Class selector 2	• `family inet6`

Table 40: Nonterminating Actions for Firewall Filters *(Continued)*

Nonterminating Action	Description	Protocol Families
	• `cs3`—Class selector 3 • `cs4`—Class selector 4 • `cs5`—Class selector 5 • `cs6`—Class selector 6 • `cs7`—Class selector 7 • `ef`—Expedited forwarding	

RELATED DOCUMENTATION

Firewall Filter Terminating Actions | **368**

Firewall Filter Match Conditions for IPv4 Traffic | **334**

Firewall Filter Match Conditions for IPv6 Traffic | **348**

Firewall Filter Terminating Actions

Firewall filters support a set of terminating actions for each protocol family. A filter-terminating action halts all evaluation of a firewall filter for a specific packet. The router performs the specified action, and no additional terms are examined.

NOTE: You cannot configure the **next term** action with a *terminating* action in the same filter term. However, you can configure the **next term** action with another *nonterminating* action in the same filter term.

For MX Series routers with MPCs, you need to initialize the filter counter for Trio-only match filters by walking the corresponding SNMP MIB, for example, `show snmp mib walk` *`name`* `ascii`. This forces Junos to learn the filter counters and ensure that the filter statistics are displayed. This guidance applies to all enhanced mode firewall filters, filters with flexible conditions, and filters with the certain terminating actions. See those topics, listed under Related Documentation, for details.

Table 39 on page 349 describes the terminating actions you can specify in a firewall filter term.

Table 41: Terminating Actions for Firewall Filters

Terminating Action	Description	Protocols
accept	Accept the packet.	• `family any` • `family inet` • `family inet6`
`discard`	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.	• `family any` • `family inet` • `family inet6`

Table 41: Terminating Actions for Firewall Filters *(Continued)*

Terminating Action	Description	Protocols
`exclude-accounting`	Exclude the packet from being included in accurate accounting statistics for tunneled subscribers on an L2TP LAC. Typically used in filters that match DHCPv6 or ICMPv6 control traffic Failure to exclude these packets results in the idle-timeout detection mechanism considering these packets as data traffic, causing the timeout to never expire. (The idle timeout is configured with the `client-idle-timeout` and `client-idle-timeout-ingress-only` statements in the access profile session options.) The term excludes packets from being included in counts for both family accurate accounting and service accurate accounting. The packets are still included in the session interface statistics.	• `family inet` • `family inet6`

Table 41: Terminating Actions for Firewall Filters *(Continued)*

Terminating Action	Description	Protocols
	The term is available for both `inet` and `inet6` families, but is used only for `inet6`.	

Table 41: Terminating Actions for Firewall Filters *(Continued)*

Terminating Action	Description	Protocols
`reject` *`message-type`*	Reject the packet and return an ICMPv4 or ICMPv6 message: - If no *`message-type`* is specified, a `destination unreachable` message is returned by default. - If `tcp-reset` is specified as the *`message-type`*, `tcp-reset` is returned only if the packet is a TCP packet. Otherwise, the `administratively-prohibited` message, which has a value of 13, is returned. - If any other *`message-type`* is specified, that message is returned. The *`message-type`* can be one of the following values: `address-unreachable`, `administratively-prohibited`, `bad-host-`	- `family inet` - `family inet6`

Table 41: Terminating Actions for Firewall Filters *(Continued)*

Terminating Action	Description	Protocols
	`tos`, `bad-network-tos`, `beyond-scope`, `fragmentation-needed`, `host-prohibited`, `host-unknown`, `host-unreachable`, `network-prohibited`, `network-unknown`, `network-unreachable`, `no-route`, `port-unreachable`, `precedence-cutoff`, `precedence-violation`, `protocol-unreachable`, `source-host-isolated`, `source-route-failed`, or `tcp-reset`.	

RELATED DOCUMENTATION

Firewall Filter Nonterminating Actions | 361

Firewall Filter Match Conditions for IPv4 Traffic | 334

Firewall Filter Match Conditions for IPv6 Traffic | 348

aaa-options (Access Profile)

IN THIS SECTION

Syntax

```
aaa-options aaa-options-name {
    aaa-context aaa-context-name;
    access-profile profile-name;
    subscriber-context subscriber-context-name
}
```

Hierarchy Level

```
[edit access]
```

Description

Define a set of AAA options for authorizing and configuring a subscriber or set of subscribers with a subscriber access profile.

Options

aaa-options-name	Name of the set of options.

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

access—To view this statement in the configuration.

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

access-profile-name (Duplicate Accounting)

IN THIS SECTION

Syntax

```
access-profile-name [profile-name];
```

Hierarchy Level

```
[edit access profile profile-name accounting duplication-vrf]
```

Description

Specify up to five access profiles, all in the same nondefault VRF (LS:RI combination), each of which lists one or more RADIUS accounting servers to which duplication accounting information is sent.

Options

profile-name Name of an access profile that lists RADIUS accounting servers for duplicate reporting.

Required Privilege Level

admin—To view this statement in the configuration.

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

accounting-profile

IN THIS SECTION

Syntax

```
accounting-profile name;
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range name]
```

Description

Enable collection of accounting data for the specified physical or logical interface or interface range.

Options

name—Name of the accounting profile.

Required Privilege Level

interface—To view this statement in the configuration.

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

accounting (Access Profile)

IN THIS SECTION

Syntax

```
accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    address-change-immediate-update;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    duplication-vrf {
        access-profile-name profile-name;
        vrf-name vrf-name;
    }
    immediate-update;
    order [accounting-method];
    send-acct-status-on-config-change
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately. Search for a statement in CLI Explorer for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

action

IN THIS SECTION

Syntax

```
action {
    loss-priority high then discard;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer name],
[edit logical-systems logical-system-name firewall three-color-policer name]
```

Description

Discard traffic on a logical interface using tricolor marking policing.

NOTE: This statement is supported only on IQ2 interfaces.

The remaining statement is explained separately. Search for a statement in CLI Explorer for details.

Required Privilege Level

firewall—To view this statement in the configuration.

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

address

IN THIS SECTION

Syntax MX Series (dynamic-profiles)

```
address (ip-address | ipv6-address);
```

MX Series (dynamic-profiles)

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family],
[edit dynamic-profiles profile-name interfaces pp0 unit  “$junos-interface-unit” family family],
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family family]
```

Description

Configure the interface address.

Options

ip-address—IPv4 address of the interface.

ipv6-address—IPv6 address of the interface. When configuring an IPv6 address on a dynamically created interface, use the *$junos-ipv6-address* dynamic variable.

Required Privilege Level

interface—To view this statement in the configuration.

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

address-assignment (Address-Assignment Pools)

IN THIS SECTION

Syntax

```
address-assignment {

    neighbor-discovery-router-advertisement ndra-pool-name;
    pool pool-name {
        active-drain;
        family family {
            dhcp-attributes {
                protocol-specific attributes;
            }
            excluded-address ip-address;
            excluded-range name low minimum-value high maximum-value;
            host hostname {
                hardware-address mac-address;
                ip-address ip-address;
            }
            network ip-prefix/<prefix-length>;
            prefix ipv6-prefix;
            range range-name {
                high upper-limit;
                low lower-limit;
                prefix-length prefix-length;
            }
        }
```

```
            hold-down;
            link pool-name;
        }
    }
```

Hierarchy Level

```
[edit access]
```

Description

Configure address-assignment pools that can be used by different client applications.

Options

neighbor-discovery-router-advertisement Configure the name of the address-assignment pool used to assign the router advertisement prefix.

- **Values:** *ndra-pool-name*—Name of the address-assignment pool.

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

admin—To view this statement in the configuration.

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

adf (Dynamic Firewalls)

IN THIS SECTION

Syntax

```
adf {
    counter;
    input-precedence precedence;
    not-mandatory;
    output-precedence precedence;
    rule rule-value;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter]
```

Description

Configure an Ascend-Data-Filter that the dynamic profile applies to a subscriber session.

Options

counter—Enable a counter that increments each time the Ascend-Data-Filter rule is used. Typically used for testing purposes.

not-mandatory—Suppress router from reporting an error when the RADIUS reply message does not include the $junos-adf-rule-v4 or $junos-adf-rule-v6 variable that is configured for the Ascend-Data-Filter in the dynamic profile. In this circumstance, the Ascend-Data-Filter is not created.

precedence—Precedence value that sets the order in which dynamic service filters are applied on the interface. The lower the precedence value, the higher the precedence that is given. The precedence setting is used in conjunction with the precedence settings of all dynamic service filters configured (not only Ascend-Data-Filters) on the same interface to establish the order. For example, the order also includes any configured input *filter-name* precedence *precedence* and output *filter-name* precedence *precedence* statements.

- **Range:** 0 through 255
- **Default:** 0

rule-value—Ascend-Data-Filter rule. You can specify either a Junos predefined variable that maps the Ascend-Data-Filter actions to Junos filter functionality or you can manually configure the Ascend-Data-Filter rule. The router supports two predefined variables depending on family type: $junos-adf-rule-v4 for family inet and $junos-adf-rule-v6 for family inet6.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

No Link Title

No Link Title

No Link Title

Guidelines for Configuring Service Filters

adjustment-control-profiles

IN THIS SECTION

Syntax

```
adjustment-control-profiles {
    profile-name {
        application {
            (ancp | dhcp-tags | pppoe-tags | radius-coa)
                priority priority;
                algorithm algorithm;
            }
        }
    }
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure the CoS adjustment control profile.

Options

profile-name	Name of the adjustment control profile.

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

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

aggregate

IN THIS SECTION

Syntax

```
aggregate {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
    }
    then {
        discard;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer name],
[edit firewall hierarchical-policer]
```

Description

On MX Series routers with Enhanced Intelligent Queuing (IQE) PICs, configure an aggregate hierarchical policer.

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

firewall—To view this statement in the configuration.

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

ancp (Adjustment Control Profiles)

IN THIS SECTION

Syntax

```
ancp {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the ANCP application.

Options

priority Priority of the ANCP application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 1

algorithm Rate adjustment algorithm used by the ANCP application.

- Values:
    - adjust-never—Do not perform rate adjustments.
    - adjust-always—Adjust the shaping rate unconditionally.
    - adjust-less—Adjust the shaping rate if it is less than the configured value.
    - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
    - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
    - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-always

Required Privilege Level

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

application (Adjustment Control Profiles)

IN THIS SECTION

Syntax

```
application {
    (ancp | dhcp-tags | pppoe-tags | radius-coa)
        priority priority;
        algorithm algorithm;
    }
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name]
```

Description

Configure which applications in the adjustment control profile can make shaping rate adjustments.

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

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

auto-configure (Demux)

IN THIS SECTION

Syntax

```
auto-configure {
    address-ranges {
     }
        dynamic-profile profile-name {
            network ip-address {
                range name {
                    low lower-limit;
                    high upper-limit;
                }
            }
        }
    }

}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet]
[edit interfaces interface-name unit unit-number demux inet6]
```

Description

Enable the configuration of dynamic, auto-sensed subscriber interfaces for the demultiplexing (demux) interface options. The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

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

auto-configure

IN THIS SECTION

Syntax

```
auto-configure {
    vlan-ranges {
        access-profile profile-name;
     }
        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;
        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.

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.

bandwidth-limit (Hierarchical Policer)

IN THIS SECTION

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]
```

Description

Configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.

Options

bps—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- • 32,000 through 18,446,744,073,709,551,615 on MX Series routers

> **NOTE**: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

Required Privilege Level

firewall—To view this statement in the configuration.

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

bandwidth-limit (Policer)

IN THIS SECTION

Syntax

```
bandwidth-limit bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding],
[edit firewall policer policer-name if-exceeding],
[edit logical-systems logical-system-name policer policer-name if-exceeding]
```

Description

For a single-rate two-color policer, configure the bandwidth limit as a number of bits per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the bandwidth limit as an absolute number of bits per second. Alternatively, for single-rate two-color policers only, you can use the `bandwidth-percent` *`percentage`* statement to specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

bps—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- • (MX Series routers) 8000 through 18,446,744,073,709,551,615

  > **NOTE**: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

- **Default:** None.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

bandwidth-percent

IN THIS SECTION

Syntax

```
bandwidth-percent percentage;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding],
[edit firewall policer policer-name if-exceeding],
[edit logical-systems logical-system-name policer policer-name if-exceeding]
```

Description

For a single-rate two-color policer, configure the bandwidth limit as a percentage value. Single-rate two-color policing uses the *single token bucket algorithm* to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate. Alternatively, you can use the `bandwidth-limit` *`bps`* statement to specify the bandwidth limit as an absolute number of bits per second.

The function of the bandwidth limit is extended by the burst size (configured using the `burst-size-limit` *`bytes`* statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

`percentage`—Traffic rate as a percentage of either the physical interface media rate or the logical interface configured shaping rate. You can configure a shaping rate on a logical interface by using class-of-service statement.

NOTE: The bandwidth percentage policer cannot be used to rate-limit tunnel or software interfaces, or for forwarding table filters. It is only valid for interface-specific filters. When used for matching bandwidth or burst-size on aggregated Ethernet or SONET bundles, bandwidth percentage policers must be used in conjunction with `shared-bandwidth-policer`.

- **Range:** 0 through 100
- **Default:** None.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

burst-size-limit (Hierarchical Policer)

IN THIS SECTION

Syntax

```
burst-size-limit bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]
```

Description

On MPCs hosted on MX Series routers configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.

Options

bytes—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 2,147,450,880 (1500 through 100,000,000,000 on MPCs hosted on MX Series routers)

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

captive-portal-content-delivery (Captive Portal Content Delivery)

IN THIS SECTION

Syntax

```
captive-portal-content-delivery {
    profile name
        dynamic;
    }
    rule rule-name {
        match-direction (input | output | input-output);
        term term-name {
            then {
                accept;
                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 dynamic-profiles profile-name services],
[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 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.

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.

classifiers (CoS)

IN THIS SECTION

Syntax

```
classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
        forwarding-class forwarding-class-name {
        loss-priority (high  | low  | medium-high  | medium-low) {
            code-point alias-or-bit-string ;
        }
        import (default | user-defined;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure a user-defined behavior aggregate (BA) classifier.

Options

- *classifier-name*—User-defined name for the classifier.
- import (default | *user-defined*)—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type `dscp` and you specify `import default`, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify `import mymap`, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named `mymap`.
- forwarding-class *class-name*—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
- loss-priority *level*—Specify a loss priority for this forwarding class: `high`, `low`, `medium-high`, `medium-low`.

- code-points (*alias* | *bits*)—Specify a code-point alias or the code points that map to this forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

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

chassis (Subscriber Limits)

IN THIS SECTION

Syntax

```
chassis {
    limit limit;
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type (Subscriber Limits) name]
```

Description

Configure the maximum number of subscribers of the specified client type allowed to be logged in on the chassis. When that number is reached, subsequent logins on the chassis are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MIC, and per MPC.

Options

limit Maximum number of subscribers.

- **Range:** `1` through `1,000,000`

Required Privilege Level

system—To view this statement in the configuration.

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

client-type (Subscriber Limits)

IN THIS SECTION

Syntax

```
client-type (any | dhcp | l2tp | pppoe) {
    chassis (Subscriber Limits) {
        limit limit;
    }
    fpc (Subscriber Limits) slot-number {
        limit limit;
        pic (Subscriber Limits) number {
            limit limit;
            port (Subscriber Limits) number {
                limit limit;
            }
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit (Resource Monitor)]
```

Description

Configure the maximum number of subscribers of the client type that are allowed to be logged in. You can configure limits for subscribers per chassis, per MPC, per MIC, and per port. When the configured maximum number of subscribers is logged in for any level, subsequent logins at that level are denied until the current number of subscribers drops below the maximum allowed.

Options

name   Type of client for which subscriber limits are configured.

- Values:
  - any—Apply the limit to the sum of all DHCP, L2TP, and PPPoE clients.
  - dhcp—Apply the limit to DHCP clients.
  - l2tp—Apply the limit to L2TP clients.
  - pppoe—Apply the limit to PPPoE clients.

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

system—To view this statement in the configuration.

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

color-aware

IN THIS SECTION

Syntax

```
color-aware;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the way preclassified packets are metered. In color-aware mode, the local router can assign a higher packet loss priority, but cannot assign a lower packet loss priority.

For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.

- If the local router applies color-aware policing to the packet, the router *cannot* change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface.

- If the local router applies color-blind policing to the packet, the router *can* change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.

NOTE: A color-aware policer cannot be applied to Layer 2 traffic.

Default

If you omit the `color-aware` statement, the default behavior is color-aware mode.

Required Privilege Level

firewall—To view this statement in the configuration.

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

color-blind

IN THIS SECTION

Syntax

```
color-blind;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the way preclassified packets are metered. In color-blind mode, the local router ignores the preclassification of packets and can assign a higher or lower packet loss priority.

For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.

- If the local router applies color-aware policing to the packet, the router *cannot* change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface.

  NOTE: A color-aware policer cannot be applied to Layer 2 traffic.

- If the local router applies color-blind policing to the packet, the router *can* change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.

Default

If you omit the `color-blind` statement, the default behavior is color-aware mode.

Required Privilege Level

firewall—To view this statement in the configuration.

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

committed-burst-size

IN THIS SECTION

Syntax

```
committed-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the committed burst size (CBS) as a number of bytes.

NOTE: When you include the `committed-burst-size` statement in the configuration, you must also include the `committed-information-rate` statement at the same hierarchy level.

In three-color policing, a committed information rate (CIR) defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the CBS. Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the `excess-burst-size` statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the `peak-information-rate` and `peak-burst-rate` statements included in the policer configuration.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

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

committed-information-rate

IN THIS SECTION

Syntax

```
committed-information-rate bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name single-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a three-color policer, configure the committed information rate as a number of bits per second. The committed information rate (CIR) is the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions.

NOTE: When you include the `committed-information-rate` statement in the configuration, you must also include the `committed-burst-size` statement at the same hierarchy level.

In three-color policing, a CIR defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the committed burst size (CBS). Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the `excess-burst-size` statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the `peak-information-rate` and `peak-burst-rate` statements included in the policer configuration.

Options

bps—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

Range:

- • 1500 through 18,446,744,073,709,551,615 bps on MX Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

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

destination-profile

IN THIS SECTION

Syntax

```
destination-profile name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address],
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-address
interface-name destination address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet unnumbered-address interface-name destination address]
```

Description

For interfaces with PPP encapsulation, assign PPP properties to the remote destination end. You define the profile at the `[edit access group-profile name ppp]` hierarchy level.

Options

name—Profile name defined at the `[edit access group-profile name ppp]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

device-count (Pseudowire Subscriber Interfaces)

IN THIS SECTION

Syntax

```
device-count number;
```

Hierarchy Level

```
[edit chassis pseudowire-service]
```

Description

Configure the number of pseudowire logical devices available to the router. The statement also defines the available interface names for the pseudowire interfaces.

NOTE: When you subsequently configure the pseudowire interfaces, you must specify the interface names in the range from ps0 up to ps*(device-count - 1)*. For example, if you set the maximum number of devices to 5, then you can only configure interfaces ps0, ps1, ps2, ps3, and ps4. If you specify an interface name outside that range, the pseudowire interface is not created.

Options

number Number of devices.

- **Range:** 1 through 7000, 1 through 18000 for MX2010 and MX2020 routers with the MX2K-MPC9E or MX2K-MPC11E line card

Required Privilege Level

interface—To view this statement in the configuration.

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

drop-profiles

IN THIS SECTION

Syntax

```
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability [values];
            fill-level [values]
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Define drop profiles for RED.

For a packet to be dropped, it must match the drop profile. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the packet.

Options

profile-name—Name of the drop profile.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

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

dynamic-profiles

IN THIS SECTION

Syntax

```
dynamic-profiles {
    profile-name {
        class-of-service {
            dynamic-class-of-service-options {
                vendor-specific-tags  tag;
            }
            interfaces {
                interface-name ;
                }
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-
profile);
```

```
                    report-ingress-shaping-rate bps;
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                        }
                    }
                }
            }
            scheduler-maps {
                map-name {
                    forwarding-class class-name scheduler scheduler-name;
                }
            }
            schedulers {
                (scheduler-name) {
                    buffer-size (seconds | percent percentage | remainder | temporal
microseconds);
                    drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
protocol (any | non-tcp | tcp) drop-profile profile-name;
                    excess-priority (low | high | $junos-cos-scheduler-excess-priority);
                    excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
                    overhead-accounting (shaping-mode) <bytes (byte-value>;
                    priority priority-level;
                    shaping-rate (rate | predefined-variable);
                    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
                }
            }
            traffic-control-profiles profile-name {
                adjust-minimum rate;
                delay-buffer-rate (percent percentage | rate);
                excess-rate (percent percentage | proportion value | percent $junos-cos-excess-
rate);
                excess-rate-high (percent percentage | proportion value);
                excess-rate-low (percent percentage | proportion value);
                guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
                max-burst-size cells;
                overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
                peak-rate rate;
                scheduler-map map-name;
                shaping-rate (percent percentage | rate | predefined-variable) <burst-size
bytes>;
```

```
            shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
            shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
            shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
            shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
            shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
            shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
            shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
            shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
            shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
            sustained-rate rate;
        }
    }
    firewall {
        family family {
            fast-update-filter filter-name {
                interface-specific;
                match-order [match-order];
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                    only-at-create;
                }
            }
            filter filter-name {
                enhanced-mode-override;
                instance-shared;
                interface-shared;
            interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                only-at-create;
            filter filter-name {
```

```
        interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
    hierarchical-policer uid {
        aggregate {
            if-exceeding {
                bandwidth-limit-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
        premium {
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
    }
    policer uid {
        filter-specific;
        if-exceeding {
            (bandwidth-limit bps | bandwidth-percent percentage);
            burst-size-limit bytes;
        }
        logical-bandwidth-policer;
        logical-interface-policer;
        physical-interface-policer;
        then {
            policer-action;
        }
    }
```

```
        three-color-policer uid {
            action {
                loss-priority high then discard;
            }
            logical-interface-policer;
            single-rate {
                (color-aware | color-blind);
                committed-burst-size bytes;
                committed-information-rate bps;
                excess-burst-size bytes;
            }
            two-rate {
                (color-aware | color-blind);
                committed-burst-size bytes;
                committed-information-rate bps;
                peak-burst-size bytes;
                peak-information-rate bps;
                }
            }
        }
    }
    interfaces interface-name {
        interface-set interface-set-name {
            interface interface-name {
                unit logical unit number {
                    advisory-options {
                        downstream-rate rate;
                        upstream-rate rate;
                    }
                }
            }
        }
        unit logical-unit-number {
            actual-transit-statistics;
            auto-configure {
                agent-circuit-identifier {
                    dynamic-profile profile-name;
                }
                line-identity {
                    include {
                        accept-no-ids;
                        circuit-id;
                        remote-id;
                    }
                    dynamic-profile profile-name;
                }
            }
            encapsulation ppp-over-ether;
            family family {
                address address;
                filter {
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
                rpf-check {
                    fail-filter filter-name;
                    mode loose;
                }
                service {
                     input {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    input-vlan-map {
                        inner-tag-protocol-id tpid;
                        inner-vlan-id number;
                        (push | swap);
                        tag-protocol-id tpid;
                        vlan-id number;
                    }
                    output {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                    output-vlan-map {
                        inner-tag-protocol-id tpid;
                        inner-vlan-id number;
                        (pop | swap);
                        tag-protocol-id tpid;
                        vlan-id number;
                    }
                    pcef  pcef-profile-name {
                        activate rule-name | activate-all;
                    }
                }
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name (
                    shared-name filter-shared-name;
                }
                output filter-name {
                    shared-name filter-shared-name;
                }
            }
            host-prefix-only;
            ppp-options {
                aaa-options aaa-options-name;
                authentication [ authentication-protocols ];
                chap {
                    challenge-length minimum minimum-length maximum maximum-length;
                    local-name name;
                }
                ignore-magic-number-mismatch;
                initiate-ncp (dual-stack-passive | ipv6 | ip)
                ipcp-suggest-dns-option;
                mru size;
                mtu (size | use-lower-layer);
                on-demand-ip-address;
                pap;
                peer-ip-address-optional;
                local-authentication {
                    password password;
                    username-include {
                        circuit-id;
                        delimiter character;
                        domain-name name;
                        mac-address;
                        remote-id;
                    }
                }
            }
            reassemble-packets;
            targeted-options {
                backup backup;
                group group;
                primary primary;
                weight ($junos-interface-target-weight | weight-value);
            }
        }
    }
    interfaces {
        demux0 {...}
    }
    interfaces {
        pp0 {...}
    }
    policy-options {
        prefix-list uid {
            ip-addresses;
            dynamic-db;
        }
    }
    predefined-variable-defaults predefined-variable <variable-option> default-value;
    profile-type remote-device-service;
    protocols {
        router-advertisement {
            interface interface-name {
                 current-hop-limit number;
                default-lifetime seconds;
                dns-server-address
                (managed-configuration | no-managed-configuration);
                max-advertisement-interval seconds;
                min-advertisement-interval seconds;
                (other-stateful-configuration | no-other-stateful-configuration);
                prefixprefix {
                    (autonomous | no-autonomous);
                    (on-link | no-on-link);
                    preferred-lifetime seconds;
                    valid-lifetime seconds;
                }
                reachable-time milliseconds;
                retransmit-timer milliseconds;
            }
        }
    }
    routing-instances routing-instance-name {
        interface interface-name;
        routing-options {
            access {
                route prefix {
                    next-hop next-hop;
                    metric route-cost;
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
             }
         }
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
      }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
    }
    services {
    variables {
        variable-name {
```

```
                default-value default-value;
                equals expression;
                mandatory;
                uid;
                uid-reference;
            }
        }
        version-alias profile-alias-string;
    }
}
```

Hierarchy Level

```
[edit]
```

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name Name of the dynamic profile; string of up to 80 alphanumeric characters.

Required Privilege Level

routing—To view this statement in the configuration.

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

dynamic-profile (Demux)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
    network ip-address {
        range name {
            low lower-limit;
            high upper-limit;
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit unit-number demux inet auto-configure address-ranges]
[edit interfaces interface-name unit unit-number demux inet6 auto-configure address-ranges]
```

Description

Assign a dynamic profile and specify address options for the demultiplexing (demux) interface options.

Options

profile-name Name of the dynamic profile for the demultiplexing (demux) interface options.

network *ip-address* Configure an IPv4 or IPv6 address for a dynamic profile for the demultiplexing (demux) interface options.

range <range-name> Configure an IP name range used within an address-assignment pool for the demultiplexing (demux) interface options.

- `low` *lower-limit*—Lower limit of IPv4 or IPv6 address range.
- `high` *upper-limit*—Upper limit of IPv4 or IPv6 address range.

Required Privilege Level

interface—To view this statement in the configuration.

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

dynamic-profile (DHCP Local Server)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
    use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name interface interface-name],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server group group-name interface interface-name],
[edit logical-systems logical-system-name system services dhcp-local-server ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services dhcp-local-server ...],
[edit routing-instances routing-instance-name system services dhcp-local-server ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.

Options

profile-name—Name of the dynamic profile.

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

system—To view this statement in the configuration.

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

dynamic-profile (DHCP Relay Agent)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
use-primary primary-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Description

Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.

Options

profile-name—Name of the dynamic profile.

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.

dynamic-profile (Domain Map)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit access domain map domain-map-name]
```

Description

Dynamic profile that is used for subscriber sessions associated with the domain map.

Options

profile-name—Name of dynamic profile.

Required Privilege Level

admin—To view this statement in the configuration.

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

dynamic-profile (Dynamic PPPoE)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
pppoe-underlying-options]
```

Description

Attach a PPPoE dynamic profile to an underlying Ethernet interface. This underlying interface is configured with either the `encapsulation ppp-over-ether` statement or the `family pppoe` statement; the two statements are mutually exclusive. When the router creates a dynamic PPPoE logical interface on the underlying interface, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

Options

profile-name—Name of a previously configured PPPoE dynamic profile, up to 64 characters in length, defined at the `[edit dynamic-profiles` *`profile-name`* `interfaces pp0]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

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.

dynamic-profile (Static Subscribers)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
    aggregate-clients  (merge | replace);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers],
[edit logical-systems logical-system-name routing-instances routing-instances-name system
services static-subscribers group group-name],
[edit logical-systems logical-system-name system services static-subscribers],
[edit logical-systems logical-system-name system services static-subscribers group group-name],
[edit routing-instances routing-instances-name system services static-subscribers],
[edit routing-instances routing-instances-name system services static-subscribers group group-
name],
[edit system services static-subscribers],
[edit system services static-subscribers group group-name]
```

Description

Specify the dynamic client profile that is instantiated at login and de-instantiated at logout for all static subscribers on interfaces configured at the `[edit system services static-subscribers interface]` hierarchy level or for the static subscribers in a specific group. The group version of the statement takes precedence over the global version.

NOTE: Do not specify a dynamic profile that creates a dynamic interface.

Default

By default, the *junos-default-profile* is used when you do not specify a global dynamic profile with this statement.

Options

`profile-name`—Name of the dynamic client profile profile.

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

access—To view this statement in the configuration.

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

dynamic-profile (VLAN)

IN THIS SECTION

Syntax

```
dynamic-profile profile-name {
    accept  (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
    accept-out-of-band protocol;
    access-profilevlan-dynamic-profile-name;
    ranges (any | low-tag)-(any | high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges]
```

Description

Configure a dynamic profile for use when configuring dynamic VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic 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.

dynamic-profile-options

IN THIS SECTION

Syntax

```
dynamic-profile-options {
    versioning;
}
```

Hierarchy Level

```
[edit system]
```

Description

Configure global dynamic profile options.

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

dhcp-local-server

IN THIS SECTION

Syntax

```
dhcp-local-server {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name ;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dhcpv6 {
        access-profile profile-name;


                authentication {
            ...
        }
        duplicate-clients incoming-interface;
        group group-name {
            access-profile profile-name;
            authentication {
                ...
            }
            interface interface-name {
                access-profile profile-name;
                exclude;
                overrides {
                    delay-advertise {
                        based-on (option-15 | option-16 | option-18 | option-37) {
                            equals {
                                ascii ascii-string;
                                hexadecimal hexadecimal-string;
                            }
                            not-equals {
                                ascii ascii-string;
                                hexadecimal hexadecimal-string;
                            }
                            starts-with {
                                ascii ascii-string;
                                hexadecimal hexadecimal-string;
                            }
                        }
                        delay-time seconds;
                    }
                    dual-stack dual-stack-group-name;
                    interface-client-limit number;
                    multi-address-embedded-option-response;
                    process-inform {
                        pool pool-name;
                    }
                    protocol-attributes attribute-set-name;
                    rapid-commit;
                }
                service-profile dynamic-profile-name;
                short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
                trace;
                upto upto-interface-name;
            }
            overrides {
                delay-advertise {
                    based-on (option-15 | option-16 | option-18 | option-37) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                    }
                    delay-time seconds;
                }
                delegated-pool;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
                protocol-attributes attribute-set-name;
                rapid-commit;
            }
            route-suppression;
            server-duid-type type;
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        }
        overrides {
            delay-advertise {
                based-on (option-15 | option-16 | option-18 | option-37) {
                    equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    not-equals {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                    starts-with {
                        ascii ascii-string;
                        hexadecimal hexadecimal-string;
                    }
                }
                delay-time seconds;
            }
            delegated-pool;
            dual-stack dual-stack-group-name;
            include-option-82 {
                forcerenew;
                nak;
            }
            interface-client-limit number;
            multi-address-embedded-option-response;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
            rapid-commit;
        }
        reconfigure {
            attempts attempt-count;
            clear-on-terminate;
            strict;
            support-option-pd-exclude;
            timeout timeout-value;
            token token-value;
             trigger {
                radius-disconnect;
            }
        }
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        requested-ip-network-match subnet-mask;
        route-suppression;
        server-duid-type type;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group name {
        access-profile access-profile;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name ;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }

        on-demand-address-allocation;
        protocol-primary (inet | inet6);
        reauthenticate (<lease-renewal> <remote-id-mismatch >);
        service-profile service-profile;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    duplicate-clients-in-subnet  (incoming-interface | option-82);
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-
profile-name>;
    forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
    group group-name {
        authentication {
            ...
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-
profile-name>;
        interface interface-name {
            exclude;
            overrides {
                client-discover-match (option60-and-option82 | incoming-interface);
                delay-offer {
                    based-on (option-60 | option-77 | option-82) {
                        equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        not-equals {
                            ascii ascii-string;
                            hexadecimal hexadecimal-string;
                        }
                        starts-with {
```

```
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        include-option-82 {
            forcerenew;
            nak;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    trace;
    upto upto-interface-name;
}

overrides {
    client-discover-match  (option60-and-option82 | incoming-interface);
    delay-offer {
        based-on (option-60 | option-77 | option-82) {
            equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            not-equals {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
            starts-with {
                ascii ascii-string;
                hexadecimal hexadecimal-string;
            }
        }
        delay-time seconds;
    }
    include-option-82 {
```

```
                forcerenew;
                nak;
            }
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            protocol-attributes attribute-set-name;
        }
        requested-ip-network-match subnet-mask
        route-suppression;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }

    on-demand-address-allocation;
    overrides {
        client-discover-match  <option60-and-option82 | incoming-interface>;
        delay-offer {
            based-on (option-60 | option-77 | option-82) {
                equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                not-equals {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
                starts-with {
                    ascii ascii-string;
                    hexadecimal hexadecimal-string;
                }
            }
            delay-time seconds;
        }
        dual-stack dual-stack-group-name;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        protocol-attributes attribute-set-name;
    }
    pool-match-order {
        external-authority;
        ip-address-first;
        option-82;
    }
    protocol-primary;
    reauthenticate (<lease-renewal> <remote-id-mismatch >);
    reconfigure {
        attempts attempt-count;
        clear-on-terminate;
        strict;
        timeout timeout-value;
        token token-value;
         trigger {
            radius-disconnect;
        }
    }
    requested-ip-network-match subnet-mask;
    route-suppression;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system
services],
[edit logical-systems logical-system-name system services],
[edit routing-instances routing-instance-name system services],
[edit system services]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the BNG CUPS Control Plane to enable the BNG CUPS Control Plane to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The `dhcpv6` stanza configures the BNG CUPS Control Plane to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.

NOTE: When you configure the `dhcp-local-server` statement at the routing instance hierarchy level, you must use a routing instance type of `virtual-router`.

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

system—To view this statement in the configuration.

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

dhcp-relay

IN THIS SECTION

Syntax

```
dhcp-relay {
    access-profile profile-name;
        active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }

    dhcpv6 {
        access-profile profile-name;
                active-server-group server-group-name;
        }
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name interface-name;
                logical-system-name;
                mac-address mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }

        duplicate-clients incoming-interface;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        forward-only {
            logical-system <current | default | logical-system-name>;
            routing-instance <current | default | routing-instance-name>;
        }
        forward-only-replies;
        }
        forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-
interfaces);
        group group-name {
            access-profile profile-name;
            active-server-group server-group-name;
            authentication {
                password password-string;
                username-include {
                    circuit-type;
                    client-id;
                    delimiter delimiter-character;
                    domain-name domain-name-string;
                    interface-description  (device-interface | logical-interface);
                    interface-name interface-name;
                    logical-system-name;
                    mac-address mac-address;
                    relay-agent-interface-id;
                    relay-agent-remote-id;
                    relay-agent-subscriber-id;
                    routing-instance-name;
                    user-prefix user-prefix-string;
                    vlan-tags;
                }
            }
            dynamic-profile profile-name {
                aggregate-clients (merge | replace);
                use-primary primary-profile-name;
            }
            forward-only {
                logical-system <current | default | logical-system-name>;
                routing-instance <current | default | routing-instance-name>;
            }
            interface interface-name {
                access-profile profile-name;
                dynamic-profile profile-name {
                    aggregate-clients (merge | replace);
                    use-primary primary-profile-name;
                }
                exclude;
                overrides {
                    allow-snooped-clients;
                    asymmetric-lease-time seconds;
                    asymmetric-prefix-lease-time seconds;
                    client-negotiation-match incoming-interface;
                    delay-authentication;
                    delete-binding-on-renegotiation;
                    dual-stack dual-stack-group-name;
                    interface-client-limit number;
                    no-allow-snooped-clients;
                    no-bind-on-request;
                    relay-source interface-name;
                    send-release-on-delete;
                }
                service-profile dynamic-profile-name;
                short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
                trace;
                upto upto-interface-name;
            }
            }
            lease-time-validation {
                lease-time-threshold seconds;
                violation-action action;
            }

            overrides {
                allow-snooped-clients;
                asymmetric-lease-time seconds;
                asymmetric-prefix-lease-time seconds;
                client-negotiation-match incoming-interface;
```

```
                delay-authentication;
                delete-binding-on-renegotiation;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                no-allow-snooped-clients;
                no-bind-on-request;
                relay-source interface-name;
                send-release-on-delete;
            }
            relay-agent-interface-id {
                include-irb-and-l2;
                keep-incoming-interface-id ;
                no-vlan-interface-name;
                prefix prefix;
                use-interface-description (logical | device);
                use-option-82 <strict>;
                use-vlan-id;
            }
            relay-agent-remote-id {
                include-irb-and-l2;
                keep-incoming-interface-id ;
                no-vlan-interface-name;
                prefix prefix;
                use-interface-description (logical | device);
                use-option-82 <strict>;
                use-vlan-id;
            }
            relay-option {
                option-number option-number;
                default-action {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
                equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
                starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                    drop;
                    forward-only;
                    relay-server-group relay-server-group;
                }
            }
            remote-id-mismatch disconnect;
            route-suppression;
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
        }

        lease-time-validation {
            lease-time-threshold seconds;
            violation-action action;
        }

        no-snoop;
        overrides {
            allow-snooped-clients;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-negotiation-match incoming-interface;
            delay-authentication;
            delete-binding-on-renegotiation;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            no-allow-snooped-clients;
            no-bind-on-request;
            relay-source interface-name;
            send-release-on-delete;
        }
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        elay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-remote-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
        }
        relay-option-vendor-specific{
            host-name;
            location;
        remote-id-mismatch disconnect;
        route-suppression;
        server-group {
            server-group-name {
                server-ip-address;
            }
        }
        server-response-time seconds;
        service-profile dynamic-profile-name;
        short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    }
    dual-stack-group dual-stack-group-name {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-description  (device-interface | logical-interface);
                interface-name;
                logical-system-name;
                mac-address;
                relay-agent-interface-id;
                relay-agent-remote-id;
                routing-instance-name;
                user-prefix user-prefix-string;
                vlan-tags;
            }
        }
        classification-key {
            circuit-id circuit-id;
            mac-address mac-address;
            remote-id remote-id;
        }
        dual-stack-interface-client-limit number;
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }


        protocol-primary (inet | inet6);
        relay-agent-interface-id {
            include-irb-and-l2;
            keep-incoming-interface-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        relay-agent-remote-id {
            include-irb-and-l2;
            keep-incoming-remote-id ;
            no-vlan-interface-name;
            prefix prefix;
            use-interface-description (logical | device);
            use-option-82 <strict>;
            use-vlan-id;
        }
        service-profile dynamic-profile-name;
```

```
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);
group group-name {
    access-profile profile-name;
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-description  (device-interface | logical-interface);
            interface-name interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 [circuit-id] [remote-id];
            routing-instance-name;
            user-prefix user-prefix-string;
            }
        vlan-tags;
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
        routing-instance <current | default | routing-instance-name>;
    }
    forward-only {
        logical-system <current | default | logical-system-name>;
```

```
            routing-instance <current | default | routing-instance-name>;
        }
        interface interface-name {
            access-profile profile-name;
            exclude;

            overrides {
                allow-no-end-option;
                allow-snooped-clients;
                always-write-giaddr;
                always-write-option-82;
                asymmetric-lease-time seconds;
                client-discover-match <option60-and-option82 | incoming-interface>;
                delay-authentication;
                delete-binding-on-renegotiation;
                disable-relay;
                dual-stack dual-stack-group-name;
                interface-client-limit number;
                layer2-unicast-replies;
                no-allow-snooped-clients;
                no-bind-on-request;
                proxy-mode;
                relay-source
                replace-ip-source-with;
                send-release-on-delete;
                trust-option-82;
            }
            service-profile dynamic-profile-name;
            short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
            trace;
            upto upto-interface-name;
        }
        overrides {
            allow-no-end-option
            allow-snooped-clients;
            always-write-giaddr;
            always-write-option-82;
            asymmetric-lease-time seconds;
            asymmetric-prefix-lease-time seconds;
            client-discover-match (option60-and-option82 | incoming-interface);
            delay-authentication;
            delete-binding-on-renegotiation;
            disable-relay;
            dual-stack dual-stack-group-name;
            interface-client-limit number;
            layer2-unicast-replies;
            no-allow-snooped-clients;
            no-bind-on-request;
            proxy-mode;
            relay-source
            replace-ip-source-with;
            send-release-on-delete;
            trust-option-82;
        }
        relay-option {
            option-number option-number;
            default-action {
                drop;
                forward-only;
                relay-server-group group-name;
            }
            equals (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                relay-server-group relay-server-group;
            }
            starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
                drop;
                forward-only;
                local-server-group local-server-group;
                relay-server-group relay-server-group;
            }
        }
        relay-option-82 {
            circuit-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            remote-id {
                prefix prefix;
                use-interface-description (logical | device);
            }
            server-id-override
        }
        remote-id-mismatch disconnect;
        route-suppression:
```

```
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
    lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}

no-snoop;
overrides {
    allow-no-end-option
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match (option60-and-option82 | incoming-interface);
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group group-name;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
```

```
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        remote-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        server-id-override
    }
    }
    remote-id-mismatch disconnect;
    route-suppression:
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    server-response-time seconds;
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the `dhcp-relay` and `dhcpv6` statements are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

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.

dhcp-service

IN THIS SECTION

Syntax

```
dhcp-service {
        dhcp-snooping-file(local_pathname | remote_URL) {
        write-interval interval;
    }
    dhcpv6-snooping-file {
        location;
        write-interval seconds;
    }
    (disable | enable);
    interface-traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level  (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
    log {
        session {
            client;
            all;
            dhcpv6 {
                client;
                server;
                relay;
                dynamic-server;
                all;
            }
            server;
            relay;
        }
    }
    ltv-syslog-interval seconds;
    }
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        level  (all | error | info | notice | verbose | warning);
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can inprove performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.

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

admin—To view this statement in the configuration.

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

dhcp-tags (Adjustment Control Profiles)

IN THIS SECTION

Syntax

```
dhcp-tags {
    algorithm algorithm;
    priority priority;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the DHCP tags application. DHCP tags are supported for session negotiation for DHCPv4 and DHCPv6 over IP demux and VLAN demux interfaces. This means that shaping rates received from DHCP option 82 in the DISCOVER message or DHCPv6 option 17 in the SOLICIT can be used to apply a shaping rate to the logical interface.

NOTE: Single-session dual-stack DHCP is not fully supported; for example, when rates vary between the individual DHCP sessions during negotiation.

Options

algorithm Rate adjustment algorithm used by the DHCP Tags application.

- Values:
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-never—Do not perform rate adjustments.
- **Default:** adjust-less

priority Priority of the DHCP tags application in the adjustment control profile.

- **Range:** 1 through 10; 1 is the highest priority.
- **Default:** 2

Required Privilege Level

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

excess-burst-size

IN THIS SECTION

Syntax

```
excess-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name single-rate],
[edit firewall three-color-policer policer-name single-rate]
```

Description

For a single-rate three-color policer, configure the excess burst size (EBS) as a number of bytes. The EBS allows for moderate periods of bursting traffic that exceeds both the committed information rate (CIR) and the committed burst size (CBS).

NOTE: When you include the `excess-burst-size` statement in the configuration, you must also include the `committed-burst-size` and `committed-information-rate` statements at the same hierarchy level.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policing uses a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red based on the `excess-burst-size` statement included in the policer configuration.

During periods of traffic that conforms to the CIR, any unused portion of the guaranteed bandwidth capacity accumulates in the first token bucket, up to the maximum number of bytes defined by the CBS. If any accumulated bandwidth capacity overflows the first bucket, the excess accumulates in a second token bucket, up to the maximum number of bytes defined by the EBS.

A nonconforming traffic flow is categorized yellow if its size conforms to bandwidth capacity accumulated in the first token bucket. Packets in a yellow flow are marked with `medium-high` packet loss priority (PLP) and then passed through the interface.

A nonconforming traffic flow is categorized red if its size exceeds the bandwidth capacity accumulated in the second token bucket. Packets in a red traffic flow are marked with `high` PLP and then either passed through the interface or optionally discarded.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation `k` (1000), `m` (1,000,000), or `g` (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

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

fail-filter (Dynamic Profiles)

IN THIS SECTION

Syntax

```
fail-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family rpf-
check],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family rpf-check]
```

Description

Specify a filter that evaluates packets that fail a unicast RPF check. The filter determines what action to take with the failed packets. If the fail filter is not configured, the failed packets are silently discarded.

Options

filter-name Name of the filter that evaluates packets that fail the RPF check.

Required Privilege Level

interface—To view this statement in the configuration.

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

family (Dynamic Firewalls)

IN THIS SECTION

Syntax

```
family family {
    fast-update-filter filter-name {
        interface-specific;
        match-order [match-order];
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
            only-at-create;
        }
    }
    filter filter-name {
        enhanced-mode-override;
        instance-shared;
        interface-shared;
        interface-specific;
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall]
```

Description

Configure fast update filters or parameterized filters for a protocol family in a dynamic client profile or a dynamic service profile.

Options

family—Protocol family:

- any—Filter packets based on protocol-independent match conditions.
- inet—Filter Internet Protocol version 4 suite packets.
- inet6—filter Internet Protocol version 6 suite packets.

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.

family (Firewall Filter)

IN THIS SECTION

Syntax

```
family family-name {
    filter filter-name {
        interface-specific;
        term term-name {
            from {
                match-conditions;
            }
            then {
                action;
                action-modifiers;
            }
        }
    }
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a firewall filter for IP version 4 or IP version 6.

Options

family-name—Version or type of addressing protocol:

- **any**—Filter packets based on protocol-independent match conditions.
- **ethernet-switching**—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **inet**—Filter IPv4 packets.
- **inet6**—Filter IPv6 packets.

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.

firewall (Dynamic Firewalls)

IN THIS SECTION

Syntax

```
firewall {
    family family {
        filter filter-name {
            enhanced-mode-override;
            instance-shared;
            interface-shared;
            interface-specific;
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
```

```
hierarchical-policer uid {
    aggregate {
        if-exceeding {
            bandwidth-limit-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            policer-action;
        }
    }
}
policer uid {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
        policer-action;
    }
}
three-color-policer uid {
    action {
        loss-priority high then discard;
    }
    logical-interface-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure firewall filters and policers in a dynamic client profile or a dynamic service profile.

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.

filter (Dynamic Profiles Filter Attachment)

IN THIS SECTION

Syntax

```
filter {
    adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
    }
    input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family]
```

Description

Apply a dynamic filter to an interface. You can configure filters for `family any` ,`family inet`, or `family inet6`. The filters can be classic filters, fast update filters, or (for the `adf` statement) Ascend-Data-Filters.

Options

`input` *`filter-name`*—Name of one filter to evaluate when packets are received on the interface.

`output` *`filter-name`*—Name of one filter to evaluate when packets are transmitted on the interface.

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.

filter (Dynamic Profiles Filter Creation)

IN THIS SECTION

Syntax

```
filter filter-name {
    enhanced-mode-override;
    instance-shared;
    interface-shared;
    interface-specific;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family]
```

Description

Create firewall filters to be applied by dynamic profile.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). The name can also be a predefined variable.

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.

filter (Firewall Filters)

IN THIS SECTION

Syntax

```
filter filter-name {
    interface-specific;
    term term-name {
        from {
            match-conditions;
        }
        then {
            action;
            action-modifiers;
        }
    }
}
```

Hierarchy Level

```
[edit firewall family family-name]
```

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

filter-specific

IN THIS SECTION

Syntax

```
filter-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall family inet prefix-action name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name],
[edit logical-systems logical-system-name firewall family inet prefix-action name]
```

Description

By default, a policer operates in *term-specific* mode, which means that for a given *firewall filter*the Junos OS creates a separate policer instance for every filter term that references the policer. You can, however, use a common policer instance for all terms within the same firewall filter by setting the *filter-specific* option in the policer. In addition, for IPv4 firewall filters with multiple terms that reference the same policer, filter-specific mode counts and monitors the activity of the policer at the firewall filter level.

NOTE: Both filter-specific and term-specific apply to prefix-specific policer sets.

Required Privilege Level

interface—To view this statement in the configuration.

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

fpc (Subscriber Limits)

IN THIS SECTION

Syntax

```
fpc slot-number {
    limit limit;
    pic (Subscriber Limits) number {
        limit limit;
        port (Subscriber Limits) number {
            limit limit;
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type (Subscriber Limits) name]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the MPC in the specified slot. When that number is reached, subsequent logins on the card are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MIC, and per chassis.

Options

limit Maximum number of subscribers.

- **Range:** 1 through 256,000

slot-number Number of the MPC slot in the chassis.

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

system—To view this statement in the configuration.

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

from

IN THIS SECTION

Syntax

```
from {
                                        match-conditions;
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name]
```

Description

Match packet fields to values specified in a match condition. If the `from` statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

Options

match-conditions —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the `then` statement to be taken.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

group-profile (Access)

IN THIS SECTION

Syntax

```
group-profile profile-name {
    ppp {
        cell-overhead;
        encapsulated-overhead;
        framed-pool address-pool-name;
        idle-timeout seconds;
        interface-id interface-identifier;
        keepalive seconds;
        primary-dns IP address;
        primary-wins IP address;
        secondary-dns IP address;
        secondary-dns IP address;
    }
}
```

Hierarchy Level

```
[edit access]
```

Description

Configure a group profile to define Point-to-Point Protocol (PPP) attributes. Any client referencing the configured group profile inherits all the group profile attributes.

Options

- `ppp`—Configure Point-to-Point Protocol (PPP) attributes.
- `cell-overhead`—Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping.
- `framed-pool` *`pool-name`*—Configure a framed-pool.

- `idle-timeout`—Configure the idle timeout for a user.
- `interface-id`—Configure the interface identifier.
- `keep-alive`—Configure the keepalive interval for an L2TP tunnel.
- `primary-dns`—Specify the primary-dns IP address.
- `secondary-dns`—Specify the secondary-dns IP address.
- `primary-wins`—Specify the primary-wins IP address.
- `secondary-wins`—Specify the secondary-wins IP address.

Required Privilege Level

access—To view this statement in the configuration.

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

hierarchical-policer

IN THIS SECTION

Syntax (Bandwidth-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
    aggregate {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            discard;
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            discard;
        }
    }
}
```

Syntax (Packets-Per-Second (pps)-Based)

```
hierarchical-policer hierarchical-policer-name | uid {
    aggregate {
        if-exceeding-pps {
            pps-limit pps;
            packet-burst packets;
        }
        then {
            discard;
        }
    }
    premium {
        if-exceeding-pps (Hierarchical Policer) {
            pps-limit (Hierarchical Policer) pps;
            packet-burst (Hierarchical Policer) packets;
        }
        then {
            discard;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall]
```

Description

Use a hierarchical policer to rate-limit ingress Layer 2 traffic at a physical or logical interface and apply different policing actions based on whether the packets are classified as `premium` for expedited forwarding (EF) or `aggregate` for a lower priority. The two policers defined within the hierarchical policer are `aggregate` and `premium`.

NOTE:

- The `if-exceeding-pps` statement is only supported on MX Series routers with MPCs.
- The `if-exceeding` and `if-exceeding-pps` statements are mutually exclusive and, therefore, cannot be applied at the same time.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

hierarchical-policer-name—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose the name in quotation marks (“ ”).

uid—When you configure a hierarchical policer at the `[edit dynamic-profiles` *`profile name`* `firewall]` hierarchy level, you must assign a variable UID as the policer name.

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

firewall—To view this statement in the configuration.

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

high-threshold (Resource Monitor)

IN THIS SECTION

Syntax

```
high-threshold number;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the high threshold value. The value is a percentage of resources. If resource usage of any line card exceeds the limit, no new subscribers are allowed to login and no new service attachments are allowed in the corresponding FPC.

Options

number High threshold percentage for memory resource utilization

- **Default:**70
- **Range:** 1 through 99

Required Privilege Level

system—To view this statement in the configuration.

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

if-exceeding (Hierarchical Policer)

IN THIS SECTION

Syntax

```
if-exceeding {
    bandwidth-limit bps;
    burst-size-limit bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles  profile-name firewall hierarchical-policer aggregate],
[edit dynamic-profiles  profile-name firewall hierarchical-policer premium],
[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]
```

Description

Specify bandwidth and burst limits for a premium or aggregate component of a hierarchical policer.

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

firewall—To view this statement in the configuration.

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

input (Dynamic Service Sets)

IN THIS SECTION

Syntax

```
input {
     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],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family
service]
```

Description

Define the input service sets and filters to be applied to traffic by a dynamic profile.

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

Command introduced in Juniper BNG CUPS Release 23.1R1.

interfaces (Static and Dynamic Subscribers)

IN THIS SECTION

Syntax

```
interfaces {
    interface-name {
        unit logical-unit-number {
            actual-transit-statistics;
              line-identity {
                    include {
                        accept-no-ids;
                        circuit-id;
                        remote-id;
                    }
                    dynamic-profile profile-name;
                }
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                filter {
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name {
                        precedence precedence;
                        shared-name  filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name  filter-shared-name;
                    }
                }
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    mode loose;
                }
                service {
                     input {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                         service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name (
                    precedence precedence;
                    shared-name  filter-shared-name;
                }
                output filter-name {
                    precedence precedence;
                    shared-name  filter-shared-name;
                }
            }
            host-prefix-only;
            ppp-options {
                chap;
                pap;
            }
            proxy-arp;
            service {
                pcef  pcef-profile-name {
                    activate rule-name | activate-all;
                }
            }
            targeted-options {
                backup backup;
                group group;
                primary primary;
                weight ($junos-interface-target-weight | weight-value);
            }
            vlan-id;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
        vlan-tagging;
    }
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
        }
        pppoe-underlying-options {
            max-sessions number;
        }
    }
    demux0 {
        unit logical-unit-number {
            demux-options {
                underlying-interface interface-name
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                demux-source {
                    source-prefix;
                }
                filter {
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
                mac-validate (loose | strict):
                max-sessions number;
                max-sessions-vsa-ignore;
                rpf-check {
                    fail-filter filter-name;
                    mode loose;
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
            filter {
                input filter-name;
                output filter-name;
            }
            vlan-id number;
            vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
    }
    pp0 {
        unit logical-unit-number {
            keepalives interval seconds;
            no-keepalives;
            pppoe-options {
                underlying-interface interface-name;
                server;
            }
            ppp-options {
                aaa-options aaa-options-name;
                authentication [ authentication-protocols ];
                chap {
                    challenge-length minimum minimum-length maximum maximum-length;
                    local-name name;
                }
                ignore-magic-number-mismatch;
                initiate-ncp (dual-stack-passive | ipv6 | ip)
                ipcp-suggest-dns-option;
                mru size;
                mtu (size | use-lower-layer);
                on-demand-ip-address;
                pap;
                peer-ip-address-optional;
                local-authentication {
                    password password;
                    username-include {
                        circuit-id;
                        delimiter character;
                        domain-name name;
                        mac-address;
                        remote-id;
                    }
                }
            }
            family inet {
                unnumbered-address interface-name;
                address address;
                service {
                    input {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                filter {
                    input filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                    output filter-name {
                        precedence precedence;
                        shared-name filter-shared-name;
                    }
                }
            }
        }
    }
    stacked-interface-set {
        interface-set-name interface-set-name {
            interface-set-name interface-set-name;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Define interfaces for dynamic client profiles.

Options

interface-name—The interface variable (`$junos-interface-ifd-name`). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

> **NOTE**: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

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

routing—To view this statement in the configuration.

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

interface-set (Dynamic Profiles)

IN THIS SECTION

Syntax

```
interface-set interface-set-name {
    interface interface-name {
        unit logical-unit-number;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces]
```

Description

For MX Series routers with enhanced queuing DPCs or MPC/MIC modules, configure an interface set for dynamic CoS.

Options

interface-set-name—Name of the interface set to be configured or one of the following Junos OS predefined variables:

- $junos-interface-set-name—Predefined variable that, when used, is replaced with the interface-set obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.
- $junos-phy-ifd-interface-set-name—Locally generated interface set name associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

  Another use case for this predefined variable is to conserve CoS resources in a mixed business and residential topology by collecting the residential subscribers into an interface set associated with the physical interface, so that a level 2 node is used for the interface set rather than for each residential interface. Otherwise, because the business and residential subscribers share the same interface and business subscribers require three levels of CoS, then three levels are configured for each residential subscriber. That results in an unnecessary level 2 node being consumed for each residential connection, wasting CoS resources.
- $junos-pon-id-interface-set-name—Locally generated interface set name extracted from the DHCPv4 (Option 82, suboption 2) or DHCPv6 (Option 37) agent remote ID string inserted by an optical line terminal (OLT) in a passive optical network (PON). The OLT must format the agent remote ID string with a pipe symbol (|) as the delimiter between substrings. The substring extracted for the interface set name consists of the characters following the last delimiter in the agent remote ID string.

  The extracted substring identifies individual customer circuits in the PON to be aggregated into the interface set. You determine the format and contents of the substring, and configure your OLT to insert the information. Typically, the substring may include the name and port of the OLT accessed by the CPE optical network terminal (ONT).
- $junos-svlan-interface-set-name—Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is *physical_interface_name - outer_VLAN_tag*.
- $junos-tagged-vlan-interface-set-name—Locally generated interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type. For dual-tagged (client) VLANs, the format of the generated variable is *physical_interface_name - outer_VLAN_tag - inner_VLAN_tag*. For single tagged (service) VLAN, the format of the generated variable is *physical_interface_name - VLAN_tag*.

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.

interface-shared

IN THIS SECTION

Syntax

```
interface-shared;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall family family-name filter filter-name],
```

Description

Set the interface-shared attribute for a firewall filter.

> **NOTE**: A firewall filter cannot be both interface-specific and interface-shared.

Required Privilege Level

interface—To view this statement in the configuration.

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

interface-specific (Firewall Filters)

IN THIS SECTION

Syntax

```
interface-specific;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit No Link Title family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name]
```

Description

Configure interface-specific names for firewall counters.

NOTE: A firewall filter cannot be both interface-specific and interface-shared.

Required Privilege Level

interface—To view this statement in the configuration.

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

logical-bandwidth-policer

IN THIS SECTION

Syntax

```
logical-bandwidth-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

For a policer with a bandwidth limit configured as a percentage (using the `bandwidth-percent` statement), specify that the percentage be based on the shaping rate defined on the logical interface, rather than on the media rate of the physical interface.

Required Privilege Level

interface—To view this statement in the configuration.

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

logical-interface-fpc-redundancy

IN THIS SECTION

Syntax

```
logical-interface-fpc-redundancy;
```

Hierarchy Level

```
[edit interfaces aenumber aggregated-ether-options targeted-options]
```

Description

Provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution. Backup links for a subscriber are chosen on a different EQ DPC or MPC from the primary link, based on the link with the fewest number of subscribers among the links on different modules. If all links are on a single module when this is configured, backup links are not provisioned. The command is only available on the BNG User Planes.

By default, link redundancy is provided for the aggregated Ethernet bundle.

Required Privilege Level

interface—To view this statement in the configuration.

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

logical-interface-policer

IN THIS SECTION

Syntax

```
logical-interface-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall atm-policeratm-policer-name],
[edit firewall policer policer-name],
[edit firewall policer policer-template-name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer name]
```

Description

Configure a logical interface policer.

To configure the aggregate policer, configure the firewall policer you want to use as `logical-interface-policer`. And at the `firewall family` *`family-name`* `filter` *`filter-name`* hierarchy level where you will reference the policer, make the policer an `interface-specific` firewall filter action.

The sample configuration shows the relationship.

```
firewall {
    policer Shared_Policer {
      logical-interface-policer;
        if-exceeding {
            bandwidth-limit 100m;
            burst-size-limit 500k;
        }
        then {
            discard;
        }
    }
}
```

```
family inet {
    filter filter_name{
       interface-specific;
        term term_name {
            then {
                policer Shared_Policer;
                count cinet;
            }
        }
    }
}
```

Required Privilege Level

firewall—To view this statement in the configuration.

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

loss-priority (Firewall Filter)

IN THIS SECTION

Syntax

Hierarchy Level

```
[edit dynamic-profiles name firewall family inet filter name term name from],
[edit dynamic-profiles name firewall filter name term name from],
[edit firewall family inet filter              name term name from],
[edit firewall filter name term name from],
[edit logical-systems name firewall family inet filter name term name from],
[edit logical-systems name firewall filtername term name from]
```

Description

Set the loss priority of incoming packets, which governs the likelihood of the system dropping packets in the event of congestion. For example, to ensure delivery of critical traffic, you might want to set the loss priority of non-critical flows to high or medium-high to intentionally sacrifice those packets in favor of the preferred traffic whenever there is contention of resources.

Options

high — Highest probability of being dropped at times of congestion

medium-high — Second highest probability of being dropped at times of congestion

medium-low — Third highest probability of being dropped at times of congestion

low — Lowest probability of being dropped at times of congestion

Required Privilege Level

firewall

loss-priority high then discard (Three-Color Policer)

IN THIS SECTION

Syntax

```
loss-priority high then discard;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name action],
[edit firewall three-color-policer policer-name action],
[edit logical-systems logical-system-name firewall three-color-policer policer-name action]
```

Description

For packets with high loss priority, discard the packets. The loss priority setting is implicit and is not configurable. Include this statement if you do not want the local router to forward packets that have high packet loss priority.

For single-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.

For two-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.

Required Privilege Level

firewall—To view this statement in the configuration.

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

map (Domain Map)

IN THIS SECTION

Syntax

```
map domain-map-name {
    aaa-logical-system logical-system-name {
        aaa-routing-instance routing-instance-name;
    }
    aaa-routing-instance routing-instance-name;
    access-profile profile-name;
    address-pool pool-name;
    dynamic-profile profile-name;
    strip-domain;
    strip-username (left-to-right | right-to-left);
    sub-domain name {
   (
   aaa-logical-system name {aaa-routing-instance (default | name)
   } | aaa-routing-instance (default | name));
   (
   target-logical-system name {target-routing-instance (default | name)
   } | target-routing-instance (default | name));
   access-profile access-profile;
   address-pool address-pool;
   dynamic-profile dynamic-profile;
   override-chap-password override-chap-password;
   override-password override-password;
   qualifier {
    vlan-id-list [ vlan-id-list ... ];
   }
   strip-domain;
   strip-username (left-to-right | right-to-left);
   tunnel-profile tunnel-profile;
   using-user-password;
  }
    override-password password;
    target-logical-system logical-system-name {
        target-routing-instance routing-instance-name;
    }
    target-routing-instance routing-instance-name;
    tunnel-profile profile-name;
    tunnel-switch-profile profile-name;
}
```

Hierarchy Level

```
[edit access domain]
```

Description

Specify the domain map to use to map options and parameters to subscriber sessions based on the subscriber domain.

Options

domain-map-name—Name of the domain map. The name is the same as the subscriber domain to which it will apply. For example, for the username `user1@example.com`, the domain map name is `example.com`.

- `*` —Use the asterisk wildcard character in the *domain-map-name* to specify a wildcard domain map, which enables mapping based on a partial match (for example, `xyz*northern.example.com`).The router performs the wildcard lookup when there is no exact match for the subscriber domain name. The wildcard can appear anywhere within the domain name string, and can match zero or more characters. The asterisk is the only wildcard character, and only one wildcard is supported in a domain map name. If you include multiple asterisks, the first asterisk is treated as the wildcard character and the others are treated as non-wildcard characters.

- `default`—Use a domain map name of `default` to specify the domain map that the router uses when there is no exact or wildcard match for the domain or realm name in the subscriber username.

- `none`—Use a domain map name of `none` to specify the domain map the router uses when a subscriber username does not have a domain or realm name.

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

admin—To view this statement in the configuration.

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

minimum-links (Interfaces)

IN THIS SECTION

Syntax

```
minimum-linksn number;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces configured as 802.3ad redundant Ethernet interface link aggregation groups (LAGs) in a chassis cluster only, set the required minimum number of physical child links on the primary node that must be working to prevent the interface from being down. Interfaces configured as redundant Ethernet interface LAGs typically have between 4 and 16 physical interfaces, but only half, those on the primary node, are relevant to the minimum-links setting.

If the number of operating interfaces on the primary node falls below the configured value, it will cause the interface to be down even if some of the interfaces are still working.

For an aggregated ethernet interface, you cannot configure all three configuration options, `bfd-liveness-detection`, `minimum-links`, and `sync-reset` at the same time.

Options

number—For redundant Ethernet interface link aggregation group links, specify the number of physical child links on the primary node in the redundant Ethernet interface that must be working. The default **minimum-links** value is 1. The maximum value is half of the total number of physical child interfaces bound to the redundant Ethernet interface being configured or 8, whichever is smaller.

no-load-throttle (Resource Monitor)

IN THIS SECTION

Syntax

```
no-load-throttle;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

The no-load-throttle statement disables line card load-based throttling. Load-based throttling is also disabled when you configure the no-throttle statement.

Required Privilege Level

system—To view this statement in the configuration.

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

no-throttle (Resource Monitor)

IN THIS SECTION

Syntax

```
no-throttle;
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Disable the throttling of subscriber services and sessions. When throttling is disabled, if resource usage of any line card exceeds the safe limit, new subscribers logins are not blocked.

Options

`no-throttle` Disable the throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels. The throttling capability is enabled by default.

Required Privilege Level

system—To view this statement in the configuration.

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

output (Dynamic Service Sets)

IN THIS SECTION

Syntax

```
output {
    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],
[edit dynamic-profiles profile-name interfaces pp0 unit  “$junos-interface-unit” family family
service]
```

Description

Define the output service sets and filters to be applied to traffic by a dynamic profile.

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

service-set-name—Name of the service set.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

packet-burst (Hierarchical Policer)

IN THIS SECTION

Syntax

```
packet-burst packets;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name
aggregate if-exceeding-pps],
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name premium
if-exceeding-pps],
[edit No Link Title hierarchical-policer hierarchical-policer-name aggregate if-exceeding-pps],
[edit No Link Title hierarchical-policer hierarchical-policer-name premium if-exceeding-pps]
```

Description

On MPCs hosted on MX Series routers, configure the packet burst limit for premium or aggregate traffic in a hierarchical policer. When used in combination with the if-exceeding-ppsand pps-limit statements, you can control the number of packets that will be allowed over a configured packets-per-second limit when traffic is in burst state.

Options

packets—Packet burst limit in packets. You can specify the number of packets either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 1 through 24414062

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

packet-burst (Policer)

IN THIS SECTION

Syntax

```
packet-burst packets;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding-pps],
[edit firewall policer policer-name if-exceeding-pps],
[edit logical-systems logical-system-name firewall policer policer-name if-exceeding-pps]
```

Description

For a single-rate two-color policer, configure the `packet-burst` as a number of packets. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the pps-limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the packet burst limit as an absolute number of packets.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allows bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

packets—Specify the number of packets either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 1 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

peak-burst-size

IN THIS SECTION

Syntax

```
peak-burst-size bytes;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a two-rate three-color policer, configure the peak burst size (PBS) as a number of bytes. The PBS defines the maximum number of bytes of unused peak bandwidth capacity that can be accumulated. The accumulated bandwidth allows for moderate periods of bursting traffic that exceeds the peak information rate (PIR) and the committed burst size (CBS).

NOTE: When you include the peak-burst-size statement in the configuration, you must also include the committed-burst-size and peak-information-rate statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the committed information rate (CIR) and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with medium-high packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with high PLP and then either passed through the interface or optionally discarded.

Options

bytes—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

- **Range:** 1500 through 100,000,000,000 bytes

Required Privilege Level

firewall—To view this statement in the configuration.

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

peak-information-rate

IN THIS SECTION

Syntax

```
peak-information-rate bps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name two-rate],
[edit firewall three-color-policer policer-name two-rate]
```

Description

For a two-rate three-color policer, configure the peak information rate (PIR) as a number of bits per second. The PIR is the maximum rate for traffic arriving at or departing from the interface under peak line conditions. Traffic that exceeds the committed information rate (CIR) and the committed burst size (CBS) is metered to the PIR.

NOTE: When you include the peak-information-rate statement in the configuration, you must also include the committed-information-rate and peak-burst-size statements at the same hierarchy level.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits.

- A traffic flow is categorized green if it conforms to both the CIR and the CBS-bounded accumulation of available committed bandwidth capacity.
- A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with medium-high packet loss priority (PLP) and then passed through the interface.
- A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with high PLP and then either passed through the interface or optionally discarded.

Options

bps—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Range:

- • 1500 through 18,446,744,073,709,551,615 bps on MX Series routers

Required Privilege Level

firewall—To view this statement in the configuration.

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

physical-interface-policer

IN THIS SECTION

Syntax

```
physical-interface-policer;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name],
[edit firewall policer policer-name],
[edit firewall three-color-policer policer-name],
[edit logical-system logical-system-name firewall policer policer-name],
[edit logical-system logical-system-name three-color-policer policer-name],
[edit routing-instances routing-instance-name firewall policer policer-name],
[edit routing-instances routing-instance-name firewall three-color-policer policer-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name firewall
policer policer-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name firewall three-
color-policer policer-name]
```

Description

Configure an aggregate policer for a physical interface.

A physical interface policer can be a two-color or three-color policer. When you apply physical interface policer, to different protocol families on the same logical interface, the protocol families share the same policer instance. This means that rate limiting is performed in aggregate for the protocol families for which the policer is applied. This feature enables you to use a single policer instance to perform aggregate policing for different protocol families on the same physical interface. If you want a policer instance to be associated with a protocol family, the corresponding physical interface filter needs to be applied to that protocol family. The policer is not automatically applied to all protocol families configured on the physical interface.

In contrast, with logical interface policers there are multiple separate policer instances.

Required Privilege Level

interface—To view this statement in the configuration.

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

pic (Subscriber Limits)

IN THIS SECTION

Syntax

```
pic number {
    limit limit;
    port (Subscriber Limits) number {
        limit limit;
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name fpc (Subscriber Limits)
slot-number]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the specified MIC. When that number is reached, subsequent logins on the MIC are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per port, per MPC, and per chassis.

Options

number MIC number.

- **Range:** 0 through 3

limit Maximum number of subscribers.

- **Range:** `1` through `256,000`

The remaining statement is explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

policer (Configuring)

IN THIS SECTION

Syntax

```
policer policer-name {
    filter-specific;
    counter {
        counter-id counter-index;}
    if-exceeding {
        bandwidth-limit bps;
        bandwidth-percent number;
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    shared-bandwidth-policer;
    then {
        policer-action;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure policer rate limits and actions. When included at the `[edit firewall]` hierarchy level, the `policer` statement creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the `policer-action` modifier in the `then` statement in a firewall filter term or on an interface.

You can configure the policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

policer-action

One or more actions to take:

- `discard`—Discard traffic that exceeds the rate limits.
- `forwarding-class` *`class-name`*—Specify the particular forwarding class.
- `loss-priority`—Set the packet loss priority (PLP) to `low`, `medium-low`, `medium-high`, or `high`.

policer-name

Name that identifies the policer. 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 (“ ”). Policer names cannot begin with an underscore in the form `__.*`.

then

Actions to take on matching packets.

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

firewall—To view this statement in the configuration.

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

policy-options (Dynamic Profiles)

IN THIS SECTION

Syntax

```
policy-options {
    prefix-list uid {
        ip-addresses;
        dynamic-db;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Define a list of IPv4 or IPv6 address prefixes for use in a dynamic firewall filter or in an HTTP redirect configuration.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

You can configure policy options in a dynamic client profile or a dynamic service profile.

Options

uid Unique identifier of the prefix list. You must assign a UID as the prefix list name.

ip-addresses List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.

dynamic-db Specify that the routing policy and policy objects reference policies configured in the dynamic database at the `[edit dynamic]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

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

port (Subscriber Limits)

IN THIS SECTION

Syntax

```
port number {
    limit limit;
}
```

Hierarchy Level

```
[edit system services resource-monitor subscribers-limit client-type name fpc slot-number pic
(Subscriber Limits) number]
```

Description

Configure the maximum number of subscribers of a client type allowed to be logged in on the specified port. When that number is reached, subsequent logins on the port are denied until the current number of subscribers drops below the maximum allowed. You can also specify the maximum number of subscribers of a client type allowed per MIC, per MPC, and per chassis.

Options

number Port number.

limit Maximum number of subscribers.

- **Range:** `1` through `256,000`

Required Privilege Level

system—To view this statement in the configuration.

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

Release Information

ppp (Group Profile)

IN THIS SECTION

Syntax

```
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    ppp-options {
        aaa-options aaa-options-name;
        chap;
        ignore-magic-number-mismatch;
        initiate-ncp (ip | ipv6 | dual-stack-passive)
        ipcp-suggest-dns-option;
        mru;
        mtu;
        pap;
        peer-ip-address-optional;
    }
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

Hierarchy Level

```
[edit access group-profile profile-name]
```

Description

Configure PPP properties for a group profile.

Options

cell-overhead

Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

encapsulation-overhead

Configure the encapsulation overhead for class-of-service calculations.

- **Values:** *bytes*—The number of bytes used as encapsulation overhead for the session.

framed-pool

Configure the address pool.

- **Values:** *framed-pool*—References a configured address pool.

idle-timeout

Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.
- **Values:** *seconds*—Number of seconds a user can remain idle before the session is terminated.
- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0

interface-id ***interface-id***

Configure the interface identifier.

- **Values:** *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the `[edit interfaces` *`interface-name`* `unit` *`local-unit-number`* `dial-options]` hierarchy level. For more information about the interface ID, see Services Interface Naming Overview.

keepalive

Configure the keepalive interval for an L2TP tunnel.

- **Values:** *seconds*—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.

The minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.

- **Range:** 0 through 32,767 seconds
- **Default:** 30 seconds

primary-dns Configure the primary Domain Name System (DNS) server.

- **Values:** *primary-dns*—An IPv4 address.

primary-wins Configure the primary Windows Internet name server.

- **Values:** *primary-wins*—An IPv4 address.

secondary-wins Configure the secondary Windows Internet name server.

- **Values:** *secondary-wins*—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

ppp (Profile)

IN THIS SECTION

Syntax

```
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
```

Hierarchy Level

```
[edit access profile profile-name client client-name]
```

Description

Configure PPP properties for a client profile.

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

admin—To view this statement in the configuration.

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

ppp-over-ether

IN THIS SECTION

Syntax

```
ppp-over-ether;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number encapsulation]
```

Description

This encapsulation is used for underlying interfaces of pp0 interfaces. This encapsulation is supported on Fast Ethernet interface, Gigabit Ethernet interface, and Redundant Ethernet interface. When Redundant Ethernet interface is used as underlying interface, an existing pppoe session can be continued in case of failover.

Required Privilege Level

interface—To view this statement in the configuration.

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

ppp-options

IN THIS SECTION

Syntax

```
ppp-options {
    authentication [ authentication-protocols ];
        mru size;
    mtu (size | use-lower-layer);
    chap {
        access-profile name;
        challenge-length minimum minimum-length maximum maximum-length;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    lcp-max-conf-req number
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
```

```
      ncp-max-conf-req number
      ncp-restart-timer milliseconds;
      on-demand-ip-address
      pap {
          access-profile name;
          default-pap-password password;
          local-name name;
          local-password password;
          passive;
      }
  }
```

Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

On interfaces with PPP encapsulation, configure PPP-specific interface properties.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

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.

ppp-options (Dynamic PPP)

IN THIS SECTION

Syntax

```
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    lcp-connection-update;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit “$junos-
interface-unit”].
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit”]
```

Description

Configure PPP-specific interface properties in a dynamic profile.

NOTE: PPP options can also be configured in a group profile with the `ppp-options (L2TP)` statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Options

lcp-connection-update

Enable PPP to act on a Connection-Status-Message VSA (26–218) received by authd in either a RADIUS Access-Accept message or a CoA message. PPP conveys the contents of the VSA in an LCP Connection-Update-Request message to the remote peer, such as a home gateway. This action requires the following to be true:

- At least the first address family has been successfully negotiated and the session is active.
- The router LCP is in the Opened state.

Otherwise PPP takes no action on the VSA. If you do not enable the `lcp-connection-update` option, PPP processes the notification from authd, but takes no action.

- **Default:** Disabled

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.

ppp-options (L2TP)

IN THIS SECTION

Syntax

```
ppp-options {
    aaa-options aaa-options-name;
    chap;
    ignore-magic-number-mismatch;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    mru;
    mtu;
    pap;
    peer-ip-address-optional;
}
```

Hierarchy Level

```
[edit access group-profile profile-name ppp]
```

Description

Configure PPP-specific properties in a group profile that applies to tunneled PPP subscribers at the LNS.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: PPP options can also be configured for an inline service interface within a dynamic profile with the `No Link Title` statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.

- When PPP options are configured in both the dynamic profile and the group profile, the group profile options are applied to the subscriber only when the dynamic profile PPP options do not include any of the following attributes: aaa-options, chap, ipcp-suggest-dns-option, mru, mtu, pap, and peer-ip-address-optional. When any of these attributes is present, the dynamic profile is applied to the subscriber.

  When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Required Privilege Level

interface—To view this statement in the configuration.

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

pppoe-tags (Adjustment Control Profiles)

IN THIS SECTION

Syntax

```
pppoe-tags {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the Point-to-Point Protocol over Ethernet (PPPoE) Tags application.

Options

priority Priority of the Point to Point Protocol over Ethernet IA Tags application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 2

algorithm Rate adjustment algorithm used by the Point to Point Protocol over Ethernet (PPPoE) IA Tags application.

- Values:
  - adjust-never—Do not perform rate adjustments.

  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-less

Required Privilege Level

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

pppoe-underlying-options (Static and Dynamic Subscribers)

IN THIS SECTION

Syntax

```
pppoe-underlying-options {
    access-concentrator name;
    dynamic-profile profile-name;
    direct-connect
    duplicate-protection;
    max-sessions number;
    max-sessions-vsa-ignore;
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max maximum-seconds>
<filter [aci]>;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (`ppp-over-ether`) encapsulation.

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.

pps-limit (Hierarchical Policer)

IN THIS SECTION

Syntax

```
pps-limit pps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name
aggregate if-exceeding],
[edit dynamic-profiles profile-name firewall hierarchical-policer hierarchical-policer-name premium
if-exceeding],
[edit No Link Title hierarchical-policer hierarchical-policer-name aggregate if-exceeding],
[edit No Link Title hierarchical-policer hierarchical-policer-name premium if-exceeding]
```

Description

Configure the maximum bandwidth in packets per second (pps) for premium or aggregate traffic in a hierarchical policer.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short periods and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

pps—Specify the number of packets per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 2 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.

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

pps-limit (Policer)

IN THIS SECTION

Syntax

```
pps-limit pps;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall policer policer-name if-exceeding-pps],
[edit firewall policer policer-name if-exceeding-pps],
[edit logical-systems logical-system-name firewallpolicer policer-name if-exceeding-pps]
```

Description

For a single-rate two-color policer, configure the packets-per-second (pps) limit as a number of packets per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.

Traffic at the interface that conforms to the pps limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with `low` packet loss priority (PLP) and then passed through the interface.

Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.

NOTE: This statement specifies the pps limit as an absolute number of packets per second. You cannot use the pps limit as a percentage of interface bandwidth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allow more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short periods and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options

pps—Specify the number of packets per second either as a decimal number or as a decimal number followed by the abbreviation `k` (1000), or `m` (1000000).

- **Range:** 2 through 24414062
- **Default:** None

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

prefix-list

IN THIS SECTION

Syntax

```
prefix-list name {
    ip-addresses;
    apply-path path;
}
```

Hierarchy Level

```
[edit dynamic policy-options],
[edit logical-systems logical-system-name policy-options],
[edit policy-options]
```

Description

Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement, or a list of IPv6 addresses or address prefixes for use in an IPv6 RA guard policy.

You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.

Options

name—Name that identifies the list of IPv4 or IPv6 addresses or address prefixes.

ip-addresses—These are the IPv4 or IPv6 prefixes specified as prefix/prefix-length. If you omit prefix-length for an IPv4 prefix, the default is /32prefix-length. If you omit prefix-length for an IPv6 prefix, the default is /128.

The remaining statement is 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.

premium (Hierarchical Policer)

IN THIS SECTION

Syntax

```
premium {
    if-exceeding {
        bandwidth-limit bandwidth;
        burst-size-limit burst;
    }
    then {
        discard;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall hierarchical-policer],
[edit firewall hierarchical-policer]
```

Description

Specify a premium level for a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

firewall—To view this statement in the configuration.

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

pseudowire-service (Pseudowire Subscriber Interfaces)

IN THIS SECTION

Syntax

```
pseudowire-service {
    device-count number;
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure properties for the pseudowire devices on the router.

The remaining statement is 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.

radius-coa (Adjustment Control Profiles)

IN THIS SECTION

Syntax

```
radius-coa {
    priority priority;
    algorithm algorithm;
}
```

Hierarchy Level

```
[edit class-of-service adjustment-control-profiles profile-name application]
```

Description

Configure the shaping rate adjustment controls for the RADIUS CoA application.

Options

priority Priority of the RADIUS CoA application in the adjustment control profile.

- **Range:** 1 through 10; 1 being the highest priority.
- **Default:** 1

algorithm Rate adjustment algorithm used by the RADIUS CoA application.

- Values:
  - adjust-never—Do not perform rate adjustments.
  - adjust-always—Adjust the shaping rate unconditionally.
  - adjust-less—Adjust the shaping rate if it is less than the configured value.
  - adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
  - adjust-greater—Adjust the shaping rate if it is greater than the configured value.
  - adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.
- **Default:** adjust-always

Required Privilege Level

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

rebalance-periodic

IN THIS SECTION

Syntax

```
rebalance-periodic start-time hour:minute interval hours
```

Hierarchy Level

```
[edit interfaces ae number aggregated-ether-options targeted-options]
```

Description

Configure periodic rebalancing of distribution of subscribers on an aggregated Ethernet bundle.

Options

hour:minute	Time at which the rebalancing occurs, in military time.
hours	Interval at which the rebalancing occurs, in hours. Default: 24 hours.

Required Privilege Level

interface—To view this statement in the configuration.

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

rebalance-subscriber-granularity

IN THIS SECTION

Syntax

```
rebalance-subscriber-granularity <subscriber-granularity-value>;
```

Hierarchy Level

```
[edit interfaces ae number aggregated-ether-options targeted-options]
```

Description

Rebalancing takes place when the member links have more subscribers than the configured value on the rebalance-subscriber-granularity option. Changing the value to a value less than the current active value does not force IFLs to be reassigned to a different IFL.

BEST PRACTICE: We recommend that you do not configure a low granularity value. A low value can have undesirable effects, such as the router running out of pseudo logical interfaces or an increase in the convergence time for rebalancing.

Leaving subscriber granularity at the default value of 500 subscribers is sufficient in most cases. Whenever more than one member link is active, targeted distribution places new subscribers on a link with fewer subscribers than other member links on the interface.

Default

The default subscriber granularity value is 500 subscribers.

Required Privilege Level

interface—To view this statement in the configuration.

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

redundancy-group (Chassis - MX Series)

IN THIS SECTION

Syntax

```
redundancy-group {
    interface-type {
        redundant-logical-tunnel {
            device count;
        }
        redundant-virtual-tunnel {
            device count;
        }
    }
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure redundant logical tunnels, redundant virtual tunnels, or both on MX Series 5G Universal Routing Platforms.

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.

redundancy-group (Redundant Tunnel)

IN THIS SECTION

Syntax

```
redundancy-group {
    member-interface interface-name {
        (active | backup);
    minimum-links number-of-links;
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure member tunnels of redundant logical or virtual tunnels.

Options

active	Set the interface to the active mode.
backup	Set the interface to the backup mode.

minimum-links Specify the minimum number of active links required for the interface to remain up.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To view this statement in the configuration.

resource-monitor

IN THIS SECTION

Syntax

```
resource-monitor {
    free-fw-memory-watermark number;
    free-heap-memory-watermark number;
    free-nh-memory-watermark number;
    high-cos-queue-threshold number;
    high-threshold (Resource Monitor) number;
    no-logging;
    no-throttle (Resource Monitor);
    resource-category jtree {
        resource-type (contiguous-pages | free-dwords | free-pages) {
            low-watermark number;
            high-watermark number;
        }
    }
    subscribers-limit (Resource Monitor) {
        client-type (Subscriber Limits) (any | dhcp | l2tp | pppoe) {
            chassis (Subscriber Limits) {
                limit limit;
            }
            fpc (Subscriber Limits) slot-number {
                limit limit;
                pic (Subscriber Limits) number {
                    limit limit;
                    port (Subscriber Limits) number {
                        limit limit;
                    }
                }
            }
        }
    }
    traceoptions (Resource Monitor) {
        file filename <files number> <match regular-expression> <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit system services]
```

Description

Enable the resource monitoring capability to provision sufficient headroom (memory space limits that are set for the application or virtual router) for monitoring the health and operating efficiency of DPCs and MPCs. This feature also enables the memory resource monitoring mechanism to avoid the system operations from compromising on the health and traffic-handling stability of the line cards by generating error logs when a specified watermark value for memory regions and threshold value for the jtree memory region are exceeded. A trade-off on the system performance can be detrimental for supporting live traffic and protocols.

The variable *number* in the Syntax section represents a percentage.

You can only configure the resource-monitoring capability on the BNG User Planes.

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

system—To view this statement in the configuration.

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

retry (RADIUS Server)

IN THIS SECTION

Syntax

```
retry attempts;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure a limit to the number of times the MX Series router can resend a request to the RADIUS server when no response from the RADIUS server is received. If the number of retries reaches this limit, the RADIUS server is marked as dead, and the MX Series router begins to send requests to other RADIUS servers in the network element.

Options

attempts Number of attempts allowed.

- **Range:** 1 through 10
- **Default:** 3

Required Privilege Level

access—To view this statement in the configuration.

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

rewrite-rules (CoS)

IN THIS SECTION

Syntax

```
rewrite-rules {
    type rewrite-name{
        import (rewrite-name | default);
        forwarding-class class-name {
            loss-priority level code-point [ aliases ] [ 6-bit-patterns ];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

Options

- *rewrite-name*—Name of a **rewrite-rules** mapping.
- *type*—Traffic type.
- **Values: dscp**, **dscp-ipv6**, **exp**, **frame-relay-de** (J Series only), **ieee-802.1**, **ieee-802.1ad**, **inet-precedence**

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.

routing-instances

IN THIS SECTION

Syntax

```
routing-instances routing-instance-name { ... }
```

Hierarchy Level

```
[edit]
```

Description

Configure an additional routing instance on the BNG CUPS Controller for DHCP.

Default

Routing instances are disabled for the router or switch.

Options

routing-instance-name—Name of the routing instance, a maximum of 31 characters.

Required Privilege Level

routing—To view this statement in the configuration.

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

routing-instances (Dynamic Profiles)

IN THIS SECTION

Syntax

```
routing-instances routing-instance-name {
    interface interface-name;
    multicast-snooping-options {
    }
    routing-options {
        access {
            route prefix {
                metric route-cost;
                next-hop next-hop;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
          }
        }
        multicast {
            interface interface-name {
                no-qos-adjust;
            }
        }
        rib routing-table-name {
            access {
                route prefix {
                    metric route-cost;
                    next-hop next-hop;
                    preference route-distance;
                    tag route-tag;
                    tag2 route-tag2;
                }
            }
            }
          }
      }
  }
```

Hierarchy Level

```
[edit dynamic-profiles]
[edit logical-systems logical-system-name ]
```

Description

Dynamically configure an additional routing entity for a router in a dynamic client profile or a dynamic service profile.

Options

routing-instance-name—The routing instance variable (*$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

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

routing—To view this statement in the configuration.

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

secret (RADIUS Server)

IN THIS SECTION

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure a shared secret to be used by the MX Series router and the RADIUS server.

Options

password Shared secret to use.

- **Range:** 1 through 64 characters

Required Privilege Level

access—To view this statement in the configuration.

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

service (Dynamic Service Sets)

IN THIS SECTION

Syntax

```
service {
     input {
         service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    output {
         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],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family]
```

Description

Define the service sets and filters to be applied to an interface. This statement is not supported for `family inet6`.

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

Command introduced in Juniper BNG CUPS Release 23.1R1.

service-filter (Dynamic Service Sets)

IN THIS SECTION

Syntax

```
service-filter filter-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service input service-set service-set-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family service output service-set service-set-name],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family
service input service-set service-set-name],
[edit dynamic-profiles profile-name interfaces pp0 unit “$junos-interface-unit” family family
service output service-set service-set-name]
```

Description

Define the filter to be applied to traffic before it is accepted for service processing. You can use the predefined dynamic interface variables `$junos-input-service-filter`, `$junos-output-service-filter`, `$junos-input-ipv6-service-filter`, and `$junos-output-ipv6-service-filter`. Configuration of a service filter is optional; if you include the `service-set` statement without a `service-filter` definition, the router software assumes that the match condition is true and selects the service set for processing automatically.

Options

`filter-name`—Identifies the filter to be applied in service processing.

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-interface (Services Interfaces)

IN THIS SECTION

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service]
```

Description

Specify the name for the services interface associated with an interface-wide service set.

Options

`interface-name`	Identifier of the service interface.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

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: 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 Juniper BNG CUPS Release 23.1R1.

service-set-options

IN THIS SECTION

Syntax

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

Hierarchy Level

```
[edit services service-set service-set-name]
[edit system services subscriber-management mode control-plane user-plane bng-user-plane-name
service-set service-set-name]
```

Description

Specify the service set options to apply to a service set. Use the statement at the `[edit services service-set *service-set-name*]` hierarchy on BNG User Planes. Use the statement at the `[edit system services subscriber-management mode control-plane user-plane bng-user-plane-name service-set *service-set-name*]` hierarchy level on BNG CUPS Controllers.

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 Juniper BNG CUPS Release 23.1R1.

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

Define the captive portal content delivery set of the rules statements to be applied to traffic. Supports converged CPCD services.

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.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

shared-bandwidth-policer (Configuring)

IN THIS SECTION

Syntax

```
shared-bandwidth-policer;
```

Hierarchy Level

```
[edit firewall policer policer-name],
[edit firewall three-color-policer policer-name],
[edit firewall hierarchical-policer policer-name]
```

Description

Policer instances share bandwidth. This enables configuration of interface-specific policers applied on an aggregated Ethernet bundle or an aggregated SONET bundle to match the effective bandwidth and burst-size to user-configured values.

Required Privilege Level

interface—To view this statement in the configuration.

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

shared-name

IN THIS SECTION

Syntax

```
shared-name filter-shared-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter input filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family
family filter output filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter input
filter-name],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number filter
output filter-name]
```

Description

Apply a filter shared name to a dynamic filter.

Options

filter-shared-name— Name of the specific shared filter or $junos-interface-set-name.

Required Privilege Level

interface—To view this statement in the configuration.

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

shmlog (Shared Memory Log)

IN THIS SECTION

Syntax

```
shmlog {
    disable;
    file filename <files maximum-no-files> <size maximum-file-size>;
    filtering enable;
    log-name {
        all;
        logname {
            <brief | detail | extensive | none | terse>;
            <file-logging |no-file-logging>;
        }
        }
    log-type (debug | info | notice);
|
```

Hierarchy Level

```
[edit system services subscriber-management overrides]
```

Description

Junos OS uses a shared memory space to store log entries for subscriber service daemons including jpppd, jdhcpd, jl2tpd, autoconfd, bbe-smgd, authd, cosd, and dfwd. Shared memory logging is enabled by default and occurs at the client level. You can view the shmlogs on a per subscriber basis, or use filters to retrieve logs according to a variety of different parameters such as interface name, IP address, session ID, subnet, and VLAN in addition to the Client Identifier or Client DUID. Filtering is disabled by default. To see a complete list of supported filters, use this command:

```
user@rdevice> show shmlog entries logname all ?
```

When viewing logs you can limit results on the basis of event flags that include interface events, routing process interaction events, l2tp tunneling events, and ldap authentication events. To see a complete list of supported flags, use this command:

```
user@rdevice> show shmlog entries logname all flag-name ?
```

NOTE: Some platforms other than MX Series routers use shared memory logs for internal processes. These logs are not intended for customer use.

Options

disable — Name of the command to override the default behavior. Use this option to disable shared memory logging; it is always enabled otherwise.

file — Name of the file containing the shmlogs. Use this option to redirect shmlogs to a file for file-based logging. Specify the file name, define the number of files (from 2 to 1000), and set the maximum file size (from 10240 to 1073741824 bytes). Data will be written to the **/var/log/**

shmlog/ directory. Files follow this naming convention: **<cfg-file-name>-<daemon>-<severity>.log**. The shmlog files are not human-readable, so to access the logs you must first run the following command to generate a file in the **/var/log/<file-name>/** directory with logs from all daemons:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<file-name>* logname all
```

If you then want to view logs from a specific daemon, you need to run the following command to generate a file under the **/var/log/<file-name>/** directory with complete logs:

```
user@rdevice> show shmlog entries filename /var/log/shmlog/<filename>  logname authd*
```

filtering

Command to enable filtering. Filtering is subscriber centric and is useful for debugging and troubleshooting. It is disabled by default so you must use this option to enable it.

For example, if you want to quickly view the transmit packet logs for subscribers with interface-name pp0.100, you could use the following command to display only the relevant results:

```
user@rdevice> show shmlog entries logname jpppd* interface-name pp0.100 flag transmit-
packets
```

To debug sessions according to the interface name, use this command:

```
user@rdevice> show shmlog entries logname all interface-name pp0.100
```

To debug sessions that are logging in via VLAN 7 on physical-interface ge-0/0/0, use this command:

```
user@rdevice> show shmlog entries logname all vlan 7 physical-interface ge-0/0/0
```

log-name

Name of the file containing the log output. Use this option to override all logs or a specified log, and to set the verbosity level (brief, detail, extensive, none, or terse). For example, to configure **bbe-autoconf-info** for detailed file logging, you would use the following command:

```
user@rdevice> [edit system services subscriber-management overrides shmlog]
user@rdevice> set log-name bbe-autoconf-info detail file-logging
```

log-type   Severity level of the collected logs. Use this option to configure the severity level for captured logs (notice, info, or debug).

Required Privilege Level

system—To view this statement in the configuration.

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

schedulers (CoS)

IN THIS SECTION

Syntax

```
schedulers {
    scheduler-name {
        adjust-minimum rate;
        adjust-percent percentage;
        buffer-size (seconds | percent percentage | remainder | temporal microseconds);
        drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any |
non-tcp | tcp) drop-profile profile-name;
        excess-priority [ low | medium-low | medium-high | high | none];
        excess-rate (percent percentage | proportion value);
        priority priority-level;
        shaping-rate (percent percentage | rate);
        transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify the scheduler name and parameter values.

Options

scheduler-name—Name of the scheduler to be configured.

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.

scheduler-maps (For Most Interface Types)

IN THIS SECTION

Syntax

```
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.

Options

map-name—Name of the scheduler map.

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.

single-rate

IN THIS SECTION

Syntax

```
single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer policer-name]
```

Description

Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

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

firewall—To view this statement in the configuration.

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

smg-service (Enhanced Subscriber Management)

IN THIS SECTION

Syntax

```
smg-service {
    failover other-routing-engine;
    traceoptions {
        file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
        flag flag <disable>;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Configure system services, including tracing operations and Routing Engine failover, for the main enhanced subscriber management session management process, `smg-service`.

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

trace—To view this statement in the configuration.

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

stacked-vlan-ranges

IN THIS SECTION

Syntax

```
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;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
            vlan-tags;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | inet);
        access-profile vlan-dynamic-profile-name;
        ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
    }
    override;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure]
```

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

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

routing—To view this statement in the configuration.

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

statistics (Access Profile)

IN THIS SECTION

Syntax

```
statistics (time | volume-time);
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options

`time`—Collect uptime statistics only.

`volume-time`—Collect both volume and uptime statistics.

Required Privilege Level

admin—To view this statement in the configuration.

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

subscriber (Access Profile)

IN THIS SECTION

Syntax

```
subscriber username {
    delegated-pool delegated-pool-name;
    framed-ip-address ipv4-address;
    framed-ipv6-pool ipv6-pool-name;
    framed-pool ipv4-pool-name;
    password password;
    target-logical-system logical-system-name<(target-routing-instance (default | routing-
instance-name)>;
    target-routing-instance (default | routing-instance-name);
}
```

Hierarchy Level

```
[edit access profile profile-name]
```

Description

Enable local authentication for subscribers by configuring a password to match the subscriber. Local authentication can take the form of either user password authentication or Challenge Handshake Authentication Protocol' (CHAP) authentication. For user password authentication, the configured password is used to verify the subscriber's login password. For CHAP authentication, the configured password acts as the challenge secret to verify the subscriber's challenge password and challenge response credential.

NOTE: Local authentication and authorization also requires the `password` option to be configured as an `authentication-order` method for the access profile.

You can also optionally configure several attributes, such as an address, address pool, logical system, or routing instance, to be authorized locally for the subscriber when authentication is successful.

Local authentication supports all subscriber types that are currently supported by subscriber management and services on MX Series routers.

Local authentication is useful when you do not want to use external authentication servers. The associated local authorization similarly is useful when you do not want to use external authorization servers. Another use case might be when you are migrating a network from E Series routers running JunosE software to MX Series routers running Junos OS. You may also want to configure local authentication and authorization as a backup for RADIUS authentication.

If you do not configure an address or address pool for local authorization, address assignment is based on network matching or the first address pool assigned to the routing instance.

NOTE: Local authentication and authorization supports a chassis-wide maximum of 100 subscribers. If subscribers are configured in access profiles where `authentication-order password` is not configured, local authentication does not occur, but these subscriber count against the system limit of 100 subscribers for local authentication.

Options

delegated-pool *delegated-pool-name*

(Optional) Specify the name of an address pool used to locally allocate a delegated IPv6 prefix for the subscriber. Corresponds to RADIUS standard attribute Delegated-IPv6-Prefix (123).

framed-ip-address *ipv4-address*

(Optional) Specify the IP address to be configured for the subscriber. Corresponds to RADIUS standard attribute Framed-IP-Address (8).

framed-ipv6-pool *ipv6-pool-name*

(Optional) Specify the name of an address pool used to assign a router advertisement IPv6 prefix or a DHCPv6 IA_NA/128 address for the subscriber. Corresponds to RADIUS standard attribute Framed-IPv6-Pool (100).

framed-pool *ipv4-pool-name*

(Optional) Specify the name of an address pool used to assign an IPv4 address for the subscriber. Corresponds to RADIUS standard attribute Framed-Pool (88).

password *password*

Specify the password used to authenticate the subscriber locally. Corresponds to RADIUS standard attributes User-Password (2) or CHAP-Password (3).

target-logical-system *logical-system-name*

(Optional) Specify the name of the logical system assigned to the subscriber.

target-routing-instance (default | *routing-instance-name*)

(Optional) Specify the name of the routing instance assigned to the subscriber; either the default routing instance or a nondefault routing instance.

Required Privilege Level

admin—To view this statement in the configuration.

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

subscribers-limit (Resource Monitor)

IN THIS SECTION

Syntax

```
subscribers-limit {
    client-type (Subscriber Limits) (any | dhcp | l2tp | pppoe) {
        chassis (Subscriber Limits) {
            limit limit;
        }
        fpc (Subscriber Limits) slot-number {
            limit limit;
            pic (Subscriber Limits) number {
                limit limit;
                port (Subscriber Limits) number {
                    limit limit;
                }
            }
        }
    }
}
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Configure the maximum number of subscribers of a specified client type allowed to be logged in on the chassis, per MPC, per MIC, and per port. When that number is reached, subsequent logins are denied until the current number of subscribers drops below the maximum allowed.

Limit the number of subscribers allowed to log in per chassis, MPC, MIC, or port.

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

system—To view this statement in the configuration.

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

targeted-distribution

IN THIS SECTION

Syntax

```
targeted-distribution;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
[edit dynamic-profiles profile-name interfaces interface-name interface-set interface-set-name]
[edit interfaces interface-set interface-set-name targeted-distribution
[edit interfaces interface-name unit logical-unit-number targeted-distribution
```

Description

Configure egress data for a dynamic or static logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Required Privilege Level

interface—To view this statement in the configuration.

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

targeted-distribution (Dynamic Demux Interfaces over Aggregated Ethernet)

IN THIS SECTION

Syntax

```
targeted-distribution;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number]
```

Description

Configure egress data for a dynamic logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Required Privilege Level

interface—To view this statement in the configuration.

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

targeted-distribution (Static Interfaces over Aggregated Ethernet)

IN THIS SECTION

Syntax

```
targeted-distribution (replicate | scale);
```

Hierarchy Level

```
[edit interfaces demux0 unit logical-unit-number],
[edit interfaces pp0 unit logical-unit-number]
```

Description

Configure egress data for a logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Default

By default, if you do not include the targeted-scheduler statement, scheduler parameters are applied to the targeted links in the scale mode.

Options

replicate—Scheduler parameters are copied to each of the targeted aggregated interface links.

scale—Scheduler parameters are scaled based on number of targeted links and applied each of the aggregated interface targeted links.

Required Privilege Level

interface—To view this statement in the configuration.

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

targeted-options

IN THIS SECTION

Syntax

```
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number]
[edit dynamic-profiles profile-name interfaces interface-name interface-set interface-set-name]
[edit interfaces interface-set interface-set-name targeted-distribution
[edit interfaces interface-name unit logical-unit-number targeted-distribution
```

Description

Configure primary and backup links, group similar subscribers, and specify a subscriber weight for manual targeting to distribute subscribers across aggregated Ethernet member links. The command is only available on the BNG User Planes.

Options

backup

(Optional) Specify a backup member link per subscriber when you configure manual targeting.

group

(Optional) Assign a group name for subscribers with similar bandwidth usage. Subscribers that are configured for targeted distribution without a group name are added to the `default` group and distributed evenly across member links. Grouping of subscribers is supported only for static subscribers.

- **Default:** default

primary

Specify a primary member link per subscriber when you configure manual targeting. You must always configure a primary link when you configure manual targeting.

weight ($junos-interface-target-weight | *weight-value*)

Specify the weight for targeted subscribers like PPPoe, demux, and conventional VLANs based on factors such as customer preferences, class of service (CoS), or bandwidth requirement. Member links for logical interfaces of aggregated Ethernet logical interfaces are assigned based on the value of the weight. When a new VLAN is added to the same aggregated Ethernet bundle, then the primary member link selected for targeting is the one with the minimum primary load and the backup link selected for targeting is the one with the minimum overall load.

The $junos-interface-target-weight predefined variable is supported for dynamic configuration only. When you configure this predefined variable, the weight value is sourced from VSA 26-213 in the RADIUS Access-Accept message when a dynamic subscriber is authenticated.

- **Range:** 1 through 1000

Required Privilege Level

interface—To view this statement in the configuration.

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

targeted-options (PS Interface)

IN THIS SECTION

Syntax

```
targeted-options {

    logical-interface-chassis-redundancy;
    logical-interface-fpc-redundancy;
    rebalance-periodic {
        interval interval;
        start-time start-time;
    }
    rebalance-subscriber-granularity;
    single-targeted-link;
    type;
}
```

Hierarchy Level

```
[edit interfaces ps name]
```

Description

Configure the option to achieve targeted distribution only for subscriber-interface on pseudowire (PS) IFDs when the PS IFD anchor point is a redundant logical tunnel (RLT), which has multiple LT links in active-active mode (not in link-protection mode).

Required Privilege Level

interface—To view this statement in the configuration.

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

term (Captive Portal Content Delivery)

IN THIS SECTION

Syntax

```
term term-name{
    from {
        destination-address address <except>;
    }
    then {
        accept;
        redirect url;
        rewrite destination-address address <destination-port port-number>;
        syslog;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name],
```

Description

Define the term match and action properties for the captive portal content delivery rule.

Options

term-name—Identifier for the term.

The remaining statement is 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 Juniper BNG CUPS Release 23.1R1.

term (Dynamic Profiles)

IN THIS SECTION

Syntax

```
term term-name {
    from {
        match-conditions;
    }
    then {
        action;
        action-modifiers;
    }
    only-at-create;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name fast-update-filter filter-name],
[edit dynamic-profiles profile-name firewall family family-name  filter  filter-name]
```

Description

Define terms for fast update filters.

Options

action—(Optional) An action to take if conditions match. If you do not specify an action, the packets that match the conditions in the `from` statement are accepted.

action-modifiers—(Optional) One or more actions to perform on a packet.

`from`—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

match-conditions—One or more conditions to make a match.

`only-at-create`—(Optional) Specify that the term is added only when the fast update filter is first created. No subsequent changes can be made to the term in the filter. Use this option only for terms that do not include subscriber-specific data in their match conditions, such as common or default terms (for example, counting the default drop packets).

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`then`—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the `from` statement, the packet is accepted.

Required Privilege Level

interface—To view this statement in the configuration.

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

term (Firewall Filter)

IN THIS SECTION

Syntax

```
term term-name {
    from {
        match-conditions;
        vxlan {
            vni vni-id
                flags value mask-in-hex value
                reserved1 value
                reserved2 value
            }
        ip-version ipv4 {
            match-conditions-mpls-ipv4-address;
                protocol (tcp | udp) {
                    match conditions-mpls-ipv4-port;
                }
        }
    }
    then {
        actions;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall family family-name filter filter-name],
[edit firewall family family-name filter filter-name],
[edit firewall family family-name service-filter filter-name],
[edit firewall family family-name simple-filter filter-name],
[edit logical-systems logical-system-name firewall family family-name filter filter-name],
[edit logical-systems logical-system-name firewall family family-name service-filter filter-name],
[edit logical-systems logical-system-name firewall family family-name simple-filter filter-name]
```

Description

Define a firewall filter term.

Options

`actions`—(Optional) Actions to perform on the packet if conditions match. You can specify one *terminating action* supported for the specified filter type. If you do not specify a terminating action, the packets that match the conditions in the `from` statement are accepted by default. As an option, you can specify one or more *nonterminating actions* supported for the specified filter type.

`filter-name`—(Optional) For `family` *`family-name`* `filter` *`filter-name`* only, reference another standard stateless firewall filter from within this term.

`from`—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the `then` statement are taken.

`match-conditions`—One or more conditions to use to make a match on a packet.

match-conditions-mpls-ipv4-address—(MPLS-tagged IPv4 traffic only) One or more IP address match conditions to match on the IPv4 packet header. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

match-conditions-mpls-ipv4-port—(MPLS-tagged IPv4 traffic only) One or more UDP or TCP port match conditions to use to match a packet in an MPLS flow. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

vxlan—(Optional) Match packets belonging to a particular VXLAN Network Identifier (VNI).

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

`then`—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the `from` statement, the packet is accepted.

The Firewall Filer Match Conditions for the different protocols are explained separately:

- "Firewall Filter Match Conditions for IPv4 Traffic" on page 334
- "Firewall Filter Match Conditions for IPv6 Traffic" on page 348

Required Privilege Level

firewall—To view this statement in the configuration.

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

terminate-code

IN THIS SECTION

Syntax

```
terminate-code (aaa (deny | service-shutdown | shutdown) | dhcp | l2tp | ppp | vlan) term-reason
radius value;
```

Hierarchy Level

```
[edit access]
```

Description

Customize the mapping between a termination cause (the internal termination identifier) and a numerical code value for the cause that is reported in the RADIUS Acct-Terminate-Cause attribute (49).

When a RADIUS Acct-Stop message is issued as a result of the termination of a subscriber or service session, the RADIUS Acct-Terminate-Cause attribute (49) reports the cause or reason for the termination. This attribute is included only in RADIUS Acct-Stop messages. The termination cause is conveyed as a code value in the attribute. *RFC 2866, RADIUS Accounting,* defines the standard mapping between 18 code values and termination causes.

Junos OS defines a set of internal termination cause codes for AAA, DHCP, L2TP, PPP, and VLAN subscriber and service session failures. By default, these internal cause codes are mapped to the RFC-defined code values. When a subscriber or service session is terminated, the router logs a message for the internal termination cause and logs another message for the RADIUS Acct-Terminate-Cause attribute. You can use the logged information to help monitor and troubleshoot the events.

Because there are many different Junos OS internal identifiers for termination causes and only 18 supported, RFC-defined standard code values, by default a given code value can map to multiple identifiers. Instead of using the default code values, you can use the `terminate-code` statement to map any of the internally defined termination causes to any 32-bit number (1 through 4,294,967,295). The flexibility of customized mapping greatly increases the possibilities for fine-grained analytics and failure tracking.

Options

aaa — Map internal identifiers for AAA-specific termination causes to a numerical value.

deny — Limit selection of termination causes to those associated with denial of subscriber access.

dhcp — Map internal identifiers for DHCP-specific termination causes to a numerical value.

l2tp — Map internal identifiers for L2TP-specific termination causes to a numerical value.

ppp — Map internal identifiers for PPP-specific termination causes to a numerical value.

radius *value* — Number that represents the termination cause in the RADIUS Acct-Terminate-Cause attribute (49).

- **Range:** 1 through 4,294,967,295

service-shutdown — Limit selection of termination causes to those associated with established service sessions independent of the parent subscriber session.

shutdown — Limit selection of termination causes to those associated with established subscriber sessions.

vlan — Map internal identifiers for VLAN-specific termination causes to a numerical value.

term-reason — Internal identifier for the termination causes defined for the specified protocol type.

For protocol-specific termination causes, see the following topics:

- No Link Title
- No Link Title
- No Link Title
- No Link Title
- No Link Title

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 Junos OS Release 11.4.

`vlan` option added in Junos OS Release 16.1.

RELATED DOCUMENTATION

No Link Title

No Link Title

then (Captive Portal Content Delivery)

IN THIS SECTION

Syntax

```
then {
    accept;
    redirect url;
    rewrite destination-address address <destination-port port-number>;
    syslog;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name services captive-portal-content-delivery rule rule-name term
term-name],
```

Description

Define the term actions and any optional action modifiers for the captive portal content delivery rule.

Options

action Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

- `accept`—Accept the packets and all subsequent packets in flows that match the rules.
- `redirect`—Redirect the packet and all subsequent packets in flows that match the rules. You can optionally configure the following action modifier:
    - *`url`*— URL destination for the redirected packet. The URL must begin with `http://` or `https://`.

rewrite Rewrite the packet and all subsequent packets in flows that match the rules. You can optionally configure one or both of the following action modifiers:

- `destination-address` *`address`*—Destination address for the rewritten packet.
- `destination-port` *`port-number`*—(Optional) Destination port for the rewritten packet.

syslog Log information about the packet to a system log file.

action Actions to accept, redirect, or rewrite packets and all subsequent packets in flows that match the rules.

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

then (Policer Action)

IN THIS SECTION

Syntax

```
then {
                        policer-action;
}
```

Hierarchy Level

```
[edit firewall policer policer-name]
[edit logical-systems logical-system-name firewall policer policer-name]
```

Description

Configure a policer action.

Options

policer-action—Actions to take are:

- discard—Discard traffic that exceeds the rate limits defined by the policer.
- forwarding-class *class-name*—Classify traffic that exceeds the rate limits defined by the policer.
- loss-priority—Set the loss priority for traffic that exceeds the rate limits defined by the policer.

Required Privilege Level

firewall—To view this statement in the configuration.

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

three-color-policer (Applying)

IN THIS SECTION

Syntax

```
three-color-policer {
    (single-rate | two-rate) policer-name;
}
```

Hierarchy Level

```
[edit firewall family family-name filter filter-name term term-name then]
[edit logical-systems logical-system-name firewall family family-name filter filter-name term
term-name then]
```

Description

Apply a tricolor marking policer.

Options

`single-rate`—Named tricolor policer is a single-rate policer.

`two-rate`—Named tricolor policer is a two-rate policer.

`policer-name`—Name of a tricolor policer.

Required Privilege Level

firewall—To view this statement in the configuration.

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

three-color-policer (Configuring)

IN THIS SECTION

Syntax

```
three-color-policer policer-name | uid {
    action {
        loss-priority high then discard;
    }
    filter-specific;
    logical-interface-policer;
    physical-interface-policer;
    shared-bandwidth-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall],
[edit firewall],
[edit logical-systems logical-system-name firewall]
```

Description

Configure a three-color policer in static firewall filters or dynamic firewall filters in a dynamic client profile or a dynamic service profile.

Options

policer-name—Name of the three-color policer. Reference this name when you apply the policer to an interface.

uid—When you configure a policer at the `[edit dynamic-profiles]` hierarchy level, you must assign a variable UID as the policer name.

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

firewall—To view this statement in the configuration.

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

timeout (RADIUS Server)

IN THIS SECTION

Syntax

```
timeout seconds;
```

Hierarchy Level

```
[edit access radius servers name]
```

Description

Configure the amount of time that the MX Series router waits to receive a response from a RADIUS server before retrying the request.

Options

seconds	Number of seconds to wait. • **Range:** 1 through 90 • **Default:** 3

Required Privilege Level

access—To view this statement in the configuration.

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

traceoptions (Protocols PPP Service)

IN THIS SECTION

Syntax

```
traceoptions {
    file <filename> <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
        user user@domain;
    }
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols ppp-service]
```

Description

Define tracing operations for PPP service processes.

Options

`file` *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

`files` *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000

- **Default:** 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

BEST PRACTICE: Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- aci *regular-expression*—Regular expression to match the agent circuit identifier provided by PPP client.

- ari *regular-expression*—Regular expression to match the agent remote identifier provided by PPP client.

- service *regular-expression*—Regular expression to match the name of PPPoE service.

- underlying-interface *interface-name*—Name of a PPP underlying interface. You cannot use a regular expression for this filter option.

- user *user@domain*—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- accounting-statistics—Trace accounting statistics events.

- all—Trace all operations.

- authentication—Trace authentication events.

- chap—Trace CHAP events.

- events—Trace interface events.

- gres—Trace GRES events.

- init—Trace daemon initialization events.

- interface-db—Trace interface database events.

- lcp—Trace LCP state machine events.

- memory—Trace memory processing events.

- `ncp`—Trace NCP state machine events.
- `packet-error`—Trace packet error events.
- `pap`—Trace PAP events.
- `parse`—Trace parsing events.
- `profile`—Trace libdynamic profile events.
- `receive-packets`—Trace received PPP packets.
- `routing-process`—Trace routing process interactions.
- `rtp`—Trace real-time priority events.
- `rtsock`—Trace routing socket events.
- `session-db`—Trace session database interactions.
- `smi-services-sentry`—Trace SMI services requests and retries.
- `states`—Trace state machine events.
- `transmit-packets`—Trace transmitted PPP packets.
- `tunnel`—Trace L2TP tunneling events.

`level`—Level of tracing to perform. You can specify any of the following levels:

- `all`—Match all levels.
- `error`—Match error conditions.
- `info`—Match informational messages.
- `notice`—Match notice messages about conditions requiring special handling.
- `verbose`—Match verbose messages.
- `warning`—Match warning messages.
- **Default:** `error`

`match` *`regular-expression`*—(Optional) Refine the output to include lines that contain the regular expression.

`no-remote-trace`—Disable remote tracing.

`no-world-readable`—(Optional) Disable unrestricted file access.

`size` *`maximum-file-size`*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Syntax:** *`size`*`k` to specify KB, *`size`*`m` to specify MB, or *`size`*`g` to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

`world-readable`—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

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

traceoptions (Resource Monitor)

IN THIS SECTION

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-
readable | no-world-readable>;
    flag flag;
}
```

Hierarchy Level

```
[edit system services resource-monitor]
```

Description

Define tracing operations for the memory resource utilization processes.

Options

file *filename*

Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**.

- **Default: rmopd**

files *number*

(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the `size` option.

- **Range:** 2 through 1000
- **Default:** 3 files

match *regular-expression*

(Optional) Refine the output to include lines that contain the regular expression.

`size` *`maximum-file-size`*

(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

`world-readable` (Optional) Enable unrestricted file access.

`no-world-readable` (Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

flag *flag* Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- `all`—Trace all operations.

Required Privilege Level

trace—To view this statement in the configuration.

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

traffic-control-profiles

IN THIS SECTION

Syntax

```
traffic-control-profiles profile-name {
    adjust-minimum rate;
    atm-service (cbr | rtvbr | nrtvbr);
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage  | proportion value );
    excess-rate-high (percent percentage | proportion value);
    excess-rate-low (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate) <burst-size bytes>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode | frame-mode-bytes | cell-mode-bytes) <bytes (byte-value)>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    strict-priority-scheduler;
    sustained-rate rate;
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.

Options

profile-name—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.

The remaining statements are explained separately. See 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.

two-rate

IN THIS SECTION

Syntax

```
two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name firewall three-color-policer name],
[edit firewall three-color-policer policer-name],
[edit logical-systems logical-system-name firewall three-color-policer policer-name]
```

Description

Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).

Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).

Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.

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

firewall—To view this statement in the configuration.

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

update-interval

IN THIS SECTION

Syntax

```
update-interval minutes;
```

Hierarchy Level

```
[edit access profile profile-name accounting]
```

Description

Enable interim accounting updates and configure the amount of time that the router or switch waits before sending a new accounting update.

Interim accounting updates are included in the exchange of messages between the client and the accounting server. In RADIUS accounting, the client is the network access server (NAS), which can be the router or switch. The NAS sends Accounting-Request messages to the server, which acknowledges receipt of the requests with Accounting-Response messages. Interim accounting updates are sent in Accounting-Request packets with the Acct-Status-Type attribute set to Interim-Update.

When a user is authenticated, the authentication server issues an Access-Accept message in response to a successful Access-Request message. The interval between interim updates can be configured directly on the server using the Acct-Interim-Interval attribute of the Access-Accept message. However, if the update interval is configured on the NAS using `update-interval`, the system prefers the attributes returned by RADIUS and overrides the locally configured values.

> **NOTE**: All information in an interim update message is cumulative from the beginning of the session, not from the last interim update message.

Default

No interim updates are sent from the client to the accounting server.

Options

`minutes`—Amount of time between updates, in minutes. All values are rounded to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

- **Range:** 10 through 1440 minutes

Required Privilege Level

admin—To view this statement in the configuration.

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

variables (Dynamic Service Profiles)

IN THIS SECTION

Syntax

```
variables variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    uid;
    uid-reference;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name]
```

Description

Configure user-defined variables in a dynamic service profile. The values that the system uses for these variables are provided by the RADIUS server and applied when the subscriber authenticates. You can configure default values that are used when RADIUS does not return a value. Alternatively, you can specify that the profile fails if RADIUS does not return a value for a variable.

> **NOTE**: Do not use this statement in a dynamic client profile.

Options

variable-name—Name of the variable.

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

Command introduced in Juniper BNG CUPS Release 23.1R1.

6

CHAPTER

Junos OS CLI Operational Commands

test aaa dhcp user | 1061

test aaa ppp user | 1067

clear dhcp relay binding

IN THIS SECTION

Syntax

```
clear dhcp relay binding
<address>
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.

Options

address (Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCP clients.

dual-stack (Optional) Clear the binding state for DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

interface *interface-name* (Optional) Clear the binding state for DHCP clients on the specified interface.

interfaces-vlan (Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard (Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).

logical-system *logical-system-name* (Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance *routing-instance-name* (Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the `clear dhcp relay binding` command is issued.

```
user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
198.51.100.32     00:00:5e:00:53:01  active   2007-02-08 16:41:17 EST
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST



user@host> clear dhcp relay binding 198.51.100.32

user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
192.168.14.8      00:00:5e:00:53:02  active   2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding dual-stack all

The following command clears all DHCP relay agent bindings for all DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

```
user@host> clear dhcp relay binding dual-stack all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

clear dhcp relay lockout-entries

IN THIS SECTION

Syntax

```
clear dhcp relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all — Clear all client entries from the lockout database.

index *index* — Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of the output fields.

Sample Output

clear dhcp relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp relay lockout-entries index 2

user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp relay lockout-entries` command is issued for all entries.

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp relay lockout-entries all

user@host> show dhcp relay lockout-entries all
```

clear dhcp relay statistics

IN THIS SECTION

Syntax

```
clear dhcp relay statistics

<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

logical-system *logical-system-name*	(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.

routing-instance *routing-instance-name* (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

See

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the `clear dhcp relay statistics` command is issued.

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                      1
    Lease Time Violated        1

Messages received:
    BOOTREQUEST                116
    DHCPDECLINE                0
    DHCPDISCOVER               11
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                105

Messages sent:
    BOOTREPLY                  44
    DHCPOFFER                  11
    DHCPACK                    11
    DHCPNAK                    11


user@host> clear dhcp relay statistics

user@host> show dhcp relay statistics
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

clear dhcp server binding

IN THIS SECTION

Syntax

```
clear dhcp server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Options

address (Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCP clients.

interface *interface-name* (Optional) Clear the binding state for DHCP clients on the specified interface.

> **NOTE**: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *`ip-address`* or *`mac-address`* options.

interfaces-vlan

(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard

(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system *logical-system-name*

(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance *routing-instance-name*

(Optional) Clear the binding state for DHCP clients on the specified routing instance.

dual-stack

(Optional) Remove either both arms or single arm of dual-stack.

> **NOTE**:
>
> - The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
> - When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the `clear dhcp server binding` command is issued.

```
user@host> show dhcp server binding

2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address        Hardware address   Type     Lease expires at
198.51.100.1       00:00:5e:00:53:01  active   2007-01-17 11:38:47 PST
198.51.100.3       00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST


user@host> clear dhcp server binding 198.51.100.1

user@host> show dhcp server binding

1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address        Hardware address   Type     Lease expires at
198.51.100.3       00:00:5e:00:53:02  active   2007-01-17 11:38:41 PST
```

clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface `ae0`, which clears DHCP bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

clear dhcp server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcp server binding dual-stack all
```

clear dhcp server lockout-entries

IN THIS SECTION

Syntax

```
clear dhcp server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv4 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all — Clear all client entries from the lockout database.

index *index* — Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcp-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of the output fields.

Sample Output

clear dhcp server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1

user@host> clear dhcp server lockout-entries index 2

user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

clear dhcp server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcp server lockout-entries` command is issued for all entries.

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
```

```
2     00:00:5E:00:53:11     Grace      120          780            2
3     00:00:5E:00:53:22     Lockout    180          2300           1

user@host> clear dhcp server lockout-entries all

user@host> show dhcp server lockout-entries all
```

clear dhcp server statistics

IN THIS SECTION

Syntax

```
clear dhcp server statistics

<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

logical-system *logical-system-name*
(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.

routing-instance *routing-instance-name*
(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

Sample Output

clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the `clear dhcp server statistics` command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
    Total                      1
    Lease Time Violation       1

Messages received:
    BOOTREQUEST                89163
    DHCPDECLINE                0
    DHCPDISCOVER               8110
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                81053

Messages sent:
    BOOTREPLY                  32420
    DHCPOFFER                  8110
    DHCPACK                    8110
    DHCPNAK                    8100


user@host> clear dhcp server statistics

user@host> show dhcp server statistics
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPOFFER                  0
    DHCPACK                    0
    DHCPNAK                    0
```

clear dhcpv6 relay binding

IN THIS SECTION

Syntax

```
clear dhcpv6 relay binding
<address>
<all>
<dual-stack>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.

Options

address (Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all (Optional) Clear the binding state for all DHCPv6 clients.

dual-stack (Optional) Clear the binding state for DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.

interfaces-vlan (Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard (Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).

interface *interface-name* (Optional) Clear the binding state for DHCPv6 clients on the specified interface.

logical-system *logical-system-name* (Optional) Clear the binding state for DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name* (Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

See `No Link Title` for an explanation of output fields.

Sample Output

clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the `clear dhcpv6 relay binding` command is issued.

```
user@host> show dhcpv6 relay binding
```

```
Prefix                   Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64    1           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64    2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:db8:3c4d:15::/64
user@host> show dhcpv6 relay binding

Prefix                   Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:16::/64    2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

clear dhcpv6 relay binding dual-stack all

The following command clears all DHCPv6 relay agent bindings for all DHCPv6 clients and the associated DHCPv4 bindings in the single-session DHCP dual stack. DHCPv4 clients created in a DHCPv4-only stack are not affected.

```
user@host> clear dhcpv6 relay binding dual-stack all
```

clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 relay binding interface ae0
```

clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

clear dhcpv6 relay lockout-entries

IN THIS SECTION

Syntax

```
clear dhcpv6 relay lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 relay agent lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all — Clear all client entries from the lockout database.

index *index* — Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-relay lockout-entries` command.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of the output fields.

Sample Output

clear dhcpv6 relay lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30            5200            2
2     00:00:5E:00:53:11     Grace      120            780             2
3     00:00:5E:00:53:22     Lockout    180            2300            1

user@host> clear dhcpv6 relay lockout-entries index 2

user@host> show dhcpv6 relay lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30            5200            2
3     00:00:5E:00:53:22     Lockout    180            2300            1
```

clear dhcpv6 relay lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 relay lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key               State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30            5200            2
```

```
2     00:00:5E:00:53:11     Grace      120           780            2
3     00:00:5E:00:53:22     Lockout    180           2300           1

user@host> clear dhcpv6 relay lockout-entries all

user@host> show dhcpv6 relay lockout-entries all
```

Command introduced in Junos OS Release 18.2R1.

clear dhcpv6 relay statistics

IN THIS SECTION

Syntax

```
clear dhcpv6 relay statistics


<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

logical-system *logical-system-name*
(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.

routing-instance *routing-instance-name*
(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Output Fields

Sample Output

clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the `clear dhcpv6 relay statistics` command is issued.

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0
    Lease Time Violated         1
```

```
Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              10
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_REPL           0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_FORW           0
user@host> clear dhcpv6 relay statistics
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                       0

Messages received:
    DHCPV6_DECLINE              0
    DHCPV6_SOLICIT              0
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE              0
    DHCPV6_REQUEST              0
    DHCPV6_CONFIRM              0
    DHCPV6_RENEW                0
    DHCPV6_REBIND               0
    DHCPV6_RELAY_REPL           0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_FORW           0
```

clear dhcpv6 server binding

IN THIS SECTION

Syntax

```
clear dhcpv6 server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Description

Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.

Options

address

(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all

(Optional) Clear the binding state for all DHCPv6 clients.

interface *interface-name*

(Optional) Clear the binding state for DHCPv6 clients on the specified interface.

interfaces-vlan

(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard

(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system *logical-system-name*

(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name*

(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

dual-stack

(Optional) Remove either both arms or single arm of dual-stack.

NOTE:

- The `dual-stack` command is added in the syntax removes both arms of the dual-stack with a single command entry.
- When the `dual-stack` command is not added in the syntax, the `clear dhcpv6 server binding` command clears only the family specific arm of the dual-stack.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface `ae0`, which clears DHCPv6 bindings on all demux VLANs on top of `ae0`:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcpv6 server binding dual-stack all
```

clear dhcpv6 server lockout-entries

IN THIS SECTION

Syntax

```
clear dhcpv6 server lockout-entries (all | index index)
```

Description

Clear all client entries from the DHCPv6 local server lockout database or only the specified entries. The lockout is terminated for all affected client sessions. The lockout history for these clients is also cleared. The clients that were locked out are allowed to attempt to log in. Any subsequent short-cycle event results in a new lockout, with the initial lockout period at the low end of the range.

Options

all — Clear all client entries from the lockout database.

index *index* — Number identifying a client entry to be cleared from the lockout database. You can view the index numbers associated with all clients by issuing the `show dhcpv6-server lockout-entries` command.

Required Privilege Level

view

Output Fields

See No Link Title for an explanation of the output fields.

Sample Output

clear dhcpv6 server lockout-entries (Specific Lockout Entry)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for a specific entry.

```
user@host> show dhcpv6 server lockout-entries all
Index     Key                State    Expires(s)     Elapsed(s)      Count
```

```
1     00:00:5E:00:53:00     LT          30          5200            2
2     00:00:5E:00:53:11     GT         120          780             2
3     00:00:5E:00:53:22     LT         180          2300            1

user@host> clear dhcpv6 server lockout-entries index 2

user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)    Elapsed(s)      Count
1     00:00:5E:00:53:00     LT          30          5200            2
3     00:00:5E:00:53:22     LT         180          2300            1
```

clear dhcpv6 server lockout-entries (All Lockout Entries)

The following sample output displays the lockout entries in the database before and after `clear dhcpv6 server lockout-entries` command is issued for all entries.

```
user@host> show dhcpv6 server lockout-entries all
Index     Key              State    Expires(s)    Elapsed(s)      Count
1     00:00:5E:00:53:00     LT          30          5200            2
2     00:00:5E:00:53:11     GT         120          780             2
3     00:00:5E:00:53:22     LT         180          2300            1

user@host> clear dhcpv6 server lockout-entries all

user@host> show dhcpv6 server lockout-entries all
```

clear dhcpv6 server statistics

IN THIS SECTION

Syntax

```
clear dhcpv6 server statistics

<interface interface-name>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

logical-system *logical-system-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.
routing-instance *routing-instance-name*	(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

clear network-access aaa statistics

IN THIS SECTION

Syntax

```
clear network-access aaa statistics
<accounting>
<address-assignment (client | pool pool-name)>
<authentication>
<dynamic-requests>
<radius>
```

```
<re-authentication>
<session-limit-per-username username username access-profile profile-name>
<terminate-code>
```

Description

Clear AAA statistics.

Options

accounting

(Optional) Clear AAA accounting statistics.

address-assignment client

(Optional) Clear AAA address-assignment statistics for the client.

address-assignment pool *pool-name*

(Optional) Clear AAA address-assignment pool statistics.

authentication

(Optional) Clear AAA authentication statistics.

dynamic-requests

(Optional) Clear AAA dynamic-request statistics.

radius

(Optional) Clears the values in the `Peak` and `Exceeded` columns only.

re-authentication

(Optional) Clear AAA reauthentication statistics.

session-limit-per-username

(Optional) Clear all blocked request statistics for all access profiles from the username session-limit table. You can also specify additional options:

- `username` *`username`*—Clear the blocked request statistics for the specified username across all access profiles. A given username can be used in more than one access profile.
- `access-profile` *`profile-name`*—Clear the blocked request statistics for all usernames in the specified access profile.

> **NOTE**: This command does not clear (delete) the entry in the session-limit table. Entries in the table are added or deleted during session login or logout processing.

terminate-code (Optional) Clear AAA termination code statistics.

Required Privilege Level

maintenance

Output Fields

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

Sample Output

clear network-access aaa statistics accounting

```
user@host> clear network-access aaa statistics accounting
```

clear network-access aaa statistics address-assignment pool

```
user@host> clear network-access aaa statistics address-assignment pool isp_1
```

clear network-access aaa statistics radius

```
user@host> clear network-access aaa statistics radius
```

clear network-access aaa subscriber

IN THIS SECTION

Syntax

```
clear network-access aaa subscriber
<session-id identifier <reconnect>>
<statistics username username>
<username username <reconnect>>
```

Description

Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.

Options

reconnect (Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-

Terminate-Cause (RADIUS attribute 49) with a value of callback (16). You can apply this option to either a Layer 2 wholesale session or a conventionally auto-sensed dynamic VLAN supporting a PPPoE session.

In the latter case, this option triggers a PPPoE session logout and removal of the dynamic VLAN logical interface. This is followed by authorization of the access-line to attempt creation of a dynamic VLAN IFL supporting Layer 2 wholesale session in its place.

session-id *identifier* — (Optional) Log out the subscriber based on the subscriber session identifier.

statistics username *username* — (Optional) Clear AAA subscriber statistics and log out the subscriber.

username *username* — (Optional) Log out the AAA subscriber.

Required Privilege Level

maintenance

Output Fields

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

Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber statistics username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber username (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

clear network-access aaa subscriber session-id (Tenant systems)

```
user@host:TSYS1> clear network-access aaa subscriber session-id 1
```

clear pppoe lockout

IN THIS SECTION

Syntax

```
clear pppoe lockout
<aci circuit-id | mac-address mac-address >
<underlying-interfaces underlying-interface-name>
```

Description

Clear the lockout condition for the PPPoE client associated with the specified media access control (MAC) source address or agent circuit identifier (ACI) value.

Options

none

Clear the lockout condition for the PPPoE clients associated with all MAC source addresses on all PPPoE underlying interfaces.

aci *circuit-id*

(Optional) Clear the lockout condition for the PPPoE client associated with the specified ACI value. To clear the lockout condition by a specified ACI value, you must specify the `filter aci` option in the `short-cycle-protection` statement when you configure PPPoE subscriber session lockout. If the `filter aci` option is missing from the `short-cycle-protection` statement , no PPPoE client sessions are cleared using the ACI filter. The `aci` option and the `mac-address` option are mutually exclusive.

mac-address *mac-address*

(Optional) Clear the lockout condition for the PPPoE client associated with the specified MAC source address. The `mac-address` option and the `aci` option are mutually exclusive.

underlying-interfaces *underlying-interface-name*

(Optional) Clear the lockout condition for all PPPoE clients associated with the specified PPPoE underlying interface.

Required Privilege Level

clear

Sample Output

clear pppoe lockout (All MAC Source Addresses on All Underlying Interfaces)

```
user@host> clear pppoe lockout
```

clear pppoe lockout mac-address (Specified MAC Source Address)

```
user@host> clear pppoe lockout mac-address 00:00:5e:00:53:30
```

clear pppoe lockout mac-address underlying-interfaces (Specified MAC Source Address on Specified Underlying Interface)

```
user@host> clear pppoe lockout mac-address 00:00:5e:00:53:30 underlying-interfaces ge-1/0/0.101
```

clear pppoe lockout underlying-interfaces (All MAC Source Addresses on Specified Underlying Interface)

```
user@host> clear pppoe lockout underlying-interfaces ge-1/0/0.101
```

clear pppoe lockout underlying-interfaces aci (ACI on Specified Underlying Interface)

```
user@host> clear pppoe lockout underlying-interfaces demux0.214 aci “Relay-identifier atm 
3/0:100\.*”
```

clear pppoe lockout vlan-identifier

IN THIS SECTION

Syntax

```
clear pppoe lockout vlan-identifier device-name device-name
<aci circuit-id | mac-address mac-address >
<svlan-id svlan-identifier>
<vlan-id vlan-identifier>
```

Description

Clear the lockout condition for the PPPoE client associated with the specified VLAN encapsulation type and, optionally, media access control (MAC) source address and agent circuit identifier (ACI) value. Because the lockout condition persists even in the absence of an underlying interface or after automatic removal of the VLAN or VLAN demux interface, using the `clear pppoe lockout vlan-identifier` command enables you to clear the lockout condition for PPPoE clients by specifying VLAN identifying characteristics rather than by specifying the underlying interface name.

The following characteristics comprise the VLAN encapsulation type identifier:

- Device name (physical interface or aggregated Ethernet bundle)
- Stacked VLAN (S-VLAN) ID (also known as the *outer tag*)
- VLAN ID (also known as the *inner tag*)

You can configure PPPoE subscriber session lockout, also known as PPPoE short-cycle protection, for VLAN, VLAN demux, and PPPoE-over-ATM dynamic subscriber interfaces.

Options

circuit-id (Optional) ACI value associated with the PPPoE client for which you want to clear lockout. To clear the lockout condition by a specified ACI value, you must specify the `filter aci` option in the `short-cycle-protection` statement when you configure PPPoE subscriber session lockout. If the `filter aci` option is missing from the `short-cycle-protection` statement, no PPPoE client sessions are cleared using the ACI filter. The `aci` option and the `mac-address` option are mutually exclusive.

device-name Name of the Ethernet physical interface or aggregated Ethernet bundle associated with the PPPoE client for which you want to clear lockout.

mac-address (Optional) MAC address value associated with the PPPoE client for which you want to clear lockout. The `mac-address` option and the `aci` option are mutually exclusive.

svlan-identifier (Optional) A valid S-VLAN identifier associated with the PPPoE client for which you want to clear lockout.

- **Range:** 1 through 4094

vlan-identifier (Optional) A valid VLAN identifier associated with the PPPoE client for which you want to clear lockout.

- **Range:** 1 through 4094

Required Privilege Level

clear

Sample Output

clear pppoe lockout vlan-identifier device-name (Untagged VLAN on Aggregated Ethernet Bundle)

```
user@host> clear pppoe lockout vlan-identifier device-name ae3
```

clear pppoe lockout vlan-identifier device-name vlan-id (Single-Tagged VLAN on Gigabit Ethernet Interface)

```
user@host> clear pppoe lockout vlan-identifier device-name ge-2/0/0 vlan-id 2000
```

clear pppoe lockout vlan-identifier device-name svlan-id vlan-id aci (Dual-Tagged VLAN on 10-Gigabit Ethernet Interface Where ACI Matches Regular Expression)

```
user@host> clear pppoe lockout vlan-identifier device-name xe-1/0/0 svlan-id 10 vlan-id 20 aci
““Relay-identifier atm 1/0:100\.*”
```

clear pppoe lockout vlan-identifier device-name svlan-id vlan-id mac-address (Dual-Tagged VLAN on Aggregated Ethernet Bundle with Specified MAC Address)

```
user@host> clear pppoe lockout vlan-identifier device-name ae0 svlan-id 1 vlan-id 100 mac-
address 00:00:5e:00:53:30
```

clear pppoe statistics

IN THIS SECTION

Syntax

```
clear pppoe statistics
<interface interface-name>
<underlying-interface-name>
```

Description

Reset PPPoE session statistics information.

Options

none	Reset PPPoE statistics for all interfaces.
underlying-interface-name	(Optional) Reset PPPoE statistics for the specified underlying PPPoE interface.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear pppoe statistics

```
user@host> clear pppoe statistics
```

clear pppoe statistics

```
user@host> clear pppoe statistics ge-4/0/3.2
```

clear services l2tp destination

IN THIS SECTION

Syntax

```
clear services l2tp destination
<all | local-gateway gateway-address | peer-gateway gateway-address>
```

Description

Clear all Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers.

NOTE: You cannot issue the `clear services l2tp destination` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all

Close all L2TP destinations.

BEST PRACTICE: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment.

Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

local-gateway *gateway-address*

Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.

peer-gateway *gateway-address*

Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear services l2tp destination all

```
user@host> clear services l2tp destination all

  Destination 2 closed
```

clear services l2tp destination lockout

IN THIS SECTION

Syntax

```
clear services l2tp destination lockout
<all | local-gateway gateway-address | peer-gateway gateway-address>
```

Description

Clear the lockout timer for all or only the specified Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. Clearing the lockout timer removes the destination from the lockout list. This command is available only for LAC on MX Series routers.

NOTE: You cannot issue the `clear services l2tp destination` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics

`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all	(Optional) Unlock all L2TP destinations.
local-gateway *gateway-address*	(Optional) Unlock only the L2TP destination with the specified local gateway address.
peer-gateway *gateway-address*	(Optional) Unlock only the L2TP destination with the specified address.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear services l2tp destination lockout all

```
user@host> clear services l2tp destination lockout all
```

clear services l2tp session

IN THIS SECTION

Syntax

```
clear services l2tp session (all | interface interface-name | local-gateway gateway-address |
local-gateway-name gateway-name | local-session-id session-id  | local-tunnel-id tunnel-id |
peer-gateway gateway-address | peer-gateway-name gateway-name | routing-instance routing-
instance-name | tunnel-group group-name | user username)
```

Description

Clear L2TP sessions on LAC and LNS.

NOTE: On MX Series routers, you cannot issue the `clear services l2tp session` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all

Close all L2TP sessions.

BEST PRACTICE: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

interface *interface-name*

Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- `si-`*`fpc`*`/`*`pic`*`/`*`port`*—MPCs on MX Series routers only.

local-gateway *gateway-address*
Clear only the L2TP sessions associated with the specified local gateway address.

local-gateway-name *gateway-name*
Clear only the L2TP sessions associated with the specified local gateway name.

local-session-id *session-id*
Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id *tunnel-id*
Clear only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway *gateway-address*
Clear only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*
Clear only the L2TP sessions associated with the peer gateway with the specified name.

routing-instance *routing-instance-name*
Clear only the L2TP sessions associated with the specified routing instance.

tunnel-group *group-name*
Clear only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

user *username*
Clear only the L2TP sessions for the specified username.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear services l2tp session

```
user@host> clear services l2tp session 31694

  Session 31694 closed
```

Sample Output

clear services l2tp session interface

```
user@host> show services l2tp session Tunnel local ID: 17185
Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  5117   1        Established          1073741828        si-2/0/0
  34915  2        Established          1073741829        si-2/1/0
  6454   3        Established          1073741830        si-2/0/0
  46142  4        Established          1073741831        si-2/1/0
```

command-name

```
user@host> clear services l2tp session interface si-2/0/0
  Session  5117 closed
  Session  6454 closed
```

command-name

```
user@host> show services l2tp session Tunnel local ID: 17185
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  34915  2        Established          1073741829        si-2/1/0
  46142  4        Established          1073741831        si-2/1/0
```

clear services l2tp session statistics

IN THIS SECTION

Syntax

```
clear services l2tp session statistics (all | interface interface-name | local-gateway gateway-address | local-gateway-name gateway-name | local-session-id session-id | local-tunnel-id tunnel-id | peer-gateway gateway-address | peer-gateway-name gateway-name  | tunnel-group group-name | user username)
```

Description

Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions.

Options

all	Clear statistics for all L2TP sessions.
interface *interface-name*	Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- `si-`*fpc*/*pic*/*port*—MPCs on MX Series routers only.

local-gateway *gateway-address*

Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.

local-gateway-name *gateway-name*

Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.

local-session-id *session-id*

Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id *tunnel-id*

Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway *gateway-address*

Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*

Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.

tunnel-group *group-name*

Clear statistics for only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

user *username*

Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Output Fields

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

Sample Output

clear services l2tp session statistics all

```
user@host> clear services l2tp session statistics all
  Session 26497 statistics cleared
```

clear services l2tp tunnel

IN THIS SECTION

Syntax

```
clear services l2tp tunnel (all | interface sp-fpc/pic/port | local-gateway gateway-address |
local-gateway-name gateway-name | local-tunnel-id tunnel-id | peer-gateway gateway-address |
peer-gateway-name gateway-name | tunnel-group group-name)
```

Description

Clear Layer 2 Tunneling Protocol (L2TP) tunnels.

NOTE: You cannot issue the `clear services l2tp tunnel` command in parallel with statistics-related `show services l2tp` commands from separate terminals. If this `clear` command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the `show` commands listed in the following table:

`show services l2tp destination extensive`	`show services l2tp summary statistics`
`show services l2tp destination statistics`	`show services l2tp tunnel extensive`
`show services l2tp session extensive`	`show services l2tp tunnel statistics`
`show services l2tp session statistics`	

Options

all

Clear all L2TP tunnels.

BEST PRACTICE: The `all` option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the `all` option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

local-gateway *gateway-address*

Clear only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*

Clear only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*

Clear only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*
Clear only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*
Clear only the L2TP tunnels associated with the peer gateway with the specified name.

tunnel-group *group-name*
Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Output Fields

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

Sample Output

clear services l2tp tunnel

```
user@host> clear services l2tp tunnel 17185

  Tunnel 17185 closed
```

clear services l2tp tunnel statistics

IN THIS SECTION

Syntax

```
clear services l2tp tunnel statistics (all | interface sp-fpc/pic/port | local-gateway gateway-
address | local-gateway-name gateway-name | local-tunnel-id tunnel-id | peer-gateway gateway-
address | peer-gateway-name gateway-name | tunnel-group group-name)
```

Description

Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels (LAC only).

Options

all	Clear statistics for all L2TP tunnels.
local-gateway *gateway-address*	Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*
Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*
Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*
Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*
Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.

Required Privilege Level

clear

Output Fields

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

Sample Output

clear services l2tp tunnel statistics all

```
user@host> clear services l2tp tunnel statistics all
  Tunnel  9933 statistics cleared
```

clear system subscriber-management statistics

IN THIS SECTION

Syntax

```
clear system subscriber-management statistics
```

Description

Clear subscriber-management statistics.

Options

This command has no options.

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 subscriber-management statistics

```
user@host> clear subscriber-management statistics
```

request dhcp server reconfigure

IN THIS SECTION

Syntax

```
request dhcp server reconfigure (all | address | interface interface-name | logical-system logical-system-
name | routing-instance routing-instance-name)
```

Description

Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcp server binding` command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all — Initiate reconfiguration for all DHCP clients.

address — Initiate reconfiguration for DHCP client with the specified IP address or MAC address.

interface *interface-name* — Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).

> **NOTE**: You cannot use the `interface` *`interface-name`* option with the `request dhcp server reconfigure` command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.

logical-system *logical-system-name* — Initiate reconfiguration for all DHCP clients on the specified logical system.

routing-instance *routing-instance-name* — Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.

Required Privilege Level

view

Output Fields

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

Sample Output

request dhcp server reconfigure

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

request dhcpv6 server reconfigure

IN THIS SECTION

Syntax

```
request dhcpv6 server reconfigure (all | address | client-id | interface interface-name | logical-system
logical-system-name | routing-instance routing-instance-name | session-id)
```

Description

Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the `clear dhcpv6 server binding` command.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the `clear-on-abort` statement to configure the client to be cleared when reconfiguration fails.

Options

all
: Initiate reconfiguration for all DHCPv6 clients.

address
: Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.

client-id
: Initiate reconfiguration for DHCPv6 client with the specified client ID.

interface *interface-name*
: Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).

logical-system *logical-system-name*
: Initiate reconfiguration for all DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name*
: Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.

session-id
: Initiate reconfiguration for DHCPv6 client with the specified session ID.

Required Privilege Level

view

Output Fields

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

Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001db8::2/16
```

request interface rebalance

IN THIS SECTION

Syntax

```
request interface rebalance interface interface-name
```

Description

Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled. For example, when you add a new member link to an existing aggregated Ethernet interface, you might want to rebalance the subscribers across the new member link.

Options

interface-name Aggregated Ethernet logical interface number.

Additional Information

The `targeted-distribution` statement in a dynamic profile automatically distributes subscribers to the aggregated Ethernet links associated with the profile. If you apply the profile and then bring up a single link on an aggregated Ethernet interface, all newly connected subscribers are located on that single, active link.

When you then bring up a second link on the aggregated Ethernet interface, you can use the `request interface rebalance interface` command to rebalance the existing subscribers so that some remain on the first link and others are redistributed to the second link.

In some cases, it might appear that rebalancing does not work. By default, the router rebalances the subscribers only when the difference in subscriber numbers on the member links is greater than 500. This number is the default threshold for subscriber granularity.

Consider the following cases with two member links on the interface:

- Case 1—One link has 300 subscribers and the other link has 10 subscribers. Issuing the command does not rebalance the subscribers because the difference between the links is only 290.

- Case 2—One link has 530 subscribers and the other link has 10 subscribers. The difference between the links is 520 subscribers. Issuing the command rebalances the subscribers so that 500 subscribers are on one link and 40 subscribers are on the other link.

You can modify the subscriber granularity with the `rebalance-subscriber-granularity` command. For example, if you set the granularity to 1 subscriber, then rebalancing always takes place when the member links differ by 2 or more subscribers.

BEST PRACTICE: We recommend that you do not configure a low granularity value. A low value can have undesirable effects, such as the router running out of pseudo logical interfaces or an increase in the convergence time for rebalancing.

Leaving subscriber granularity at the default value of 500 subscribers is sufficient in most cases. Whenever more than one member link is active, targeted distribution places new subscribers on a link with fewer subscribers than other member links on the interface.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. You can compare the output of the `show interfaces targeting` command before and after the request to view the effect of the rebalancing operation. .

Sample Output

request interface rebalance (Adding a New Link)

```
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 1000

Distribution Group name: default
Number of distributed interfaces: 1000
Physical interface: xe-0/2/0, Link status: Up
Number of primary distributions: 1000     /*All subscribers are distributed on this active link*/
Number of backup distributions: 0
Physical interface: xe-0/2/1, Link status: Down   /*This link is inactive*/
Number of primary distributions: 0     /*No subscribers*/
Number of backup distributions: 1000
user@host >request interface rebalance interface ae0 (After Link xe-0/2/1 is Activated)
Rebalance operation on interface ae0 started
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 1000

Distribution Group name: default
Number of distributed interfaces: 1000
Physical interface: xe-0/2/0, Link status: Up
  Number of primary distributions: 500    /*500 subscribers were redistributed from this link*/
  Number of backup  distributions: 500
Physical interface: xe-0/2/1, Link status: Up    /*member link is now active*/
  Number of primary distributions: 500    /*500 subscribers were redistributed to this newly
active member link*/
  Number of backup  distributions: 500
```

request interface rebalance (Rebalancing Two Existing Links)

```
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 540

Distribution Group name: default
Number of distributed interfaces: 540
Physical interface: xe-0/2/0, Link status: Up
```

```
  Number of primary distributions: 530     /*This interface has 520 more subscribers than the
other*/
  Number of backup  distributions: 10
Physical interface: xe-0/2/1, Link status: Up
  Number of primary distributions: 10
  Number of backup  distributions: 530

user@host >request interface rebalance interface ae0
Rebalance operation on interface ae0 started
user@host >show interfaces targeting ae0
Aggregated interface: ae0
Targeting Type: Auto
Redundancy mode: Link Level Redundancy
Total number of distribution groups: 1
Total number of distributed interfaces: 540

Distribution Group name: default
Number of distributed interfaces: 540
Physical interface: xe-0/2/0, Link status: Up
  Number of primary distributions: 500    /*30 subscribers were redistributed from this link*/
  Number of backup  distributions: 40
Physical interface: xe-0/2/1, Link status: Up
  Number of primary distributions: 40    /*30 subscribers were redistributed to the link with
fewer subscribers*/
  Number of backup  distributions: 500
```

request network-access aaa accounting

IN THIS SECTION

Syntax

```
request network-access aaa accounting (baseline | suspend | resume)
```

Description

Suspend accounting processes; determine a baseline of the statistical details while accounting is suspended; and restart accounting operations after baselining is completed. This command is useful in service provider environments when an upgrade of the server infrastructure is critical and needed immediately. RADIUS Acct-Start, Interim-Update, and Acct-Stop messages are not generated while accounting is suspended; the router does not send any accounting messages to the RADIUS server. While accounting is suspended, subscribers can continue to log in and log out.

Options

baseline (Optional) Determine a baseline of accounting statistics for current subscriber sessions. Applies to only those subscribers for which interim accounting is configured. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline when you retrieve baseline-relative statistics after accounting resumes.

resume Restart the accounting processes for all logged-in subscriber sessions after baselining of statistics completes.

suspend Temporarily halt accounting processes for all logged-in subscriber sessions.

Required Privilege Level

view

Sample Output

request network-access aaa accounting suspend

```
user@host> request network-access aaa accounting suspend
```

request network-access aaa accounting baseline

```
user@host> request network-access aaa accounting baseline
```

request network-access aaa accounting resume

```
user@host> request network-access aaa accounting resume
```

request network-access aaa replay pending-accounting-stops

IN THIS SECTION

Syntax

```
request network-access aaa replay pending-accounting-stops
```

Description

Force the router to attempt contact with the accounting sever immediately, rather than allowing it to wait until the periodic interval has expired.

Options

This command has no options.

Required Privilege Level

view

Output Fields

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

Sample Output

request network-access aaa replay pending-accounting-stops

```
user@host> request network-access aaa replay pending-accounting-stops
replay started
```

request network-access aaa subscriber modify session-id

IN THIS SECTION

Syntax

```
request network-access aaa subscriber modify session-id subscriber-session-id predefined-
variable variable-option
```

Description

Modify a predefined variable that is applied to a subscriber who is currently logged in to the network.

Options

predefined-variable Name of the predefined variable that you want to modify.

subscriber-session-id ID of the subscriber session.

variable-option Name of the variable option that you want to apply to the predefined variable.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 42 on page 729 lists possible messages that might be returned.

Table 42: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Successful completion`	Variable was successfully modified	–
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.

Sample Output

request network-access aaa subscriber modify session-id

```
user@host> request network-access aaa subscriber modify session-id 49 junos-cos-traffic-control-
profile TCP-gold
Successful completion
```

request network-access aaa subscriber set session-id

IN THIS SECTION

Syntax

```
request network-access aaa subscriber set session-id subscriber-session-id provisioning-state
none
```

Description

Release control of the PCRF over the specified subscriber session. In response, AAA clears the subscriber’s provisioning state and sends a terminated request to the PCRF indicating the subscriber is no longer available.

Options

subscriber-session-id ID of the subscriber session.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. Table 43 on page 731 lists possible error messages that might be returned if the service activation fails.

Table 43: Service Activation/Deactivation Error Messages

Message	Description	Corrective Action
`Error: AUTHD ISSU in progress`	A unified ISSU operation is active.	Wait until the unified ISSU operation completes and then retry the service activation/deactivation.
`Service activation/ deactivation already in progress`	Another service activation/ deactivation operation is currently in progress.	Wait until the active operation completes and then retry the activation/deactivation operation.
`Session identifier is not for a subscriber session`	The session ID is incorrect.	Verify the correct session ID for the subscriber and then retry the activation/deactivation operation.

Sample Output

request network-access aaa subscriber set session-id

```
user@host> request network-access aaa subscriber set session-id session-id 49 provisioning-state
none
Successful completion
```

show accounting pending-accounting-stops

IN THIS SECTION

Syntax

```
show accounting pending-accounting-stops
<detail | terse>
<profile-name>
```

Description

Display all statistics for all pending accounting stop requests, including both service and session requests.

Options

none Display information for all access profiles.

detail | terse (Optional) Display the specified level of output.

profile-name (Optional) Particular access profile for which you want to display accounting stop statistics.

Required Privilege Level

view

Output Fields

Table 44 on page 733 lists the output fields for the `show accounting pending-accounting-stops` command. Output fields are listed in the approximate order in which they appear.

Table 44: show accounting pending-accounting-stops Output Fields

Field Name	Field Description	Level of Output
`Type`	Type of client.	All levels
`Username`	Name of the user logged in to the session.	All levels
`Logical system/Routing instance`	Logical system and routing instance used for the session.	`detail` none
`Access-profile`	Access profile used for AAA services for the session.	`detail` none
`Session ID`	ID of the subscriber session; generated when the subscriber logs in. In the `Service name` block, this is the ID of the service session.	All levels
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	`detail` none

Table 44: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`IP Address`	IP address of the subscriber.	`detail` none
`IPv6 Prefix`	IPv6 address of the subscriber.	`detail` none
`Authentication State`	State of the subscriber authentication session: `AuthInit`, `AuthStart`, `AuthChallenge`, `AuthRedirect`, `AuthClntRespWait`, `AuthAcctVolStatsAckWait`, `AuthAcctStopAckWait`, `AuthServCreateRespWait`, `AuthLogoutStart`, `AuthStateActive`, `AuthClntLogoutRespWait`, `AuthProfileUpdateWait`, `AuthProvisionRespWait`, `AuthProvisionServiceCreationWait`	`detail` none
`Accounting State`	State of the subscriber accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	`detail` none
`Service name`	Name of the attached service or policy.	`detail` none
`Service State`	State of the service provided in the subscriber session.	`detail` none
`Session uptime`	How long the session has been up, in *HH:MM:SS*.	`detail` none
`Accounting status`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time` or `volume+time`. Configured in RADIUS Service-Statistics VSA [26-69].	`detail` none
`Service accounting session ID`	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	`detail` none

Table 44: show accounting pending-accounting-stops Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Service accounting state	State of the service accounting session: Acc-Init, Acc-Start-Sent, Imm-Update-Stats-Pending, Acc-Interim-Sent, Acc-Stop-Stats-Pending, Acc-Stop-Sent, Acc-Stop-On-Fail-Deny-Sent, Acc-Stop-Ackd	detail none
Accounting interim interval	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail none
Subscriber ID	ID of the subscriber; generated when the subscriber logs in.	detail none
Service ID	ID of the subscriber service.	All levels
Service	Name of the attached service or policy.	terse

Sample Output

show accounting pending-accounting-stops detail

```
user@host> show accounting pending-accounting-stops detail
Type: pppoe
Username: vjshah29@example.com
AAA Logical system/Routing instance: default:default
Access-profile: ce-ppp-profile
Session ID: 84
Accounting Session ID: 84
IP Address: 192.168.0.25
IPv6 Prefix: 2010:db8:9999:18::/48
Authentication State: AuthAcctStopAckWait
Accounting State: Acc-Stop-Stats-Pending
Service name: cos-service
  Service State: SvcInactive
  Session ID: 94
  Session uptime: 00:08:02
  Accounting status: on/time
  Service accounting session ID: 84:94-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service
  Service State: SvcInactive
  Session ID: 93
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:93-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
Service name: filter-service6
  Service State: SvcInactive
  Session ID: 95
  Session uptime: 00:08:02
  Accounting status: on/volume+time
  Service accounting session ID: 84:95-1352294677
  Service accounting state: Acc-Stop-Stats-Pending
  Accounting interim interval: 600
```

show accounting pending-accounting-stops (Specific Profile)

```
user@host> show accounting pending-accounting-stops ce-ppp-profile
Type:         Username:               Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com       84
  pppoe        vjshah29@example.com       84              94                cos-service
  pppoe        vjshah29@example.com       84              93                filter-service
  pppoe        vjshah29@example.com       84              95                filter-service6
```

show accounting pending-accounting-stops terse

```
user@host> show accounting pending-accounting-stops terse
 Type:         Username:                   Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com           84              94                cos-service
```

```
user@host> show accounting pending-accounting-stops terse
 Type:         Username:                   Session ID:     Service ID:       Service
  pppoe        vjshah29@example.com           84
  pppoe        vjshah29@example.com           84              94                cos-service
  pppoe        vjshah29@example.com           84              93                filter-service
  pppoe        vjshah29@example.com           84              95                filter-service6
  pppoe        larry@example.com              85
  pppoe        larry@example.com              85              94                cos-service
  pppoe        larry@example.com              85              93                filter-service
  pppoe        larry@example.com              85              95                filter-service6
```

show class-of-service classifier

IN THIS SECTION

Syntax

```
show class-of-service classifier
<name name>
<type dscp | type dscp-ipv6 | type exp | type ieee-802.1 | type inet-precedence>
```

Description

For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.

Options

none Display all classifiers.

name *name* (Optional) Display named classifier.

type dscp (Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.

type dscp-ipv6 (Optional) Display all classifiers of the DSCP for IPv6 type.

type exp (Optional) Display all classifiers of the MPLS experimental (EXP) type.

type ieee-802.1 (Optional) Display all classifiers of the ieee-802.1 type.

type inet-precedence (Optional) Display all classifiers of the inet-precedence type.

Required Privilege Level

view

Output Fields

Table 45 on page 738 describes the output fields for the `show class-of-service classifier` command. Output fields are listed in the approximate order in which they appear.

Table 45: show class-of-service classifier Output Fields

Field Name	Field Description
`Classifier`	Name of the classifier.
`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.

Table 45: show class-of-service classifier Output Fields *(Continued)*

Field Name	Field Description
`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 class-of-service classifier type ieee-802.1

```
user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point         Forwarding Class                  Loss priority
  000              best-effort                         low
  001              best-effort                         high
  010              expedited-forwarding                low
  011              expedited-forwarding                high
  100              assured-forwarding                  low
  101              assured-forwarding                  medium-high
  110              network-control                     low
  111              network-control                     high

Classifier: users-ieee802.1, Code point type: ieee-802.1
  Code point         Forwarding class                    Loss priority
  100                expedited-forwarding                low
```

show class-of-service drop-profile

IN THIS SECTION

Syntax

```
show class-of-service drop-profile
<profile-name profile-name>
```

Description

Display data points for each class-of-service (CoS) random early detection (RED) drop profile.

Options

none	Display all drop profiles.
profile-name *profile-name*	(Optional) Display the specified profile only.

Required Privilege Level

view

Output Fields

Table 46 on page 741 describes the output fields for the `show class-of-service drop-profile` command. Output fields are listed in the approximate order in which they appear.

Table 46: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.
Type	Type of drop profile: • **discrete** (default) • **interpolated** (EX8200 switches, QFX Series switches, QFabric systems, EX4600 switches, OCX Series switches only)
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

Sample Output

show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
```

```
          100                   100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
    Fill level     Drop probability
             0                    0
             1                    1
             2                    2
             4                    4
             5                    5
             6                    6
             8                    8
            10                   10
            12                   15
            14                   20
            15                   23
          ... 64 entries total
            90                   96
            92                   96
            94                   97
            95                   98
            96                   98
            98                   99
            99                   99
           100                  100
```

show class-of-service interface

IN THIS SECTION

Syntax

```
show class-of-service interface <interface-name>
```

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

> **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 (Optional) Display class-of-service (CoS) associations for the specified interface.

none Display CoS associations for all physical and logical interfaces.

Required Privilege Level

view

Output Fields

Table 47 on page 744 describes the output fields for the `show class-of-service interface` command. Output fields are listed in the approximate order in which they appear.

Table 47: show 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 for MX Series routers) 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. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX-Series routers) 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. Supported only on Trio MPC/MIC interfaces on MX Series routers. (Enhanced subscriber management for MX Series routers) This field is not displayed for enhanced subscriber management.
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. The display designates between the control plane and the user plane by adding cp or up to the name. (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.

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, 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. The display designates between the control plane and the user plane by adding cp or up to the name. (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_UID1006) instead of with a subscriber interface.
Classifier	Name and type of classifiers associated with this interface. The display designates between the control plane and the user plane by adding cp or up to the name.
Forwarding-class-map	Name of the forwarding map associated with this interface.
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier, Fragmentation-map (for LSQ interfaces only), Scheduler-map, or Rewrite.
Name	Name of an object.
Type	Type of an object: dscp, dscp-ipv6, exp, ieee-802.1, ip, or inet-precedence.
Link-level type	Encapsulation on the physical interface.

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`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.
`Flow control`	Whether flow control is enabled or disabled.
`Auto-negotiation`	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
`Remote-fault`	(Gigabit Ethernet interfaces) Remote fault status. • `Online`—Autonegotiation is manually configured as online. • `Offline`—Autonegotiation is manually configured as offline.

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`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 47: show 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 47: show 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 47: show 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 47: show 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 47: show 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 47: show 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. **NOTE**: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the `Dropped packets` field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description

`SONET alarms`
`SONET defects`
	(SONET) 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`.

`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

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
`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 47: show 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 47: show 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 47: show 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 47: show 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 47: show 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) 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.
RED-dropped bytes	Number of bytes dropped because of RED. The byte counts vary by PIC type.

Table 47: show class-of-service interface Output Fields *(Continued)*

Field Name	Field Description
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.
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.

Table 47: show 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.

Table 47: show 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]. This feature is supported on MPC/MIC interfaces on MX Series routers. 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 class-of-service interface

```
user@host> show class-of-service interface up:up1:demux0.3221225472
Logical interface: up:up1:demux0.3221225472, Index: 3221225472
    Object                   Name                   Type             Index
    Traffic-control-profile  cp::tcpd_UID1024       Output            16456
    Scheduler-map            cp::smap1              Output       4294967298
    Classifier               cp::cl-ieee-1          ieee8021p         11469
```

show class-of-service interface (Physical)

```
user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Maximum usable queues: 8, Queues in use: 4
Total non—default queues created: 4
  Scheduler map: <default>, Index: 2032638653

  Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                   Name                   Type           Index
    Scheduler-map            <default>                              27
    Rewrite                  exp-default            exp             21
    Classifier               exp-default            exp             5
    Classifier               ipprec-compatibility   ip              8
    Forwarding—class—map     exp-default            exp             5
```

show class-of-service interface (Logical)

```
user@host> show class-of-service interface xe-0/2/3.0
Logical interface: xe-0/2/3.0, Index: 344
    Object                   Name                   Type             Index
    Traffic-control-profile  cp::tcp1               Output            16456
    Scheduler-map            cp::smap1              Output       4294967298
```

show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Maximum usable queues: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
    Object                  Name                   Type           Index
    Traffic-control-profile tcp-pppoe.o.pp0.1      Output         2726446535
    Classifier              ipprec-compatibility   ip             13

    Adjusting application: PPPoE
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface (DHCP Interface)

```
user@host> show class-of-service interface demux0.3221225472
Logical interface: demux0.3221225472, Index: 3221225472
    Object                  Name                   Type           Index
    Traffic-control-profile tcpd_UID1024           Output         16456
    Scheduler-map           cp::smap1              Output    4294967298
    Classifier              cp::cl-ieee-1          ieee8021p      11469
```

show class-of-service interface (PPPoE Subscriber Interface for Enhanced Subscriber Management)

```
user@host> show class-of-service interface pp0.3221225474
  Logical interface: pp0.3221225475, Index: 3221225475
Object                  Name                   Type                    Index
Traffic-control-profile TC_PROF_100_199_SERIES_UID1006 Output      4294967312
Scheduler-map           SMAP-1_UID1002         Output              4294967327
Rewrite-Output          ieee-rewrite           ieee8021p                60432
Rewrite-Output          rule1                  ip                       50463

    Adjusting application: PPPoE IA tags
      Adjustment type: absolute
      Configured shaping rate: 11000000
      Adjustment value: 5000000
      Adjustment target: node

    Adjusting application: ucac
      Adjustment type: delta
      Configured shaping rate: 5000000
      Adjustment value: 100000
      Adjustment target: node
```

show class-of-service interface-set

IN THIS SECTION

Syntax

```
show class-of-service interface-set
<interface-set-name>
```

Description

Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical class of service (CoS).

Options

none — Display CoS associations for all logical interface sets.

interface-set *interface-set-name* — (Optional) Display CoS associations for the specified interface set.

Required Privilege Level

view

Output Fields

Table 48 on page 768 describes the output fields for the `show class-of-service interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 48: show 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.
`Queues supported`	Number of queues you can configure on the interface.
`Queues in use`	Number of queues currently configured.
`Output traffic control profile`	Name of the output traffic control profile attached to the logical interface set.
`Output traffic control profile remaining`	(Enhanced subscriber management for MX Series routers) For dynamic subscriber management, name of the output traffic control profile for remaining traffic attached to the logical interface set.

Table 48: show class-of-service interface-set Output Fields *(Continued)*

Field Name	Field Description
`Adjusting application`	Name of the application that communicates shaping-rate adjustment information to the Junos OS class-of-service process (**cosd**) on the broadband services router (BSR). The BSR uses the information from this application to perform shaping-rate adjustments on the scheduler node that manages the interface set. The adjusting application appears as **ancp LS-0** which is the Junos OS Access Node Control Profile process (**ancpd**) that performs shaping-rate adjustments on schedule nodes. The nodes are logical interface sets configured to represent subscriber local loops. When the synchronization speed of the DSL line changes, **ancpd** communicates the local loop speed to **cosd** over the default logical system, **LS-0**, and then the BSR throttles the shaping rate on the scheduler node to the loop speed. 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]. This feature is supported on MPC/MIC interfaces on MX Series routers. 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 shaping-rate adjustment performed by the BSR on the scheduler node. The type of adjustment appears as **Adjustment type**, meaning that the configured shaping rate is adjusted by an absolute value as opposed to by a percentage of the configured rate.
`Configured shaping rate`	The maximum transmission rate on the physical interface as configured by the output traffic-control profile attached to the scheduler node.
`Adjustment value`	Value of the shaping-rate adjustment information sent by the adjusting application to **cosd**.
`Adjustment overhead-accounting mode`	Configured shaping mode: `frame` or `cell`.

Sample Output

show class-of-service interface-set

```
user@host> show class-of-service interface-set example-ifset-ge-4/0/0-7
Interface-set: example-ifset-ge-4/0/0-7, Index: 8
Physical interface: ge-4/0/0, Index: 270
Queues supported: 8, Queues in use: 8
  Output traffic control profile: example-tcp-basic-rate, Index: 11395
Adjusting application: ancp LS-0
  Adjustment type: absolute
  Configured shaping rate: 50000000
  Adjustment value: 888000
  Adjustment overhead-accounting mode: cell
```

show class-of-service interface-set (Enhanced Subscriber Management)

```
user@host> show class of service interface-set
Interface-set: ge-1/0/0-201-201, Index: 1
Physical interface: ge-1/0/0, Index: 142
Queues supported: 8, Queues in use: 4
  Output traffic control profile: LEVEL_2_UID1001, Index: 4294967307
  Output traffic control profile remaining: TCP_REMAIN_UID1003, Index: 4294967308
```

show class-of-service rewrite-rule

IN THIS SECTION

Syntax

```
show class-of-service rewrite-rule
<name name>
<type type>
```

Description

Display the mapping of forwarding classes and loss priority to code point values.

Options

none — Display all rewrite rules.

name *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**—For IPv4 traffic.
- **dscp-ipv6**—For IPv6 traffic.
- **exp**—For MPLS traffic.
- **frame-relay-de**—(SRX Series only) For Frame Relay traffic.
- **ieee-802.1**—For Layer 2 traffic.
- **inet-precedence**—For IPv4 traffic.

Required Privilege Level

view

Output Fields

Table 49 on page 772 describes the output fields for the `show class-of-service rewrite-rule` command. Output fields are listed in the approximate order in which they appear.

Table 49: show class-of-service rewrite-rule Output Fields

Field Name	Field Description
Rewrite rule	Name of the rewrite rule.
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.
Index	Internal index for this particular rewrite rule.
Loss priority	Loss priority for rewriting.
Code point	Code point value to rewrite.

Sample Output

show class-of-service rewrite-rule type dscp

```
user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class                    Loss priority       Code point
  gold                                high              000000
  silver                              low               110000
  silver                              high              111000
  bronze                              low               001010
  bronze                              high              001100
  lead                                high              101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
Forwarding class                    Loss priority       Code point
  gold                                low               000111
  gold                                high              001010
  silver                              low               110000
  silver                              high              111000
  bronze                              high              001100
  lead                                low               101110
  lead                                high              110111
```

show class-of-service traffic-control-profile

IN THIS SECTION

Syntax

```
show class-of-service traffic-control-profile
<profile-name>
```

Description

For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and MPC/MIC interfaces only, display traffic shaping and scheduling profiles. You can view configured items for both the control plane and the user plane.

Options

none	Display all profiles.
profile-name	(Optional) Display information about a single profile.

Required Privilege Level

view

Output Fields

Table 50 on page 774 describes the output fields for the `show class-of-service traffic-control-profile` command. Output fields are listed in the approximate order in which they appear.

Table 50: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
`Traffic control profile`	Name of the traffic control profile. You can configure objects of the same type with the same name on the user plane and the control plane. The display designates between the control plane and the user plane by adding cp or up to the name. Also, the dynamically generated UID is displayed.
`Index`	Index number of the traffic control profile.

Table 50: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`ATM Service`	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: • cbr—Constant bit rate. • rtvbr—Real time variable bit rate. • nrtvbr—Non real time variable bit rate. • ubr—Unspecified bit rate.
`Maximum Burst Size`	Configured maximum burst size, in cells.
`Peak rate`	Configured peak rate, in cps.
`Sustained rate`	Configured sustained rate, in cps.
`Shaping rate`	Configured shaping rate, in bps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
`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.
`Shaping rate priority high`	Configured shaping rate for high-priority traffic, in bps.
`Shaping rate priority medium`	Configured shaping rate for medium-priority traffic, in bps.
`Shaping rate priority low`	Configured shaping rate for low-priority traffic, in bps.
`Shaping rate excess high`	Configured shaping rate for high-priority excess traffic, in bps.

Table 50: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Shaping rate excess low`	Configured shaping rate for low-priority excess traffic, in bps.
`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.
`Delay Buffer rate`	Configured delay buffer rate, in bps.
`Excess rate`	Configured excess rate, in percent or proportion.
`Excess rate high`	Configured excess rate for high priority traffic, in percent or proportion.
`Excess rate low`	Configured excess rate for low priority traffic, in percent or proportion.
`Guaranteed rate`	Configured guaranteed rate, in bps or cps. **NOTE**: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM 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.
`Guaranteed rate burst`	Configured burst size for the guaranteed rate, in bytes.
`adjust-minimum`	Configured minimum shaping rate for an adjusted queue, in bps.
`overhead accounting mode`	Configured shaping mode: `Frame Mode` or `Cell Mode`.
`Overhead bytes`	Configured byte adjustment value.

Table 50: show class-of-service traffic-control-profile Output Fields *(Continued)*

Field Name	Field Description
`Adjust parent`	Configured shaping-rate adjustment for parent scheduler nodes. If enabled, this field appears. `flow-aware` indicates that the parent scheduler node is adjusted only once per multicast channel.

Sample Output

show class-of-service traffic-control-profile

```
user@host> show class-of-service traffic-control-profile
tcp1 {
    scheduler-map smap1;
    shaping-rate 8m;
}
cp::tcp1 {
    scheduler-map cp::smap8;
    shaping-rate 8m;
}
cp::tcpd_UID1024 {
    scheduler-map cp::smap1_UID0047;
    shaping-rate 9m;
}
```

show dhcp relay binding

IN THIS SECTION

Syntax

```
show dhcp relay binding
<address>
<brief>
<detail>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<ip-address | mac-address>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address (Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.

- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

brief

(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcp relay binding`.

detail

(Optional) Display detailed client binding information.

interface *interface-name*

(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan

(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard

(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name*

(Optional) Perform this operation on the specified logical system.

routing-instance *routing-instance-name*

(Optional) Perform this operation on the specified routing instance.

summary

(Optional) Display a summary of DHCP client information.

Required Privilege Level

view

Output Fields

Table 51 on page 780 lists the output fields for the `show dhcp relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 51: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail
State	State of the DHCP relay address binding table on the DHCP client: • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers.	brief detail

Table 51: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface`	Incoming client interface.	brief
`Lease Expires`	Date and time at which the client's IP address lease expires.	detail
`Lease Expires in`	Number of seconds in which the lease expires.	detail
`Lease Start`	Date and time at which the client's IP address lease started.	detail
`Lease time violated`	Lease time violation has occurred.	detail
`Incoming Client Interface`	Client's incoming interface.	detail
`Server IP Address`	IP address of the DHCP server.	detail
`Server Interface`	Interface of the DHCP server.	detail
`Bootp Relay Address`	IP address of BOOTP relay.	detail
`Type`	Type of DHCP packet processing performed on the router: • `active`—Router actively processes and relays DHCP packets. • `passive`—Router passively snoops DHCP packets passing through the router.	All levels
`Lease expires at`	Date and time at which the client's IP address lease expires.	All levels

Table 51: show dhcp relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Dual Stack Group`	Name of dual stack that is configured with the DHCP binding.	detail
`Dual Stack Peer Prefix`	Prefix of dual stack DHCPv6 peer.	detail
`Dual Stack Peer Address`	Address of the dual stack DHCPv6 peer.	detail

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.11     41          00:00:5e:00:53:01  86371       BOUND      ge-1/0/0.0
198.51.100.12     42          00:00:5e:00:53:02  86371       BOUND      ge-1/0/0.0
198.51.100.13     43          00:00:5e:00:53:03  86371       BOUND      ge-1/0/0.0
198.51.100.14     44          00:00:5e:00:53:04  86371       BOUND      ge-1/0/0.0
198.51.100.15     45          00:00:5e:00:53:05  86371       BOUND      ge-1/0/0.0
```

show dhcp relay binding detail

```
user@host> show dhcp relay binding detail


Client IP Address:  198.51.100.11
     Hardware Address:            00:00:5e:00:53:01
     State:                       BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:               2009-07-21 11:00:06 PDT
     Lease Expires in:            86361 seconds
     Lease Start:                 2009-07-20 11:00:06 PDT
     Lease time violated:         yes
     Last Packet Received:         2009-07-20 11:00:06 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.22
     Server Interface:             none
     Bootp Relay Address:          198.51.100.32
     Session Id:                   41
     Dual Stack Group:             dual-stack-retail6
     Dual Stack Peer Prefix:       2001:db8:0:4::/64
     Dual Stack Peer Address:      2001:db8:1:0:8003::1/128

Client IP Address:  198.51.100.12
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
     Lease Start:                  2009-07-20 11:00:06 PDT
     Last Packet Received:         2009-07-20 11:00:06 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.22
     Server Interface:             none
     Bootp Relay Address:          198.51.100.32
     Session Id:                   42
     Generated Remote ID           host:ge-1/0/0:100
```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2

IP address       Hardware address   Type     Lease expires at
198.51.100.1      00:00:5e:00:53:01  active   2007-03-27 15:06:20 EDT
```

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100

IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:94  86124       BOUND      ge-1/1/0:100
```

show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100

IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.16      7           00:00:5e:00:53:92  86124       BOUND      ge-1/1/0:10-100
```

show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 198.51.100.13
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.13      43          00:00:5e:00:53:03  86293       BOUND      ge-1/0/0.0
```

show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:00:5e:00:53:05
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      45          00:00:5e:00:53:05  86279       BOUND      ge-1/0/0.0
```

show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
 IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.11      41          00:00:5e:00:53:53  86305       BOUND      ge-1/0/0.0
```

show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17       42          00:00:5e:00:53:02  86346       BOUND      ge-1/0/0.1073741827
192.168.0.16       41          00:00:5e:00:53:01  86346       BOUND      ge-1/0/0.1073741827
```

show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
IP address          Session Id  Hardware address   Expires     State      Interface
192.168.0.9          24           00:00:5e:00:53:04  86361       BOUND       ge-1/3/0.110
192.168.0.8          23           00:00:5e:00:53:03  86361       BOUND       ge-1/3/0.110
192.168.0.7          22           00:00:5e:00:53:02  86361       BOUND       ge-1/3/0.110
```

show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)
```

show dhcp relay lockout-entries

IN THIS SECTION

Syntax

```
show dhcp relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 relay agent lockout database.

Options

all	Display all client entries in the lockout database.
index	Number identifying a client entry to be displayed.

Required Privilege Level

view

Output Fields

Table 52 on page 786 lists the output fields for the `show dhcp relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 52: show dhcp relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	Client identifier for the client in the lockout database.	`all` and `index`

Table 52: show dhcp relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	Type of lockout period for the entry: • Grace—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • Lockout—Client is currently locked out; attempts to establish a session are rejected.	all and index
Expires (s)	Number of seconds until the current lockout period expires.	all only
Elapsed (s)	Number of seconds since the current lockout or grace timer started.	all only
Count	Number of consecutive times the client has been locked out.	all only
Expires	Date and time when the current lockout period ends.	index only
Expires in	Number of seconds until the current period expires.	index only
Lockout count	Number of consecutive times client has been locked out.	index only
Next lockout time	Duration of the next lockout period for this client.	index only
Min lockout time	Minimum duration for a lockout period; the initial lockout time.	index only

Table 52: show dhcp relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp relay lockout-entries (All Entries)

```
user@host> show dhcp relay lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp relay lockout-entries (Specific Entry)

```
user@host> show dhcp relay lockout-entries index 2
     Index:                                 2
     Key:                                   default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                 Lockout
     Expires:                               2018-03-29 19:06:17 IST
     Expires in:                            87
     Lockout count:                         1
     Next lockout time:                     200
     Min lockout time:                      100
     Lockout reason:                        181
```

show dhcp relay statistics

IN THIS SECTION

Syntax

```
show dhcp relay statistics


<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

logical-system *logical-system-name*

(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

routing-instance *routing-instance-name*	(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 53 on page 791 lists the output fields for the `show dhcp relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 53: show dhcp relay statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the `Packets dropped` output. When all of the Packets dropped statistics are 0 (zero), only the `Total` field appears. • `Total`—Total number of packets discarded by the extended DHCP relay agent application. • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified. • `Bad opcode`—Number of packets discarded because an invalid operation code was specified. • `Bad options`—Number of packets discarded because invalid options were specified. • `Invalid server address`—Number of packets discarded because an invalid server address was specified. • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment. • `No interface match`—Number of packets discarded because they did not belong to a configured interface. • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance. • `No valid local address`—Number of packets discarded because there was no valid local address. • `Packet too short`—Number of packets discarded because they were too short. • `Read error`—Number of packets discarded because of a system read error. • `Send error`—Number of packets that the extended DHCP relay application could not send. • `Option 60`—Number of packets discarded containing DHCP option 60 vendor-specific information. • `Option 82`—Number of packets discarded because DHCP option 82 information could not be added.

Table 53: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages received`	Number of DHCP messages received. • `BOOTREQUEST`—Number of BOOTP protocol data units (PDUs) received • `DHCPDECLINE`—Number of DHCP PDUs of type DECLINE received • `DHCPDISCOVER`—Number of DHCP PDUs of type DISCOVER received • `DHCPINFORM`—Number of DHCP PDUs of type INFORM received • `DHCPRELEASE`—Number of DHCP PDUs of type RELEASE received • `DHCPREQUEST`—Number of DHCP PDUs of type REQUEST received • `DHCPLEASEACTIVE`—Number of active DHCP leases • `DHCPLEASEUNASSIGNED`—Number of DHCP leases that are managed by the server but have not yet been assigned • `DHCPLEASEUNKNOWN`—Number of unknown DHCP leases • `DHCPLEASEQUERYDONE`—The leasequery is complete
`Messages sent`	Number of DHCP messages sent. • `BOOTREPLY`—Number of BOOTP PDUs transmitted • `DHCPOFFER`—Number of DHCP OFFER PDUs transmitted • `DHCPACK`—Number of DHCP ACK PDUs transmitted • `DHCPNACK`—Number of DHCP NACK PDUs transmitted • `DHCPFORCERENEW`—Number of DHCP FORCERENEW PDUs transmitted • `DHCPLEASEQUERY`—Number of DHCP leasequery messages transmitted • `DHCPLEASEBULKLEASEQUERY`—Number of DHCP bulk leasequery messages transmitted
`External Server Response`	State of the external DHCP server responsiveness.

Table 53: show dhcp relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Packets forwarded`	Number of packets forwarded. • `BOOTREQUEST`—Number of BOOTREQUEST protocol data units (PDUs) forwarded • `BOOTREPLY`—Number of BOOTREPLY protocol data units (PDUs) forwarded
`External Server Response`	State of the external DHCP server responsiveness.

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                      34
    Bad hardware address       1
    Bad opcode                 1
    Bad options                3
    Invalid server address     5
    Lease Time Violation       1
    No available addresses     1
    No interface match         2
    No routing instance match  9
    No valid local address     4
    Packet too short           2
    Read error                 1
    Send error                 1
    Option 60                  1
    Option 82                  2

Messages received:
    BOOTREQUEST                116
    DHCPDECLINE                0
    DHCPDISCOVER              11
    DHCPINFORM                0
    DHCPRELEASE               0
    DHCPREQUEST               105
    DHCPLEASEACTIVE           0
    DHCPLEASEUNASSIGNED       0
    DHCPLEASEUNKNOWN          0
    DHCPLEASEQUERYDONE        0

Messages sent:
    BOOTREPLY                 0
    DHCPOFFER                 2
    DHCPACK                   1
    DHCPNAK                   0
    DHCPFORCERENEW            0
    DHCPLEASEQUERY            0
    DHCPBULKLEASEQUERY        0

Packets forwarded:
    Total                     4
    BOOTREQUEST               2
    BOOTREPLY                 2

External Server Response:
    State                     Responding
```

show dhcp server binding

IN THIS SECTION

Syntax

```
show dhcp server binding
<address>
<interfaces-vlan><brief | detail | summary>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.

> **NOTE**: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

Options

address (Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.

- *session-id*—The specified session ID.

brief | detail | summary
(Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcp server binding`.

interface *interface-name*
(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan
(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard
(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name*
(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

routing-instance *routing-instance-name*
(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

Required Privilege Level

view

Output Fields

Table 54 on page 796 lists the output fields for the `show dhcp server binding` command. Output fields are listed in the approximate order in which they appear.

Table 54: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
`number` `clients, (`*`number`* `init,` *`number`* `bound,` *`number`* `selecting,` *`number`* `requesting,` *`number`* `renewing,` *`number`* `releasing)`	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	`summary`

Table 54: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: • BOUND—Client has active IP address lease. • FORCERENEW—Client has received forcerenew message from server. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCP server. • SELECTING—Client receiving offers from DHCP servers.	brief detail
Interface	Interface on which the request was received.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail

Table 54: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Lease Expires in	Number of seconds in which lease expires.	detail
Lease Start	Date and time at which the client’s IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	detail
Last Packet Received	Date and time at which the router received the last packet.	detail
Incoming Client Interface	Client’s incoming interface.	detail
Client Interface Svlan Id	S-VLAN ID of the client’s incoming interface.	detail
Client Interface Vlan Id	VLAN ID of the client’s incoming interface.	detail
Demux Interface	Name of the IP demultiplexing (demux) interface.	detail
Server IP Address or Server Identifier	IP address of DHCP server.	detail
Server Interface	Interface of DHCP server.	detail
Client Pool Name	Name of address pool used to assign client IP address lease.	detail

Table 54: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Liveness Detection State`	State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session: **NOTE**: This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber. • `DOWN`—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The `DOWN` state is reported only during this transition period of time. • `UNKNOWN`—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The `UNKNOWN` state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the `UP` state. • `UP`—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have *both* determined that the liveness detection session for the BFD protocol is in the `UP` state. • `WENT_DOWN`—State is functionally equivalent to the `DOWN` state. A liveness detection session that was previously in an `UP` state has transitioned to a `DOWN` state implying a liveness detection failure.	`detail`

Table 54: show dhcp server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	The WENT_DOWN state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine.	
Client Profile Name	DHCP client profile name.	detail
Dual Stack Group	DHCP server profile name.	detail
Dual Stack Peer Prefix	IPv6 prefix of peer.	detail
Dual Stack Peer Address	IPv6 address of peer.	detail

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
IP address        Session Id  Hardware address   Expires     State      Interface
16.0.0.2          8           00:00:64:03:01:02  99903       BOUND      up:green-arrow:ge-0/3/5.2
```

show dhcp server binding detail

```
user@host> show dhcp server binding detail
Client IP Address:  198.51.100.15
     Hardware Address:             00:00:5e:00:53:01
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.9
     Server Interface:             none
     Session Id:                   6
     Client Pool Name:             6
     Liveness Detection State:     UP
 Client IP Address:            198.51.100.16
     Hardware Address:             00:00:5e:00:53:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Lease time violated:          yes
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            198.51.100.9
     Server Interface:             none
     Session Id:                   7
     Client Pool Name:             7
     Liveness Detection State:     UP
```

When DHCP binding is configured with dual-stack, we get the following output:

```
user@host> show dhcp server binding detail
Client IP Address:  198.51.100.10
     Hardware Address:             00:00:64:03:01:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
     Protocol-Used:                DHCP
     Lease Expires:                2016-11-07 08:30:39 PST
     Lease Expires in:             43706 seconds
     Lease Start:                  2016-11-04 11:00:37 PDT
     Last Packet Received:         2016-11-06 09:00:39 PST
     Incoming Client Interface:    ae0.3221225472
     Client Interface Svlan Id:    2000
     Client Interface Vlan Id:     1
     Server Ip Address:            198.51.100.2
     Session Id:                   2
     Client Pool Name:             my-v4-pool
     Client Profile Name:          dhcp-retail
     Dual Stack Group:             my-dual-stack
     Dual Stack Peer Prefix:       2001:db8:ffff:0:4::/64
     Dual Stack Peer Address:      2001:db8:0:8003::1/128
```

show dhcp server binding interface <vlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:100
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/1/0:100
```

show dhcp server binding interface <svlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.16      7          00:00:5e:00:53:02  86124       BOUND      ge-1/1/0:10-100
```

show dhcp server binding <ip-address>

```
user@host> show dhcp server binding 198100.19
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.19      10          00:00:5e:00:53:05  86081       BOUND      ge-1/0/0.0
```

show dhcp server binding <session-id>

```
user@host> show dhcp server binding 6
IP address        Session Id  Hardware address   Expires     State      Interface
198.51.100.15      6           00:00:5e:00:53:01  86124       BOUND      ge-1/0/0.0
```

show dhcp server binding summary

```
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcp server binding <interfaces-vlan>

```
user@host> show dhcp server binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17      42          00:00:5e:00:53:02  86346       BOUND      ge-1/0/0.1073741827
192.168.0.16      41          00:00:5e:00:53:01  86346       BOUND      ge-1/0/0.1073741827
```

show dhcp server binding <interfaces-wildcard>

```
user@host> show dhcp server binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9       24          00:00:5e:00:53:04  86361       BOUND      ge-1/3/0.110
192.168.0.8       23          00:00:5e:00:53:03  86361       BOUND      ge-1/3/0.110
192.168.0.7       22          00:00:5e:00:53:02  86361       BOUND      ge-1/3/0.110
```

show dhcp server lockout-entries

IN THIS SECTION

Syntax

```
show dhcp server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv4 local lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 55 on page 804 lists the output fields for the `show dhcp server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 55: show dhcp server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	Client identifier for the client in the lockout database.	`all` and `index`

Table 55: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	Type of lockout period for the entry: - `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. - `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only

Table 55: show dhcp server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcp server lockout-entries (All Entries)

```
user@host> show dhcp server lockout-entries all
Index     Key              State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00     Lockout     30           5200            2
2     00:00:5E:00:53:11     Grace      120           780             2
3     00:00:5E:00:53:22     Lockout    180           2300            1
```

show dhcp server lockout-entries (Specific Entry)

```
user@host> show dhcp server lockout-entries index 2
     Index:                                  2
     Key:                                    default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                  Lockout
     Expires:                                2018-03-29 19:06:17 IST
     Expires in:                             87
     Lockout count:                          1
     Next lockout time:                      200
     Min lockout time:                       100
     Lockout reason:                         181
```

show dhcp server statistics

IN THIS SECTION

Syntax

```
show dhcp server statistics
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

logical-system *logical-system-name*

(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.

routing-instance *routing-instance-name*	(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 56 on page 809 lists the output fields for the `show dhcp server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 56: show dhcp server statistics Output Fields

Field Name	Field Description
`Packets dropped`	Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • `Total`—Total number of packets discarded by the extended DHCP local server • `Authentication`—Number of packets discarded because they could not be authenticated • `Bad hardware address`—Number of packets discarded because an invalid hardware address was specified • `Bad opcode`—Number of packets discarded because an invalid operation code was specified • `Bad options`—Number of packets discarded because invalid options were specified • `Dynamic profile`—Number of packets discarded due to dynamic profile information • `Invalid server address`—Number of packets discarded because an invalid server address was specified • `Lease Time Violation`—Number of packets discarded because of a lease time violation • `No available addresses`—Number of packets discarded because there were no addresses available for assignment • `No interface match`—Number of packets discarded because they did not belong to a configured interface • `No routing instance match`—Number of packets discarded because they did not belong to a configured routing instance • `No valid local address`—Number of packets discarded because there was no valid local address • `Packet too short`—Number of packets discarded because they were too short • `Read error`—Number of packets discarded because of a system read error • `Send error`—Number of packets that the extended DHCP local server could not send

Table 56: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
`Offer Delay`	Number of DHCPv4 offer messages delayed. • `DELAYED`—Number of DHCPv4 offer packets that have been sent after being delayed. • `INPROGRESS`—Number of DHCPv4 offer packets that are in the delay queue. • `TOTAL`—Total number of delayed DHCPv4 offer messages; sum of `DELAYED` and `INPROGRESS`.
`Messages received`	Number of DHCP messages received. • `BOOTREQUEST`—Number of BOOTP protocol data units (PDUs) received • `DHCPDECLINE`—Number of DHCP PDUs of type DECLINE received • `DHCPDISCOVER`—Number of DHCP PDUs of type DISCOVER received • `DHCPINFORM`—Number of DHCP PDUs of type INFORM received • `DHCPRELEASE`—Number of DHCP PDUs of type RELEASE received • `DHCPREQUEST`—Number of DHCP PDUs of type REQUEST received • `DHCPLEASEQUERY`—Number of DHCP leasequery messages received. • `DHCPRENEW`—Number of DHCP renew messages received; subset of `DHCPREQUEST` counter. • `DHCPREBIND`—Number of DHCP rebind messages received; subset of `DHCPREQUEST` counter.

Table 56: show dhcp server statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of DHCP messages sent. • `BOOTREPLY`—Number of BOOTP PDUs transmitted • `DHCPOFFER`—Number of DHCP OFFER PDUs transmitted • `DHCPACK`—Number of DHCP ACK PDUs transmitted • `DHCPNACK`—Number of DHCP NACK PDUs transmitted • `DHCPFORCERENEW`—Number of DHCP FORCERENEW PDUs transmitted • `DHCPLEASEUNASSIGNED`—Number of DHCP leases that are managed by the server but have not yet been assigned • `DHCPLEASEUNKNOWN`—Number of unknown DHCP leases • `DHCPLEASEACTIVE`—Number of active DHCP leases • `DHCPLEASEQUERYDONE`—The leasequery is complete

Sample Output

show dhcp server statistics

```
user@host> show dhcp server statistics
Packets dropped:
    Total                      0

Offer Delay:
    DELAYED                    0
    INPROGRESS                 0
    TOTAL                      0

Messages received:
    BOOTREQUEST                4
    DHCPDECLINE                0
    DHCPDISCOVER               2
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                2
    DHCPLEASEQUERY             0
    DHCPBULKLEASEQUERY         0
    DHCPACTIVELEASEQUERY       0

Messages sent:
    BOOTREPLY                  4
    DHCPOFFER                  2
    DHCPACK                    2
    DHCPNAK                    0
    DHCPFORCERENEW             0
    DHCPLEASEUNASSIGNED        0
    DHCPLEASEUNKNOWN           0
    DHCPLEASEACTIVE            0
    DHCPLEASEQUERYDONE         0
```

show dhcp server statistics

```
user@host> show dhcp server statistics verbose
Packets dropped:
    Total                      0

Messages received:
    BOOTREQUEST                238
    DHCPDECLINE                0
    DHCPDISCOVER               1
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                237
    DHCPRENEW                  236
    DHCPREBIND                 0

Messages sent:
    BOOTREPLY                  20
    DHCPOFFER                  10
    DHCPACK                    10
    DHCPNAK                         0
    DHCPFORCERENEW                  0
```

show dhcpv6 relay binding

IN THIS SECTION

Syntax

```
show dhcpv6 relay binding
<address>
<brief>
<detail>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<summary>
```

Description

Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address (Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

brief (Optional) Display brief information about the active client bindings. This is the default, and produces the same output as `show dhcpv6 relay binding`.

detail (Optional) Display detailed client binding information.

interface *interface-name* (Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.

interfaces-vlan (Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.

interfaces-wildcard (Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name* (Optional) Perform this operation on the specified logical system.

routing-instance *routing-instance-name* (Optional) Perform this operation on the specified routing instance.

summary (Optional) Display a summary of DHCPv6 client information.

Required Privilege Level

view

Output Fields

Table 57 on page 815 lists the output fields for the `show dhcpv6 relay binding` command. Output fields are listed in the approximate order in which they appear.

Table 57: show dhcpv6 relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* rebinding, *number* releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Client IPv6 Prefix	Prefix of the DHCPv6 client.	brief detail
Client IPv6 Excluded Prefix	IPv6 Prefix of the DHCP client excluded.	detail
Client DUID	DHCP for IPv6 Unique Identifier (DUID) of the client.	brief detail
Client IPv6 Address	IPv6 address assigned to the subscriber.	detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail

Table 57: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the DHCPv6 relay address binding table on the DHCPv6 client: • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RECONFIGURE—Client is broadcasting a request to reconfigure the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCPv6 server. • SELECTING—Client is receiving offers from DHCPv6 servers.	brief detail
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client’s IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client’s IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which the client’s IPv6 prefix expires.	detail

Table 57: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lease Start`	Date and time at which the client’s IP address lease started.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Incoming Client Interface`	Client’s incoming interface.	`detail`
`Server Address`	IP address of the DHCPv6 server. Displays `unknown` for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the `Next Hop Server Facing Relay` field.	`detail`
`Next Hop Server Facing Relay`	Next-hop address in the direction of the DHCPv6 server.	`detail`
`Server Interface`	Interface of the DHCPv6 server.	`detail`
`Relay Address`	IP address of the relay.	`detail`
`Client Pool Name`	Address pool that granted the client lease.	`detail`
`Client ID Length`	Length of client ID.	All levels
`Client Id`	Client ID.	All levels
`Generated Circuit ID`	Circuit ID generated by the DHCPv6 Interface-ID option (option 18)	`detail`
`Generated Remote ID Enterprise Number`	The Juniper Networks IANA private enterprise number	`detail`

Table 57: show dhcpv6 relay binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Generated Remote ID`	Remote ID generated by the DHCPv6 Remote-ID option (option 37)	`detail`
`Dual Stack Group`	Name of the dual-stack group for the DHCPv6 binding.	`detail`
`Dual Stack Peer Address`	Address of the dual-stack DHCPv4 peer.	`detail`

Sample Output

show dhcpv6 relay binding

```
user@host> show dhcpv6 relay binding
Prefix                   Session Id  Expires  State    Interface    Client DUID
2001:db8:3c4d:15::/64    1           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:db8:3c4d:16::/64    2           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:db8:3c4d:17::/64    3           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:db8:3c4d:18::/64    4           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:db8:3c4d:19::/64    5           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:db8:3c4d:20::/64    6           83720    BOUND    ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06
```

show dhcpv6 relay binding (Address)

```
user@host> show dhcp6 relay binding 2001:db8:1111:2222::/64 detail
Session Id:  1
     Client IPv6 Prefix:                   2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Incoming Client Interface:             ge-1/0/0.0
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail (Client ID)

```
user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001 detail
Session Id:  1
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Lease time violated:                   yes
     Incoming Client Interface:             ge-1/0/0.0
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
```

show dhcpv6 relay binding detail

```
user@host> show dhcpv6 relay binding detail
Session Id:  1
     Client IPv6 Prefix:                    2001:db8:3c4d:15::/64
     Client DUID:                           LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
     State:                                 BOUND(RELAY_STATE_BOUND)
     Lease Expires:                         2011-05-25 07:12:09 PDT
     Lease Expires in:                      77115 seconds
     Preferred Lease Expires:               2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:            600 seconds
     Lease Start:                           2011-05-24 07:12:09 PDT
     Lease time violated:                   yes
     Incoming Client Interface:             ge-1/0/0.0
     Server Address:                        2001:db8:aaaa:bbbb::1
     Server Interface:                      none
     Relay Address:                         2001:db8:1111:2222::
     Client Pool Name:                      pool-25
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x4bfa26af/0x00109400/0x0001
     Generated Remote ID Enterprise Number: 1411
     Generated Remote ID:                   host:ge-1/0/0:100
```

show dhcpv6 relay binding detail (Dual-Stack)

```
user@host> show dhcpv6 relay binding detail
Session Id:  2
     Client IPv6 Prefix:                    2001:db8:ffff:0:4::/64
     Client IPv6 Address:                   2001:db8:3000:8003::1/128
     Client DUID:                           LL0x1-00:00:64:01:01:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                         2016-10-17 07:39:25 PDT
     Lease Expires in:                      3450 seconds
     Lease Start:                           2016-10-17 06:39:25 PDT
     Last Packet Received:                  2016-10-17 06:39:25 PDT
     Incoming Client Interface:             ae0.3221225472
     Client Interface Svlan Id:             2000
     Client Interface Vlan Id:              1
     Server Ip Address:                     2001:db8:3000::2
     Server Interface:                      none
     Client Profile Name:                  my-dual-stack
     Client Id Length:                     10
     Client Id:                            /0x00030001/0x00006401/0x0102
     Dual Stack Group:                     group1
     Dual Stack Peer Address:              192.0.2.4
```

show dhcpv6 relay binding detail (Multi-Relay Topology)

```
user@host > show dhcpv6 relay binding detail
Session Id:  13
     Client IPv6 Prefix:                   2001:db8:3000:0:8001::5/128
     Client DUID:                          LL0x1-00:00:65:03:01:02
     State:                                BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                        2011-11-21 06:14:50 PST
     Lease Expires in:                     293 seconds
     Preferred Lease Expires:              2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:           600 seconds
     Lease Start:                          2011-11-21 06:09:50 PST
     Incoming Client Interface:            ge-1/0/0.0
     Server Address:                       unknown
     Next Hop Server Facing Relay:         2001:db8:4000::2
     Server Interface:                     none
     Client Id Length:                     10
     Client Id:                            /0x00030001/0x00006503/0x0102
```

show dhcpv6 relay binding (Session ID)

```
user@host> show dhcpv6 relay binding 41
Prefix                  Session Id   Expires   State   Interface    Client DUID
2001:db8:3c4d:15::/64   41           78837     BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
```

show dhcpv6 relay binding (Subscriber with Multiple Addresses)

```
user@host> show dhcpv6 relay binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
2001:db8:1001::1:24/128         23          593      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
```

```
2001:db8:1001::1:1c/128          23          393      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:1001::1:14/128          23          193      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::300/120           23          293      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::200/120           23          193      BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
2001:db8:3001::100/120           23          93       BOUND    ge-9/0/9.0
LL_TIME0x1-0x55306754-00:10:94:00:00:02
```

When DHCPv6 relay binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 relay binding detail
Session Id:  6
     Hardware Address:                       00:10:94:00:00:01
     Client IPv6 Address:                    7001:2:3::d/128
     Lease Expires:                          2017-12-11 07:45:27 IST
     Lease Expires in:                       9999952 seconds
     Preferred Lease Expires:                2017-12-11 07:45:27 IST
     Preferred Lease Expires in:             9999952 seconds
     Client IPv6 Prefix:                     7001::1000:0:0:0/68
     Client IPv6 Excluded Prefix:            7001::1fff:ffff:ffff:ff00/120
     Lease Expires:                          2017-12-11 07:45:27 IST
     Lease Expires in:                       9999952 seconds
     Preferred Lease Expires:                2017-12-11 07:45:27 IST
     Preferred Lease Expires in:             9999952 seconds
     Client DUID:                            LL_TIME0x1-0x599553b0-00:10:94:00:00:01
     State:                                  BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Start:                            2017-08-17 13:58:33 IST
     Last Packet Received:                   2017-08-17 13:58:48 IST
     Incoming Client Interface:              ge-0/0/0.100
     Client Interface Vlan Id:               100
     Server Ip Address:                      7002::1
     Server Interface:                       none
     Client Id Length:                       14
     Client Id:                              /0x00010001/0x599553b0/0x00109400/0x0001
     Generated Circuit ID:                   ge-0/0/0:100
```

show dhcpv6 relay binding detail (Subscriber with Multiple Addresses)

```
user@host> **show dhcpv6 relay binding detail**
Session Id:  3
     Client IPv6 Address:                   2001:db8:1001::1:2/128
     Lease Expires:                         2015-05-15 02:34:51 PDT
     Lease Expires in:                      24 seconds
     Preferred Lease Expires:               2015-05-15 02:34:51 PDT
     Preferred Lease Expires in:            24 seconds
     Client IPv6 Address:                   2001:db8:1001::1:12/128
     Lease Expires:                         2015-05-15 02:41:31 PDT
     Lease Expires in:                      424 seconds
     Preferred Lease Expires:               2015-05-15 02:41:31 PDT
     Preferred Lease Expires in:            424 seconds
     Client IPv6 Address:                   2001:db8:1001::1:a/128
     Lease Expires:                         2015-05-15 02:38:11 PDT
     Lease Expires in:                      224 seconds
     Preferred Lease Expires:               2015-05-15 02:38:11 PDT
     Preferred Lease Expires in:            224 seconds
     Client IPv6 Prefix:                    2001:db8:3001::/120
     Lease Expires:                         2015-05-15 02:34:51 PDT
     Lease Expires in:                      24 seconds
     Preferred Lease Expires:               2015-05-15 02:34:51 PDT
     Preferred Lease Expires in:            24 seconds
     Client IPv6 Prefix:                    2001:db8:3001::200/120
     Lease Expires:                         2015-05-15 02:38:11 PDT
     Lease Expires in:                      224 seconds
     Preferred Lease Expires:               2015-05-15 02:38:11 PDT
     Preferred Lease Expires in:            224 seconds
     Client IPv6 Prefix:                    2001:db8:3001::100/120
     Lease Expires:                         2015-05-15 02:36:31 PDT
     Lease Expires in:                      124 seconds
     Preferred Lease Expires:               2015-05-15 02:36:31 PDT
     Preferred Lease Expires in:            124 seconds
     Client DUID:                           LL_TIME0x1-0x55554c6e-00:10:94:00:00:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Start:                           2015-05-15 02:34:21 PDT
     Last Packet Received:                  2015-05-15 02:34:22 PDT
     Incoming Client Interface:             ge-9/0/9.0
     Client Interface Vlan Id:              111
     Demux Interface:                       demux0.3221225475
     Server Ip Address:                     2001:db8:5001::1
     Server Interface:                      none
     Client Profile Name:                   DHCP-IPDEMUX-PROF
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x55554c6e/0x00109400/0x0002
     Generated Circuit ID:                  ge-9/0/9:111
     Generated Remote ID Enterprise Number: 1411
     Generated Remote ID:                   ge-9/0/9:111
```

show dhcpv6 relay binding (Interfaces VLAN)

```
user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:DB8::/32        11        87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:DB8::/32        30       79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding (Interfaces Wildcard)

```
user@host> show dhcpv6 relay binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:DB8::/32        22       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

```
2001:DB8:C9::/32     24       79681     BOUND     ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 relay binding summary

```
user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 relay lockout-entries

IN THIS SECTION

Syntax

```
show dhcpv6 relay lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 relay agent lockout database.

Options

all	Display all client entries in the lockout database.
index *index*	Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 58 on page 826 lists the output fields for the `show dhcpv6 relay lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 58: show dhcpv6 relay lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`

Table 58: show dhcpv6 relay lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcpv6 relay lockout-entries (All Entries)

```
user@host> show dhcpv6 relay lockout-entries all
Index     Key              State    Expires(s)    Elapsed(s)     Count
```

```
1     00:00:5E:00:53:00      Lockout     30          5200            2
2     00:00:5E:00:53:11      Grace      120          780             2
3     00:00:5E:00:53:22      Lockout    180          2300            1
```

show dhcpv6 relay lockout-entries (Specific Entry)

```
user@host> show dhcpv6 relay lockout-entries index 2
     Index:                                 2
     Key:                                   default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                 Lockout
     Expires:                               2018-03-29 19:06:17 IST
     Expires in:                            87
     Lockout count:                         1
     Next lockout time:                     200
     Min lockout time:                      100
     Lockout reason:                        181
```

show dhcpv6 relay statistics

IN THIS SECTION

Syntax

```
show dhcpv6 relay statistics

<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.

Options

logical-system *logical-system-name*
: (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

routing-instance *routing-instance-name*
: (Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 59 on page 830 lists the output fields for the `show dhcpv6 relay statistics` command. Output fields are listed in the approximate order in which they appear.

Table 59: show dhcpv6 relay statistics Output Fields

Field Name	Field Description
DHCPv6 Packets dropped	Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • Total—Total number of packets discarded by the DHCPV6 relay agent application. • Bad options—Number of packets discarded because invalid options were specified. • Bad send—Number of packets that the extended DHCP relay application could not send. • Bad src address—Number of packets discarded because the family type was not AF_INET6. • Client MAC validation—Number of packets discarded because validation of the client MAC address failed. • No client id—Number of packets discarded because they could not be matched to a client. • Lease Time Violation—Number of packets discarded because of a lease time violation • No safd—Number of packets discarded because they arrived on an unconfigured interface. • Short packet—Number of packets discarded because they were too short. • Relay hop count—Number of packets discarded because the hop count in the packet exceeded 32.

Table 59: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages received`	Number of DHCPv6 messages received. • `DHCPv6_DECLINE`—Number of DHCPv6 PDUs of type DECLINE received • `DHCPV6_SOLICIT`—Number of DHCPv6 PDUs of type SOLICIT received • `DHCPV6_INFORMATION_REQUEST`—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received • `DHCPV6_RELEASE`—Number of DHCPv6 PDUs of type RELEASE received • `DHCPV6_REQUEST`—Number of DHCPv6 PDUs of type REQUEST received • `DHCPV6_CONFIRM`—Number of DHCPv6 PDUs of type CONFIRM received `DHCPV6_RENEW`—Number of DHCPv6 PDUs of type RENEW received `DHCPV6_REBIND`—Number of DHCPv6 PDUs of type REBIND received `DHCPV6_RELAY_REPL`—Number of DHCPv6 PDUs of type RELAY-REPL received • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs received. • `DHCPV6_LEASEQUERY_REPLY`—Number of DHCPv6 replies received from the DHCPv6 sever • `DHCPV6_LEASEQUERY_DATA`—xxxx • `DHCPV6_LEASEQUERY_DONE`—The leasequery is complete
`Messages sent`	Number of DHCPv6 messages sent. • `DHCPV6_ADVERTISE`—Number of DHCPv6 ADVERTISE PDUs transmitted • `DHCP_REPLY`—Number of DHCPv6 REPLY PDUs transmitted • `DHCP_RECONFIGURE`—Number of DHCPv6 RECONFIGURE PDUs transmitted • `DHCP_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs transmitted • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs transmitted. • `DHCP6_LEASEQUERY`—Number of DHCP leasequery messages transmitted

Table 59: show dhcpv6 relay statistics Output Fields *(Continued)*

Field Name	Field Description
Packets forwarded	Number of packets forwarded by the extended DHCPv6 relay agent application. • FWD REQUEST—Number of DHCPv6 REQUEST packets forwarded • FWD REPLY—Number of DHCPv6 REPLY packets forwarded
External Server Response	State of the external DHCP server responsiveness.

Sample Output

show dhcpv6 relay statistics

```
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                      2
    Lease Time Violation       1
    Client MAC validation      1

Messages received:
    DHCPV6_DECLINE             0
    DHCPV6_SOLICIT             10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             10
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
    DHCPV6_RELAY_FORW          0
    DHCPV6_LEASEQUERY_REPLY    0
    DHCPV6_LEASEQUERY_DATA     0
    DHCPV6_LEASEQUERY_DONE     0

Messages sent:
    DHCPV6_ADVERTISE            0
    DHCPV6_REPLY                0
    DHCPV6_RECONFIGURE          0
    DHCPV6_RELAY_REPL           0
    DHCPV6_LEASEQUERY           0

Packets forwarded:
    Total                       4
    FWD REQUEST                 2
    FWD REPLY                   2

External Server Response:
    State                      Responding
```

show dhcpv6 server binding

IN THIS SECTION

Syntax

```
show dhcpv6 server binding
<address>
<brief | detail | summary>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
```

```
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.

Options

address (Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

brief | detail | summary (Optional) Display the specified level of output about active client bindings. The default is `brief`, which produces the same output as `show dhcpv6 server binding`.

interface *interface-name* (Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan (Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.

interfaces-wildcard (Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).

logical-system *logical-system-name* (Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name* (Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.

Required Privilege Level

view

Output Fields

"show dhcpv6 server binding" on page 833 lists the output fields for the `show dhcpv6 server binding` command. Output fields are listed in the approximate order in which they appear.

Table 60: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
number `clients`, (*number* `init`, *number* `bound`, *number* `selecting`, *number* `requesting`, *number* `renewing`, *number* `releasing`)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	`summary`
`Prefix`	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	`brief` `detail`
`Session Id`	Session ID of the subscriber session.	`brief` `detail`
`Expires`	Number of seconds in which lease expires.	`brief` `detail`

Table 60: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	State of the address binding table on the extended DHCPv6 local server: • `BOUND`—Client has active IP address lease. • `INIT`—Initial state. • `RECONFIGURE`—Server has sent reconfigure message to client. • `RELEASE`—Client is releasing IP address lease. • `RENEWING`—Client sending request to renew IP address lease. • `REQUESTING`—Client requesting a DHCPv6 server. • `SELECTING`—Client receiving offers from DHCPv6 servers.	`brief` `detail`
`Interface`	Interface on which the DHCPv6 request was received.	`brief`
`Client IPv6 Address`	Client's IPv6 address.	`detail`
`Client IPv6 Prefix`	Client's IPv6 prefix.	`detail`
`Client IPv6 Excluded Prefix`	IPv6 Prefix of the DHCP client excluded.	`detail`
`Client DUID`	Client's DHCP Unique Identifier (DUID).	`brief` `detail`
`Lease expires`	Date and time at which the client's IP address lease expires.	`detail`
`Lease expires in`	Number of seconds in which lease expires.	`detail`

Table 60: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Preferred Lease Expires`	Date and UTC time at which the client's IPv6 prefix expires.	`detail`
`Preferred Lease Expires in`	Number of seconds at which client's IPv6 prefix expires.	`detail`
`Lease Start`	Date and time at which the client's address lease was obtained.	`detail`
`Lease time violated`	Lease time violation has occurred.	`detail`
`Incoming Client Interface`	Client's incoming interface.	`detail`
`Server IP Address`	IP address of DHCPv6 server.	`detail`
`Server Interface`	Interface of DHCPv6 server.	`detail`
`Client Pool Name`	Address pool used to assign IPv6 address.	`detail`
`Client Prefix Pool Name`	Address pool used to assign IPv6 prefix.	`detail`
`Client Id length`	Length of the DHCPv6 client ID, in bytes.	`detail`
`Client Id`	ID of the DHCPv6 client.	`detail`
`Server Id`	DHCP unique identifier (DUID) for the DHCPv6 server.	`detail`
`Client Interface Svlan Id`	S-VLAN ID of the client's incoming interface.	`detail`

Table 60: show dhcpv6 server binding Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Client Interface Vlan Id`	VLAN ID of the client’s incoming interface.	`detail`
`Dual Stack Group`	DHCPv6 server profile name.	`detail`
`Dual Stack Peer Address`	DHCPv6 Peer IP address.	`detail`

Sample Output

show dhcpv6 server binding

```
user@host> show dhcpv6 server binding
Prefix                  Session Id  Expires  State    Interface    Client DUID
1000::3/128             9           86303    BOUND    up:green-arrow:ge-0/3/5.2
LL0x1-00:00:64:03:01:02
```

show dhcpv6 server binding detail

```
user@host> show dhcpv6 server binding detail
Session Id:  2
     Client IPv6 Prefix:                    2001:db8:ffff:0:4::/64
     Client IPv6 Address:                   2001:db8:0:8003::1/128
     Client DUID:                           LL0x1-00:00:64:01:01:02
     State:                                 BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Expires:                         2016-11-07 08:30:39 PST
     Lease Expires in:                      43706 seconds
     Preferred Lease Expires:               2016-11-07 08:30:39 PST
     Preferred Lease Expires in:            43706 seconds
     Lease Start:                           2016-11-04 11:00:37 PDT
     Last Packet Received:                  2016-11-06 09:00:39 PST
     Incoming Client Interface:             ae0.3221225472
     Client Interface Svlan Id:             2000
     Client Interface Vlan Id:              1
     Server Ip Address:                     2001:db8::2
     Server Interface:                      none
     Client Profile Name:                   my-dual-stack
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006401/0x0102
     Dual Stack Group:                      my-dual-stack
     Dual Stack Peer Address:               192.0.2.10
```

command-name

When DHCPv6 binding is configured with prefix exclude option, we get the following output:

```
user@host> show dhcpv6 server binding detail
Session Id:  5
     Client IPv6 Address:                   2001:db8:2:3::d/128
     Lease Expires:                         2017-12-11 07:45:15 IST
     Lease Expires in:                      9999995 seconds
     Preferred Lease Expires:               2017-12-11 07:45:15 IST
     Preferred Lease Expires in:            9999995 seconds
     Client IPv6 Prefix:                    2001:db8::1000:0:0/68
       Client IPv6 Excluded Prefix:           2001:db8::1fff:ffff:ff00/120
     Lease Expires:                         2017-12-11 07:45:15 IST
     Lease Expires in:                      9999995 seconds
     Preferred Lease Expires:               2017-12-11 07:45:15 IST
     Preferred Lease Expires in:            9999995 seconds
     Client DUID:                           LL_TIME0x1-0x599553b0-00:10:94:00:00:01
     State:                                 BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)
     Lease Start:                           2017-08-17 13:58:32 IST
     Last Packet Received:                  2017-08-17 13:58:36 IST
     Incoming Client Interface:             ge-0/0/0.0
     Client Interface Vlan Id:              100
     Client Pool Name:                      ia_na_pool
     Client Prefix Pool Name:               prefix_delegate_pool
     Client Id Length:                      14
     Client Id:                             /0x00010001/0x599553b0/0x00109400/0x0001
     Relay Id Length:                       31
     Relay Id:                              /0x00020000/0x05830130/0x303a3035/0x3a38363a
     Relay Id:                              /0x34343a65/0x323a6330/0x00000000/0x000000
```

show dhcpv6 server binding interface

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix                  Session Id  Expires  State    Interface     Client DUID
2001:db8:1111:2222::/64 1           86055    BOUND    ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01
```

show dhcpv6 server binding interface detail

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id:  7
     Client IPv6 Prefix:                  2001:db8:1111:2222::/64
     Client DUID:                         LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                               BOUND(bound)
     Lease Expires:                       2009-07-21 10:41:15 PDT
     Lease Expires in:                    86136 seconds
     Preferred Lease Expires:             2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:          600 seconds
     Lease Start:                         2009-07-20 10:41:15 PDT
     Incoming Client Interface:           ge-1/0/0.0
     Server Ip Address:                   0.0.0.0
     Server Interface:                    none
     Client Id Length:                    14
     Client Id:                           /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (IPv6 Prefix)

```
user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005 detail
Session Id:  7
     Client IPv6 Prefix:                  2001:db8:1111:2222::/64
     Client DUID:                         LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
     State:                               BOUND(bound)
     Lease Expires:                       2009-07-21 10:41:15 PDT
     Lease Expires in:                    86136 seconds
     Preferred Lease Expires:             2012-07-24 00:18:14 UTC
     Preferred Lease Expires in:          600 seconds
     Lease Start:                         2009-07-20 10:41:15 PDT
     Incoming Client Interface:           ge-1/0/0.0
     Server Ip Address:                   0.0.0.0
     Server Interface:                       none
     Client Id Length:                       14
     Client Id:                              /0x00010001/0x02e159c0/0x00109400/0x0002
```

show dhcpv6 server binding (Session ID)

```
user@host> show dhcpv6 server binding 8
Prefix           Session Id  Expires  State    Interface    Client DUID
2001:db8::/32    8           86235    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
```

show dhcpv6 server binding (Interfaces VLAN)

```
user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix            Session Id  Expires  State    Interface             Client DUID
2001:db8::/32     11          87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     12       87583    BOUND    ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding demux0
Prefix            Session Id  Expires  State    Interface           Client DUID
2001:db8::/32     30          79681    BOUND    demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     31       79681    BOUND    demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     32       79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding ge-1/3/*
Prefix            Session Id  Expires  State    Interface      Client DUID
2001:db8::/32     22          79681    BOUND    ge-1/3/0.110
```

```
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:19::/32     33       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:db8:C9::/32     24       79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 server lockout-entries

IN THIS SECTION

Syntax

```
show dhcpv6 server lockout-entries (all | index index)
```

Description

Display information about all client entries or detailed information about a specific client entry in the DHCPv6 local server lockout database.

Options

all — Display all client entries in the lockout database.

index *index* — Display detailed information for the specified client.

Required Privilege Level

view

Output Fields

Table 61 on page 843 lists the output fields for the `show dhcpv6 server lockout-entries` command. Output fields are listed in the approximate order in which they appear.

Table 61: show dhcpv6 server lockout-entries Output Fields

Field Name	Field Description	Level of Output
`Index`	Number identifying a specific entry in the lockout database.	`all` and `index`
`Key`	DUID identifying the client in the lockout database.	`all` and `index`

Table 61: show dhcpv6 server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`State`	Type of lockout period for the entry: • `Grace`—A previously locked out client enters the grace period when the lockout expires. If the client attempts to establish a session within in this period, the next lockout time is increased. If the grace time passes without a log in, the entry is removed from the lockout database. • `Lockout`—Client is currently locked out; attempts to establish a session are rejected.	`all` and `index`
`Expires (s)`	Number of seconds until the current lockout period expires.	`all` only
`Elapsed (s)`	Number of seconds since the current lockout or grace timer started.	`all` only
`Count`	Number of consecutive times the client has been locked out.	`all` only
`Expires`	Date and time when the current lockout period ends.	`index` only
`Expires in`	Number of seconds until the current period expires.	`index` only
`Lockout count`	Number of consecutive times client has been locked out.	`index` only
`Next lockout time`	Duration of the next lockout period for this client.	`index` only
`Min lockout time`	Minimum duration for a lockout period; the initial lockout time.	`index` only

Table 61: show dhcpv6 server lockout-entries Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Lockout reason`	Reason for the current lockout. The possible values are internal jdhcpd error codes. These values are provided for debugging by Juniper Networks technical support.	`index` only

Sample Output

show dhcpv6 server lockout-entries (All Entries)

```
user@host> show dhcpv6 server lockout-entries all
Index     Key                State    Expires(s)     Elapsed(s)      Count
1     00:00:5E:00:53:00      Lockout     30            5200            2
2     00:00:5E:00:53:11      Grace      120            780             2
3     00:00:5E:00:53:22      Lockout    180            2300            1
```

show dhcpv6 server lockout-entries (Specific Entry)

```
user@host> show dhcpv6 server lockout-entries index 2
     Index:                                 2
     Key:                                   default/00 01 00 01 5a bc e1 7b 00 10 94 00 00 06/
     State:                                 Lockout
     Expires:                               2018-03-29 19:06:17 IST
     Expires in:                            87
     Lockout count:                         1
     Next lockout time:                     200
     Min lockout time:                      100
     Lockout reason:                        181
```

show dhcpv6 server statistics

IN THIS SECTION

Syntax

```
show dhcpv6 server statistics

<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Description

Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

logical-system *logical-system-name*

(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.

routing-instance ***routing-instance-name***	(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Output Fields

Table 62 on page 848 lists the output fields for the `show dhcpv6 server statistics` command. Output fields are listed in the approximate order in which they appear.

Table 62: show dhcpv6 server statistics Output Fields

Field Name	Field Description
Packets dropped	Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • Total—Total number of packets discarded by the extended DHCPv6 local server • Strict Reconfigure—Number of solicit messages discarded because the client does not support reconfiguration • Bad hardware address—Number of packets discarded because an invalid hardware address was specified • Bad opcode—Number of packets discarded because an invalid operation code was specified • Bad options—Number of packets discarded because invalid options were specified • Client MAC validation—Number of packets discarded because validation of the client MAC address failed. • Invalid server address—Number of packets discarded because an invalid server address was specified • Lease Time Violation—Number of packets discarded because of a lease time violation • No available addresses—Number of packets discarded because there were no addresses available for assignment • No interface match—Number of packets discarded because they did not belong to a configured interface • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance • No valid local address—Number of packets discarded because there was no valid local address • Packet too short—Number of packets discarded because they were too short • Read error—Number of packets discarded because of a system read error • Send error—Number of packets that the extended DHCPv6 local server could not send

Table 62: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
`Advertise Delay`	Number of DHCP advertise messages delayed. • `DELAYED`—Number of DHCPv6 advertise packets that have been sent after being delayed. • `INPROGRESS`—Number of DHCPv6 advertise packets that are in the delay queue. • `TOTAL`—Total number of delayed DHCPv6 advertise messages; sum of `DELAYED` and `INPROGRESS`.
`Messages received`	Number of DHCPv6 messages received. • `DHCPV6_CONFIRM`—Number of DHCPv6 CONFIRM PDUs received. • `DHCPV6_DECLINE`—Number of DHCPv6 DECLINE PDUs received. • `DHCPV6_INFORMATION_REQUEST`—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • `DHCPV6_REBIND`—Number of DHCPv6 REBIND PDUs received. • `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs received. • `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs received. • `DHCPV6_RELEASE`—Number of DHCPv6 RELEASE PDUs received. • `DHCPV6_RENEW`—Number of DHCPv6 RENEW PDUs received. • `DHCPV6_REQUEST`—Number of DHCPv6 REQUEST PDUs received. • `DHCPV6_SOLICIT`—Number of DHCPv6 SOLICIT PDUs received. • `DHCPV6_LEASEQUERY`—Number of DHCPv6 leasequery messages received.

Table 62: show dhcpv6 server statistics Output Fields *(Continued)*

Field Name	Field Description
`Messages sent`	Number of DHCPv6 messages sent. - `DHCPV6_ADVERTISE`—Number of DHCPv6 ADVERTISE PDUs transmitted. - `DHCPV6_REPLY`—Number of DHCPv6 ADVERTISE PDUs transmitted. - `DHCPV6_LOGICAL_NAK`—Number of logical NAK messages sent, signifying T1 and T2 timers with values of zero; subset of `DHCPV6_REPLY` counter. (Displays only at `verbose` level. - `DHC6_RECONFIGURE`—Number of DHCPv6 RECONFIGURE PDUs transmitted. - `DHCPV6_RELAY_REPL`—Number of DHCPv6 RELAY-REPL PDUs transmitted. - `DHCPV6_RELAY_FORW`—Number of DHCPv6 RELAY-FORW PDUs transmitted. - `DHCPV6_LEASEQUERY_REPLY`—Number of DHCPv6 leasequery replies transmitted to the DHCPv6 relay agent. - `DHCPV6_LEASEQUERY_DATA`—Number of DHCPv6 LEASEQUERY-DATA packets transmitted. - `DHCPV6_LEASEQUERY_DONE`—Number of DHCPv6 LEASEQUERY-DONE packets sent.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
Dhcpv6 Packets dropped:
    Total               0

Advertise Delay:
    DELAYED                    0
    INPROGRESS                 0
    TOTAL                      0

Messages received:
    DHCPV6_DECLINE             0
    DHCPV6_SOLICIT             1
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE             0
    DHCPV6_REQUEST             1
    DHCPV6_CONFIRM             0
    DHCPV6_RENEW               0
    DHCPV6_REBIND              0
    DHCPV6_RELAY_FORW          0
    DHCPV6_LEASEQUERY          0
    DHCPV6_ACTIVELEASEQUERY    0

Messages sent:
    DHCPV6_ADVERTISE           1
    DHCPV6_REPLY               1
    DHCPV6_RECONFIGURE         0
    DHCPV6_RELAY_REPL          0
    DHCPV6_LEASEQUERY_REPLY    0
    DHCPV6_LEASEQUERY_DATA     0
    DHCPV6_LEASEQUERY_DONE     0
```

show dynamic-profile session

IN THIS SECTION

Syntax

```
show dynamic-profile session
<client-id client-id>
<profile-name profile-name>
<service-id service-id>
```

Description

Display dynamic profile (client or service) information for all subscribers or for subscribers specified by client ID or service session ID. You can filter the output by also specifying a dynamic profile.

NOTE:

- The output does not display the variable stanzas defined in the dynamic profile configuration.
- The variables in the profile configuration are replaced with subscriber specific values.
- If the conditional variable in the dynamic profile is evaluated as NULL, the subscriber value for the variable is displayed as `NONE` in the command output.
- The variable is also displayed as `NONE` when the variable (any variable and not necessarily conditional) in the dynamic profile has no value associated with it.
- The format in which the configuration is displayed looks similar, but not exactly the same as the format of the `show configuration dynamic-profiles` command.

Options

client-id *client-id* — Display dynamic profile information for subscribers associated with the specified client.

profile-name *profile-name* — (Optional) Display dynamic profile information for the specified subscriber or service profile.

service-id *service-id* Display dynamic profile information for subscribers associated with the specified service session.

Required Privilege Level

view

Output Fields

This command displays the dynamic client or service profile configuration for each subscriber.

Sample Output

show dynamic-profile session client-id (Client ID)

```
user@host>show dynamic-profile session client-id 20
pppoe {
    interfaces {
        pp0 {
            unit 1073741831 {
                ppp-options {
                    chap;
                    pap;
                }
                pppoe-options {
                    underlying-interface ge-2/0/0.0;
                    server;
                }
                family {
                    inet {
                        unnumbered-address lo0.0;
                    }
                }
            }
        }
    }
    class-of-service {
        traffic-control-profiles {
            tcp1 {
                scheduler-map smap1_UID1024;
                shaping-rate 100m;
            }
        }
        interfaces {
            pp0 {
                unit 1073741831 {
                    output-traffic-control-profile tcp1;
                }
            }
        }
        scheduler-maps {
            smap1_UID1024 {
                forwarding-class best-effort scheduler sch1_UID1023;
            }
        }
        schedulers {
            sch1_UID1023 {
                transmit-rate percent 40;
                buffer-size percent 40;
                priority low;
            }
        }
    }
}
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session client-id profile-name (Client ID and Dynamic Profile)

```
user@host>show dynamic-profile session client-id 20 profile-name cos-service
cos-service {
    class-of-service {
        scheduler-maps {
            smap2_UID1029 {
                forwarding-class assured-forwarding scheduler sch2_UID1028;
            }
        }
        schedulers {
            sch2_UID1028 {
                transmit-rate percent 60;
                buffer-size percent 60;
                priority high;
            }
        }
    }
}
```

show dynamic-profile session service-id (Service Session)

```
user@host>show dynamic-profile session service-id 21
filter-service {
    interfaces {
        pp0 {
            unit 1073741831 {
                family {
                    inet {
                        filter {
                            input input-filter_UID1026 precedence 50;
                            output output-filter_UID1027 precedence 50;
                        }
                    }
                }
            }
        }
    }
    firewall {
        family {
            inet {
                filter input-filter_UID1026 {
                    interface-specific;
                    term t1 {
                        then {
                            policer policer1_UID1025;
                            service-accounting;
                        }
                    }
                    term rest {
                        then accept;
                    }
                }
                filter output-filter_UID1027 {
                    interface-specific;
                    term rest {
                        then accept;
                    }
                }
            }
        }
        policer policer1_UID1025 {
            if-exceeding {
                bandwidth-limit 1m;
                burst-size-limit 15k;
            }
            then discard;
        }
    }
}
```

show interfaces extensive demux0

IN THIS SECTION

Syntax

```
show interfaces extensive demux0.logical-interface-number
```

Description

Display status information about the specified demux interface.

Options

none	Display standard information about the specified demux interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information.

snmp-index ***snmp-index***	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 63 on page 859 lists the output fields for the `show interfaces` *`ifl-name`* command. Output fields are listed in the approximate order in which they appear.

Table 63: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
`Link`	Status of the physical link (Up or Down).	`terse` none
`Targeting summary`	Status of LT links that are configured with targeted distribution (primary or backup)	`extensive` none
Logical Interface		
`Logical interface`	Name of the logical interface.	`brief detail extensive` none
`Index`	Index number of the logical interface, which reflects its initialization sequence.	`detail extensive` none
`SNMP ifIndex`	SNMP interface index number for the logical interface.	`detail extensive` none

Table 63: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under Common Output Fields Description.	`brief` `detail` `extensive` none
`Encapsulation`	Type of encapsulation configured on the logical interface.	`brief` `extensive` none
`Demux`	Specific IP demultiplexing (demux) values: • Underlying interface—The underlying interface that the demux interface uses. • Index—Index number of the logical interface. • Prefix—inet family prefix. • Family—Protocol family configured on the logical interface. • Source prefixes, total—Total number of source prefixes for the underlying interface.	`detail` `extensive` none

Sample Output

show interfaces extensive *demux0.logical-interface-number*

```
user@host> show interfaces demux0.3221225544
    Logical interface demux0.3221225544 (Index 536871034) (SNMP ifIndex 200000122) (Generation 
76)
    Flags: Up Encapsulation: ENET2
    Interface set: iflset100100
    Demux:
      Underlying interface: demux0.3221225527 (Index 536871017)
      Family Inet Source prefixes, total 1
          Prefix: 192.168.0.5/32
    Link:
      ge-1/0/9.32767
      ge-1/1/8.32767
      ge-1/1/9.32767
      ge-4/2/0.32767
    Targeting summary:
      ge-1/1/8, primary, Physical link is Up
      ge-4/2/8, backup, Physical link is Up
```

show interfaces interface-set

IN THIS SECTION

Syntax

```
show interfaces interface-set iflset-name
```

Description

Display information about the specified aggregated Ethernet interface set.

Options

interface-set *iflset-name*	Display standard information about the specified aggregated Ethernet interface set.
detail | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 63 on page 859 lists the output fields for the `show interfaces interface-set` command. Output fields are listed in the approximate order in which they appear.

Table 64: show interfaces interface-set Output Fields

Field Name	Field Description
`Interface set`	Name of the interface set or sets. • Interface set index—Index number of the interface set. • Interface set snmp index—SNMP interface index number for the interface set.
`Members`	Status of the physical link (Up or Down).
`Targeting summary`	Status of links that are configured with targeted distribution (primary or backup)

Sample Output

show interfaces interface-set *iflset-name*

```
user@host> show interfaces interface-set ae1-103
  Interface set: ae1-103
  Interface set index: 536870916
  Interface set snmp index: 603979780
  Members:
   demux0.3221225610
   demux0.3221225616
   demux0.3221225617
   demux0.3221225618

  Targeting summary:
    ge-4/2/8, Primary, Physical link is Up
    ge-1/0/9, Backup, Physical link is Up
```

show interfaces targeting

IN THIS SECTION

Syntax

```
show interfaces targeting aex
```

Description

Displays status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.

Options

Required Privilege Level

view

Output Fields

Table 65 on page 864 lists the output fields for the `show interfaces targeting` command. Output fields are listed in the approximate order in which they appear.

Table 65: show interfaces targeting Output Fields

Field Name	Field Description	Level of Output
Aggregated Ethernet Interface		
`Aggregated interface`	Name of the aggregated Ethernet bundle.	All levels

Table 65: show interfaces targeting Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Redundancy mode`	Redundancy mechanism on the interface: `Link Level Redundancy` or `FPC Redundancy`.	All levels
`Total number of distributed interfaces`	Number of distributed links in the bundle.	All levels
Physical Interface		
`Physical interface`	Name of the physical interface and state of the interface.	All levels
`Link status`	Status of the link on the physical interface: `up` or `down`.	
`Number of primary distributions`	Number of subscribers distributed on primary links.	All levels
`Number of backup distributions`	Number of subscribers distributed on backup links.	All levels

Sample Output

show interfaces targeting ae1

```
user@host> show interfaces targeting ae1
Aggregated interface: ae1
Redundancy mode: Link Level Redundancy
Total number of distributed interfaces: 3
Physical interface: ge-1/0/0, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-1/1/0, Link status: Up
```

```
Number of primary distributions: 200
Number of backup distributions: 199
Physical interface: ge-2/0/7, Link status: Up
Number of primary distributions: 200
Number of backup distributions: 200
Physical interface: ge-2/0/8, Link status: Up
Number of primary distributions: 199
Number of backup distributions: 200
```

show interfaces terse

IN THIS SECTION

Syntax

```
show interfaces terse
```

Description

Display summary information about interfaces.

Options

This command has no options.

Additional Information

Interfaces are always displayed in numerical order, from the lowest to the highest FPC slot number. Within that slot, the lowest PIC slot is shown first. On an individual PIC, the lowest port number is always first.

Required Privilege Level

view

Output Fields

Table 66 on page 867 lists the output fields for the `show interfaces terse` command. Output fields are listed in the approximate order in which they appear.

Table 66: show interfaces terse Output Fields

Field Name	Field Description
Interface	Interface name.
Admin	Whether the interface is turned on (up) or off (down).
Link	Link state: **up** or **down**.
Proto	Protocol family configured on the logical interface. A logical interface on a router that supports Ethernet OAM always shows the multiservice protocol.
Local	Local IP address of the logical interface.

Table 66: show interfaces terse Output Fields *(Continued)*

Field Name	Field Description
Remote	Remote IP address of the logical interface.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
t1-0/1/0:0              up    up
t1-0/1/0:0.0            up    up   inet     192.168.220.18/30
t1-0/1/0:1              up    up
t1-0/1/0:2              up    up
t1-0/1/0:3              up    up
at-1/0/0                up    up
at-1/0/1                up    up
dsc                     up    up
fxp0                    up    up
fxp0.0                  up    up   inet     192.168.71.249/21
fxp1                    up    up
fxp1.0                  up    up   inet     10.0.0.4/8
                                   tnp      4
gre                     up    up
ipip                    up    up
lo0                     up    up
lo0.0                   up    up   inet     10.0.1.4            --> 0/0
                                            127.0.0.1           --> 0/0
lo0.16385               up    up   inet

lsi                     up    up
mtun                    up    up
```

show interfaces terse (TX Matrix Plus Router)

```
user@host> show interfaces terse

Interface               Admin Link Proto    Local                 Remote
xe-0/0/0                up    up
xe-0/0/1                up    up
xe-0/0/2                up    up
xe-0/0/3                up    up
xe-6/0/0                up    up
xe-6/0/1                up    up
xe-6/0/2                up    up
xe-6/0/3                up    up
xe-6/1/0                up    up
xe-6/1/1                up    up
xe-6/1/2                up    up
xe-6/1/3                up    up
so-0/0/0                up    up
so-0/0/0.0              up    up   inet     10.1.1.1/30
ge-1/3/0.0              up    up   inet     --> 0/0
ge-7/0/0                up    up
ge-7/0/0.0              up    up   inet     10.2.1.1/30
ge-7/0/0.1              up    up   inet     10.2.1.5/30
ge-7/0/0.2              up    up   inet     10.2.1.9/30
ge-7/0/0.3              up    up   inet     10.2.1.13/30
ge-7/0/0.4              up    up   inet     10.2.1.17/30
ge-7/0/0.5              up    up   inet     10.2.1.21/30
...
em0                     up    up
em0.0                   up    up   inet     192.168.178.11/25
gre                     up    up
ipip                    up    up
ixgbe0                  up    up
ixgbe0.0                up    up   inet     10.34.0.4/8
                                            162.0.0.4/2
                                   inet6    fe80::200:ff:fe22:4/64
                                            fec0::a:22:0:4/64
                                   tnp      0x22000004
ixgbe1                  up    up
ixgbe1.0                up    up   inet     10.34.0.4/8
                                            162.0.0.4/2
                                   inet6    fe80::200:1ff:fe22:4/64
```

```
                                                fec0::a:22:0:4/64
                                       tnp      0x22000004
```

show interfaces terse (PTX Series Packet Transport Routers)

```
user@host> show interfaces em0 terse
   Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

Release Information

Command introduced on Release 24.2.

show network-access aaa accounting

IN THIS SECTION

Syntax

```
show network-access aaa accounting
```

Description

Display the state of the RADIUS Acct-On response sent from the RADIUS server.

Required Privilege Level

view

Output Fields

Table 67 on page 871 lists the output fields for the `show network-access aaa accounting` command. Output fields are listed in the approximate order in which they appear.

Table 67: show network-access aaa accounting Output Fields

Field Name	Field Description
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.
`Logical System`	Logical system associated with the access profile.
`Routing Instance`	Routing instance associated with the access profile.
`Acct-On-Response`	Status of the RADIUS Acct-On response. • ACK—ACK response for the Acct-On message is received from the RADIUS server. • ERROR—An error condition has occurred. • NONE— No Acct-On message is sent. • PENDING—Acct-On message is sent to RADIUS server, but no response has been received yet.

Sample Output

show network-access aaa accounting

```
user@host> show network-access aaa accounting
Profile         Logical System   Routing Instance    Acct-On-Response
ppp-profile     default          default             ACK
l2tp-profile    default          l2tp_RI             PENDING
```

show network-access aaa radius-servers

IN THIS SECTION

Syntax

```
show network-access aaa radius-servers
<detail>
```

Description

Display RADIUS server status and information.

Options

detail (Optional) Display detailed level of information.

Required Privilege Level

view

Output Fields

Table 68 on page 873 lists the output fields for the `show network-access aaa radius-servers` command. Output fields are listed in the approximate order in which they appear.

Table 68: show network-access aaa radius-servers Output Fields

Field Name	Field Description	Level of Output
`Profile`	Name of the profile associated with the RADIUS server. A RADIUS server can be associated with more than one profile.	All levels
`Server address`	IPv4 or IPv6 address of the RADIUS server.	All levels
`Authentication port`	RADIUS server authentication port number.	All levels
`Preauthentication port`	RADIUS server preauthentication port number.	All levels
`Accounting port`	RADIUS server accounting port number.	All levels
`Accounting retry`	Number of times the router retransmits RADIUS accounting messages when no response is received from the server.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting timeout`	Period the local router waits to receive a response from a RADIUS accounting server before retransmitting the message.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Status`	RADIUS server status, `UP` (Alive), `UNREACHABLE`, or `DOWN` (DEAD). If status is `DOWN`, the Status field includes the number of seconds configured by the `revert-interval` statement. The router does not send requests to servers in the `DOWN` state, but does send requests to servers with a status of either `UP` or `UNREACHABLE`. This field also displays the status of AAA accounting suspension or resumption, and the status of baselining of accounting statistics if you suspended or resumed accounting operations or initiated the generation of a baseline. This information is applicable only for RADIUS servers that are in the `UP` state. **NOTE**: After requests to a server or set of servers time out after 10 seconds, the status of the servers changes. The following guidelines apply to server status: • For the purpose of marking a server as `Down` (DEAD), the request includes the original request and any retries that are configured. The 10-second timeout period starts after the initial request and all retries have expired without receiving a response from the server. The amount of the timeout period that elapses before the server is marked `Down` is not always exactly 10 seconds, and can vary depending on how frequently subscribers are logging in. When subscribers are continually and rapidly logging in, the server is marked as `Down` at 10 seconds. However, if subscribers are logging in less frequently and at a slower pace, then the server is not marked `Down` until a subsequent subscriber attempts to log in. For example, if the subsequent subscriber logs in a minute after the request and all retries lapse, and the 10-second timeout starts, the actual time until the server is marked `Down` is 50 seconds after the timeout starts (the one minute between subscriber login minus the 10-second timeout). • All servers cannot be marked as `DOWN`; instead, the unresponsive servers are marked as `UNREACHABLE`.	All levels

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	For example, if only one RADIUS server is configured and that server is unresponsive, the server status is marked as `UNREACHABLE` rather than `DOWN`. • If at least one server has a status of `UP`, the status of all unresponsive servers is set to `DOWN` for the remainder of the configured revert-interval setting. • If no server has a status of `UP`, then the status of the unresponsive servers is set to `UNREACHABLE` for the remainder of the revert-interval setting or for 30 seconds, whichever is less. • The status of unresponsive servers is returned to `UP` from `DOWN` or `UNREACHABLE` at the end of the revert-interval setting (or the 30-second interval). • If no requests are sent to a server, the server's status is always `UP`.	
`RADIUS servers`	Details for specific RADIUS server, identified by IP address.	Detail
`Authentication requests`	Number of authentication requests received by the authentication server.	Detail
`Authentication rollover requests`	Number of requests coming into the server as a result of the previous server timing out.	Detail
`Authentication retransmissions`	Number of retransmissions.	Detail
`Accepts`	Number of authentication requests accepted by the authentication server.	Detail
`Rejects`	Number of authentication requests rejected by the authentication server.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Challenges`	Number of authentication requests challenged by the authentication server.	Detail
`Authentication malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Authentication bad authenticators`	Number of responses in which the authenticator is incorrect for the authentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Authentication requests pending`	Number of authentication requests waiting for a response.	Detail
`Authentication request timeouts`	Number of times an authentication request to the server timed out.	Detail
`Authentication unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Authentication packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
`Preauthentication requests`	Number of preauthentication requests received by the preauthentication server.	Detail
`Preauthentication rollover requests`	Number of preauthentication requests coming into the server as a result of the previous server timing out.	Detail
`Preauthentication retransmissions`	Number of retransmissions of preauthentication requests.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Preauthentication Accepts`	Number of preauthentication requests accepted by the preauthentication server.	Detail
`Preauthentication Rejects`	Number of preauthentication requests rejected by the preauthentication server.	Detail
`Preauthentication Challenges`	Number of preauthentication requests challenged by the preauthentication server.	Detail
`Preauthentication malformed responses`	Number of responses to preauthentication requests with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Preauthentication bad authenticators`	Number of responses in which the authenticator is incorrect for the preauthentication request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Preauthentication requests pending`	Number of preauthentication requests waiting for a response.	Detail
`Preauthentication request timeouts`	Number of times a preauthentication request to the server timed out.	Detail
`Preuthentication unknown responses`	Number of unknown responses during the preauthentication phase. The RADIUS response type in the header is invalid or unsupported.	Detail
`Preauthentication packets dropped`	Number of preauthentication packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail
`Accounting start requests`	Number of accounting start requests received.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting interim requests`	Number of accounting interim requests received.	Detail
`Accounting stop requests`	Number of accounting stop requests received.	Detail
`Accounting rollover requests`	Number of requests coming into the server as a result of the previous server timing out.	Detail
`Accounting retransmissions`	Number of retransmissions.	Detail
`Accounting start responses`	Number of accounting start responses sent by the server.	Detail
`Accounting interim responses`	Number of accounting interim responses sent by the server.	Detail
`Accounting stop responses`	Number of accounting stop responses sent by the server.	Detail
`Accounting malformed responses`	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).	Detail
`Accounting bad authenticators`	Number of responses in which the authenticator is incorrect for the accounting request. This can occur if the RADIUS secrets for the client and server do not match.	Detail
`Accounting requests pending`	Number of accounting requests waiting for a response.	Detail

Table 68: show network-access aaa radius-servers Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out.	Detail
`Accounting unknown responses`	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.	Detail
`Accounting packets dropped`	Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request.	Detail

Sample Output

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile1
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
Profile: xyz-profile2
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: DOWN ( 60 seconds )
Profile: xyz-profile11
    Server address: 2001:DB8:0:f101::2
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UP
```

show network-access aaa radius-servers

```
user@host> show network-access aaa radius-servers
Profile: xyz-profile3
    Server address: 192.168.30.188
      Authentication port: 1645
      Preauthentication port: 1810
      Accounting port: 1646
      Status: UNREACHABLE
Profile: xyz-profile3
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
      Preauthentication port: 1810
      Status: UNREACHABLE
```

show network-access aaa radius-servers detail

```
user@host> show network-access aaa radius-servers detail
Profile: xyz_profile5
    Server address: 192.168.30.188
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.190
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Accounting retry: 5
      Accounting port: 60
      Status: UP (accounting suspended, baseline in progress)
    Server address: 192.168.30.192
      Authentication port: 1812
      Preauthentication port: 1810
      Accounting port: 1813
      Status: UP
    Server address: 192.168.30.190
      Authentication port: 1812
      Accounting port: 1813
      Accounting retry: 5
      Accounting port: 60
      Status: UP
    Server address: 192.168.30.192
      Authentication port: 1812
      Accounting port: 1813
      Status: UP

RADIUS Servers
  192.168.30.188
    Authentication requests: 7658
    Authentication rollover requests: 0
    Authentication retransmissions: 3600
    Accepts: 6458
    Rejects: 0
    Challenges: 0
    Authentication malformed responses: 0
    Authentication bad authenticators: 0
    Authentication requests pending: 0
    Authentication request timeouts: 4800
    Authentication unknown responses: 0
    Authentication packets dropped: 0
    Preauthentication requests: 7658
    Preauthentication rollover requests: 0
    Preauthentication retransmissions: 3600
    Preauthentication Accepts: 6458
    Preauthentication Rejects: 0
    Preauthentication Challenges: 0
    Preauthentication malformed responses: 0
    Preauthentication bad authenticators: 0
    Preauthentication requests pending: 0
    Preauthentication request timeouts: 4800
    Preauthentication unknown responses: 0
    Preauthentication packets dropped: 0
    Accounting start requests: 1
    Accounting interim requests: 1
    Accounting stop requests: 0
    Accounting rollover requests: 0
    Accounting retransmissions: 0
    Accounting start responses: 1
    Accounting interim responses: 1
    Accounting stop responses: 0
    Accounting malformed responses: 0
    Accounting bad authenticators: 0
    Accounting requests pending: 0
    Accounting request timeouts: 0
    Accounting unknown responses: 0
    Accounting packets dropped: 0
```

show network-access aaa statistics

IN THIS SECTION

Syntax

```
show network-access aaa statistics
<accounting (detail)>
<address-assignment (client | pool pool-name)>
<dynamic-requests>
<radius>
<session-limit-per-username>
```

Description

Display AAA accounting, address-assignment, dynamic request statistics, RADIUS settings and statistics, and subscriber session limit statistics.

Options

accounting (detail)

(Optional) Display AAA accounting statistics. The `detail` keyword displays additional accounting information

address-assignment (client | pool *pool-name*)

(Optional) Display AAA address-assignment client and pool statistics.

dynamic-requests

(Optional) Display AAA dynamic requests.

radius

(Optional) Display RADIUS settings and statistics.

session-limit-per-username

Maximum number of sessions allowed for a username per access profile. Use the `brief` option to display only active users with blocked requests. Use the `detail` option to display all active users.

Required Privilege Level

view

Output Fields

Table 69 on page 884 lists the output fields for the `show network-access aaa statistics` command. Output fields are listed in the approximate order in which they appear.

Table 69: show network-access aaa statistics Output Fields

Field Name	Field Description	Level of Output
`Requests received`	• Number of accounting requests generated by the AAA framework. • Number of dynamic requests received from the external server. Does not include requests sent from backup accounting.	All levels

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting request failures`	Number of accounting requests that failed to be sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting request success`	Number of accounting requests successfully sent or queued from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Account on requests`	Number of accounting on requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting start requests`	Number of accounting start requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting interim requests`	Number of accounting interim requests sent from a client to a RADIUS accounting server.	`detail`
`Accounting stop requests`	Number of accounting stop requests sent from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting request timeouts`	Number of accounting requests to the accounting server that timed out. This field was named `Timed out requests` in releases before Junos OS Release 16.1. Does not include requests sent from backup accounting.	All levels
`Accounting Response failures`	Number of accounting requests not acknowledged (NAK) by the accounting server. Does not include requests sent from backup accounting.	All levels
`Accounting response success`	Number of accounting requests acknowledged by the accounting server. Does not include requests sent from backup accounting.	All levels
`Account on responses`	Number of accounting on requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting start responses`	Number of accounting start requests acknowledged by the RADIUS accounting server.	`detail`
`Accounting interim responses`	Number of accounting interim requests acknowledged by the RADIUS accounting server.	`detail`

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting stop responses`	Number of accounting stop requests acknowledged by the RADIUS accounting server. Does not include requests sent from backup accounting.	`detail`
`Accounting rollover requests`	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.	`detail`
`Accounting unknown requests`	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, when the header has invalid or unsupported information).	`detail`
`Accounting radius pending requests`	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.	`detail`
`Accounting malformed responses`	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.	`detail`
`Accounting retransmissions`	Number of accounting requests made by a client to the RADIUS sever that were retransmitted. Does not include requests sent from backup accounting.	`detail`
`Accounting bad authenticators`	Number of accounting responses from a RADIUS accounting server that have an incorrect authenticator (for example, the client and server RADIUS secret do not match).	`detail`

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting packets dropped`	Number of accounting responses from a RADIUS accounting server that are dropped by a client.	`detail`
`Accounting backup record creation requests`	Number of accounting stop requests from a client to a RADIUS accounting server that were forwarded to be backed up.	`detail`
`Accounting backup replay request success`	Number of backup accounting stop requests successfully created by clients after each timeout for replay to a RADIUS accounting server.	`detail`
`Accounting backup request failures`	Number of backup accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.	`detail`
`Accounting backup request success`	Number of backup accounting requests successfully sent or queued from a client to a RADIUS accounting server.	`detail`
`Accounting backup timeouts`	Number of backup accounting requests that timed out after being sent to a RADIUS accounting server.	`detail`

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting backup in-flight requests	Number of backup accounting requests that were successfully sent or queued to a RADIUS accounting server for which no response or error has been received yet. Backup requests are replayed only in the following circumstances: • When the request being replayed receives a positive response, the next request can be replayed. • When the request being replayed receives a timeout response, it can be replayed again. Consequently this intermediate timer displays 1 or 0. The value eventually drops to 0 as requests are responded to positively or fail due to error.	detail
Accounting backup responses success	Number of backup records that were successfully acknowledged with a positive response from a RADIUS accounting server.	detail
Accounting backup radius requests	Number of backup requests sent to UDP level. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. An observation that the value is increasing is more significant than the exact value of the counter.	detail

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting backup radius responses	Number of responses received at the UDP level for backup requests. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup radius timeouts	Number of backup requests that timed out after being sent to UDP. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup radius pending requests	Number of backup requests sent to a RADIUS accounting server that are waiting for a response from the server. This is an intermediate state counter that eventually drops to zero as requests are responded to or failed due to error.	detail
Accounting backup radius retransmissions	Sum of backup request retransmissions for each RADIUS accounting server. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting backup malformed responses	Sum of malformed responses received for backup requests sent to each RADIUS accounting server at the UDP level.	detail
Accounting backup bad authenticators	Sum of responses received for backup accounting requests for each RADIUS accounting server where authenticators were mismatched.	detail
Accounting backup responses dropped	Sum of responses for backup accounting requests for each RADIUS accounting server that were dropped due to various sanity checks.	detail
Accounting backup rollover requests	Sum of backup accounting requests rolled over for each RADIUS accounting server.	detail
Accounting backup unknown responses	Sum of unknown responses for backup accounting requests for each RADIUS accounting server.	detail
Client	Client type; for example, DHCP, Mobile IP, PPP.	none specified
Out of Memory	Number of times an address was not given to the client due to memory issues.	none specified
No Matches	Number of times there were no network matches for the pool.	none specified

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Pool Name`	Name of the address-assignment pool for this client.	none specified
`Out of Addresses`	Number of times there were no available addresses in the pool.	none specified
`Address total`	Number of addresses in the pool.	none specified
`Addresses in use`	Number of addresses in use.	none specified
`Addresses excluded`	Number of addresses excluded from being allocated from the pool with the `excluded-address` or `excluded-range` statements.	none specified
`Address Usage (percent)`	Percentage of total addresses in use. This value does not take excluded addresses into account.	none specified
`Pool drain configured`	Configuration state of active drain for the specified local address pool, `yes` or `no`.	none specified
`Pool Usage`	Percentage of allocated addresses in the specified address pool.	none specified
`processed successfully`	Number of dynamic requests processed successfully by the AAA framework.	All levels
`errors during processing`	Number of dynamic requests that resulted in processing errors by the AAA framework.	All levels

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link Name	Name of the secondary address-assignment pool to which the primary pool is linked.	
silently dropped	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.	All levels
RADIUS Server	IPv4 or IPv6 address of the RADIUS server to which the router is sending requests.	All levels
Profile	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.	All levels
Configured	Configured maximum number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded. The range of values is 0 through 2000 outstanding requests. The default value is 1000.	All levels
Current	Current number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded.	All levels

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Peak	Highest number of outstanding requests from the router to the RADIUS server for a specific profile at any point in time since the router was started or since the counter was last cleared. **NOTE**: If the value of this field is equal to the value of the Configured field, you may want to increase the value of the Configured field.	All levels
Exceeded	Number of times that the router attempted to send requests to the RADIUS server in excess of the configured maximum value for a specific profile. **NOTE**: If the value of this field is nonzero, you may want to increase the value of the Configured field.	All levels
Username	Username for a subscriber with one or more active sessions for an access profile.	briefdetail
Access-profile	Name of the access profile where the username is active.	briefdetail
Blocked requests	Number of session requests that have been blocked for the username for an access profile. A request is blocked when it exceeds the configured session limit.	briefdetail
Session count	Number of active sessions for the username for an access profile.	briefdetail
Total usernames	Number of active usernames for all access profiles.	none summary

Table 69: show network-access aaa statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Total usernames exceeding session limit`	Number of usernames that have attempted sessions greater than the limit configured for the username.	none `summary`
`Total blocked requests`	Number of session requests that have been blocked because the session limit is exceeded.	none `summary`

Sample Output

show network-access aaa statistics accounting

```
user@host> show network-access aaa statistics accounting
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
```

show network-access aaa statistics accounting detail

```
user@host> show network-access aaa statistics accounting detail
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request failures: 0
  Accounting request success: 5000
    Account on requests: 0
    Accounting start requests: 3000
    Accounting interim requests: 0
    Accounting stop requests: 2000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
    Account on responses: 0
    Accounting start responses: 3000
    Accounting interim responses: 0
    Accounting stop responses: 0
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting radius pending requests: 0
  Accounting malformed responses: 0
  Accounting retransmissions: 6000
  Accounting bad authenticators: 0
  Accounting packets dropped: 0

  Accounting backup record creation requests: 3000
  Accounting backup request replay success: 9808
  Accounting backup request failures: 0
  Accounting backup request success: 3006
  Accounting backup timeouts: 6
  Accounting backup in-flight requests: 0
  Accounting backup responses success: 3000
  Accounting backup radius requests: 3006
  Accounting backup radius responses: 3000
  Accounting backup radius timeouts: 99
  Accounting backup radius pending requests: 0
  Accounting backup radius retransmissions: 99
  Accounting backup malformed responses: 0
  Accounting backup bad authenticators: 0
  Accounting backup responses dropped: 0
  Accounting backup rollover requests: 0
  Accounting backup unknown responses: 0
```

show network-access aaa statistics address-assignment client

```
user@host> show network-access aaa statistics address-assignment client
Address-assignment statistics
 Client: jdhcpd
 Out of Memory: 0
 No Matches: 2
```

show network-access aaa statistics address-assignment pool

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 9
 Address total: 47
 Addresses in use: 47
 Address Usage (percent): 100
 Pool drain configured: yes
```

show network-access aaa statistics address-assignment pool (Excluded Addresses)

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 0
 Address total: 24000
 Addresses in use: 12000
 Addresses excluded: 1000
 Address Usage (percent): 50
 Pool drain configured: yes
```

show network-access aaa statistics dynamic-requests

```
user@host> show network-access aaa statistics dynamic-requests
requests received: 0
processed successfully: 0
errors during processing: 0
silently dropped: 0
```

show network-access aaa statistics radius

```
user@host> show network-access aaa statistics radius
Outstanding Requests
RADIUS Server           Profile       Configured   Current   Peak   Exceeded
198.51.100.239          prof1         1000         0         1000   14
                        prof2         500          17        432    0
198.51.100.211          myprof        200          0         200    27
203.0.113.254           pppoe-auth    111          0         1      0
2001:db8:0:f101::2      xyz-profile11 1000         10        135    0
```

show network-access aaa statistics session-limit-per-username (Users with Blocked Requests)

```
user@host> show network-access aaa statistics session-limit-per-username brief
Username           Access-profile     Blocked requests     Session count
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
```

show network-access aaa statistics session-limit-per-username (All Active Users)

```
user@host> show network-access aaa statistics session-limit-per-username detail
Username           Access-profile     Blocked requests     Session count
rkv@example.net    BNG1               0                    4
xyz@example.net    BNG1               3                    5
abc@example.net    BNG2               2                    5
pqr@example.net    BNG2               0                    1
```

show network-access aaa statistics session-limit-per-username

```
user@host> show network-access aaa statistics session-limit-per-username
Total usernames:  15
Total usernames exceeding session limit: 2
Total blocked requests:  5
```

show network-access aaa statistics authentication

IN THIS SECTION

Syntax

```
show network-access aaa statistics authentication
<detail>
```

Description

Display AAA authentication statistics.

Options

`detail` (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 70 on page 900 lists the output fields for the `show network-access aaa statistics authentication` command. Output fields are listed in the approximate order in which they appear.

Table 70: show network-access aaa statistics authentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of authentication requests received from clients.	All levels
`Accepts`	Number of authentication requests accepted by the authentication server.	All levels
`Rejects`	Number of authentication requests rejected by the authentication server.	All levels
`Challenges`	Number of authentication requests challenged by the authentication server.	All levels
`Timed out requests`	Number of authentication requests that timed out.	All levels
`RADIUS authentication failures`	Number of RADIUS authentication requests that have failed.	Detail
`Queue request deleted`	Number of queue requests that have been deleted.	Detail
`Malformed reply`	Number of malformed replies received from the RADIUS authentication server.	Detail
`No server configured`	Number of authentication requests that failed because no authentication server is configured.	Detail
`Access Profile configuration not found`	Number of authentication requests that failed because no access profile is configured.	Detail

Table 70: show network-access aaa statistics authentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Unable to create client record`	Number of times that the router is unable to create the client record for the authentication request.	Detail
`Unable to create client request`	Number of times that the router is unable to create the client request for the authentication request.	Detail
`Unable to build authentication request`	Number of times that the router is unable to build the authentication request.	Detail
`No server found`	Number of requests to the authentication server that have timed out; the server is then considered to be down.	Detail
`Unable to create handle`	Number of authentication requests that have failed because of an internal allocation failure.	Detail
`Unable to queue request`	Number of times the router was unable to queue the request to the authentication server.	Detail
`Invalid credentials`	Number of times the router did not have proper authorization to access the authentication server.	Detail
`Malformed request`	Number of times the router request to the authentication server is malformed.	Detail
`License unavailable`	Number of times the router did not have a license to access the authentication server.	Detail
`Redirect requested`	Number of authentication requests that have been redirected based on routing instance.	Detail
`Internal failure`	Number of internal failures.	Detail

Table 70: show network-access aaa statistics authentication Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Local authentication failures`	Number of times local authentication failed.	Detail
`LDAP lookup failures`	Number of times the LDAP lookup operation failed.	Detail

Sample Output

show network-access aaa statistics authentication

```
user@host> show network-access aaa statistics authentication
Authentication module statistics
  Requests received: 2118
    Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

show network-access aaa statistics authentication detail

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
    Accepts: 261
  Rejects: 975
    RADIUS authentication failures: 975
      Queue request deleted: 0
      Malformed reply: 0
      No server configured: 0
      Access Profile configuration not found: 0
      Unable to create client record: 0
      Unable to create client request: 0
      Unable to build authentication request: 0
      No server found: 975
      Unable to create handle: 0
      Unable to queue request: 0
      Invalid credentials: 0
      Malformed request: 0
      License unavailable: 0
      Redirect requested: 0
      Internal failure: 0
    Local authentication failures: 0
    LDAP lookup failures: 0
  Challenges: 0
  Timed out requests: 882
```

show network-access aaa statistics pending-accounting-stops

IN THIS SECTION

Syntax

```
show network-access aaa statistics pending-accounting-stops
```

Description

Display the number of pending accounting stop requests.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 71 on page 904 lists the output field for the `show network-access aaa statistics pending-accounting-stops` command.

Table 71: show network-access aaa statistics pending-accounting-stops Output Fields

Field Name	Field Description
`Pending accounting stops`	Total number of accounting stop messages queued.

Sample Output

show network-access aaa statistics pending-accounting-stops

```
user@host> show network-access aaa statistics pending-accounting-stops
Pending accounting stops: 10,000
```

show network-access aaa statistics preauthentication

IN THIS SECTION

Syntax

```
show network-access aaa statistics preauthentication
```

Description

Display AAA preauthentication statistics.

Options

`detail` (Optional) Displays detailed information about authentication.

Required Privilege Level

view

Output Fields

Table 72 on page 906 lists the output fields for the `show network-access aaa statistics preauthentication` command. Output fields are listed in the approximate order in which they appear.

Table 72: show network-access aaa statistics preauthentication Output Fields

Field Name	Field Description	Level of Output
`Requests received`	Number of preauthentication requests received from clients.	All levels
`Multistack requests`	Number of preauthentication requests for dual-stack subscribers.	All levels
`Accepts`	Number of preauthentication requests accepted by the preauthentication server.	All levels
`Rejects`	Number of preauthentication requests rejected by the preauthentication server.	All levels
`Challenges`	Number of preauthentication requests challenged by the preauthentication server.	All levels
`Timed out requests`	Number of preauthentication requests that timed out.	All levels

Sample Output

show network-access aaa statistics preauthentication

```
user@host> show network-access aaa statistics preauthentication
Preauthentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
  Challenges: 0
  Timed out requests: 882
```

show network-access aaa statistics re-authentication

IN THIS SECTION

Syntax

```
show network-access aaa statistics re-authentication
```

Description

Display statistics for RADIUS re-authentication, and starting in Junos OS Release 18.2R1, for local re-authentication.

Required Privilege Level

view

Output Fields

Table 73 on page 908 lists the output fields for the show network-access aaa statistics re-authentication command. Output fields are listed in the approximate order in which they appear.

Table 73: show network-access aaa statistics re-authentication Output Fields

Field Name	Field Description
`Re-authentication statistics`	Displays re-authentication statistics.
`Requests received`	Total number of re-authentication requests that the device received from clients.
`Accepts`	Total number of accepted re-authentications.
`Challenges`	Total number of re-authentication challenges.
`Internal errors`	Total number of re-authentication internal errors.
`Rejects`	Total number of re-authentications rejected.
`Timed out requests`	Total number of re-authentication accounting timeouts.

show network-access aaa statistics re-authentication

command-name

```
user@host> show network-access aaa statistics re-authentication
Re-authentication statistics
  Requests received: 0
  Accepts: 0
  Rejects: 0
  Challenges: 0
  Timed out requests: 0
```

show network-access aaa terminate-code

IN THIS SECTION

Syntax

```
show network-access aaa terminate-code
<brief | detail | summary>
<reverse>
<(aaa | dhcp | l2tp | ppp)>
```

Description

Display the count for termination cause types and the current mapping between session termination cause types and code values.

Options

none	Display all mappings.
brief | detail | summary	(Optional) Display the specified level of output. The `summary` output is displayed by default and includes base count information about mappings. The `brief` output displays mappings with non-zero usage count and custom mappings. The `detail` output displays all mappings.
aaa	(Optional) Limit display to AAA mappings only.
dhcp	(Optional) Limit display to DHCP mappings only.
l2tp	(Optional) Limit display to L2TP mappings only.
ppp	(Optional) Limit display to PPP mappings only.
reverse	(Optional) Display mapping of the code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) to the termination cause type.
vlan	(Optional) Limit display to VLAN mappings only.

Required Privilege Level

view

Output Fields

Table 74 on page 911 lists the output fields for the `show network-access aaa terminate-code` command. Output fields are listed in the approximate order in which they appear.

Table 74: show network-access aaa terminate-code Output Fields

Field Name	Field Description	Level of Output
RADIUS	RFC-defined code value conveyed in the RADIUS Acct-Terminate-Cause attribute (49) or a nonstandard, customized value that you configure with the `terminate-code aaa` statement at the `[edit access]` hierarchy level.	`brief detail` None (with `reverse` option)
Custom	Whether or not the termination cause is a customized mapping or the default mapping.	All levels
Mapping-Count	Number of mappings that occurred for a specific terminate cause type or category (standard or summary output) or per termination cause (reverse output).	`summary` None
Usage-Count	Number of times the terminate code mapping was used.	All levels
Type	Termination cause type—null, aaa, dhcp, l2tp, ppp, or vlan. **NOTE**: The null termination cause type indicates that no termination reason was provided by the subscriber and the RADIUS Acct-Terminate-Cause attribute (49) was not included in the Acct-Stop request	All levels
Code	Specific termination cause.	`brief detail`

Sample Output

show network-access aaa terminate-code

```
user@host> show network-access aaa terminate-code
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code reverse

```
user@host> show network-access aaa terminate-code reverse
Terminate-code:
  RADIUS    Custom Mapping-Count Usage-Count Type
  0         no     1             0           null
  1         no     1             0           aaa
  1         no     1             0           dhcp
  1         no     5             0           l2tp
  1         no     8             0           ppp
  1         no     2             10          vlan
  2         no     1             0           dhcp
  2         no     3             0           ppp
  2         no     2             0           vlan
  4         no     1             0           aaa
  4         no     1             0           dhcp
  4         no     1             0           l2tp
  4         no     1             0           ppp
  5         no     2             0           aaa
  5         no     1             0           l2tp
  5         no     1             0           ppp
  6         no     2             0           aaa
  6         no     13            0           l2tp
  6         no     3             0           ppp
  6         no     3             0           vlan
  8         no     3             0           l2tp
  8         no     5             0           ppp
  9         no     13            0           l2tp
  9         no     12            0           ppp
  9         no     4             0           vlan
  10        no     4             0           aaa
  10        no     1             0           dhcp
  10        no     128           0           l2tp
  10        no     171           0           ppp
  15        no     1             0           dhcp
  15        no     190           0           l2tp
  15        no     1             0           vlan
  16        no     1             0           vlan
  17        no     2             0           aaa
  17        no     10            0           l2tp
  17        no     6             0           ppp
```

show network-access aaa terminate-code dhcp

```
user@host> show network-access aaa terminate-code dhcp
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     5             0           dhcp
```

show network-access aaa terminate-code detail

```
user@host> show network-access aaa terminate-code aaa detail

Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  17         no     0           aaa  deny-server-request-timeout
  6          no     0           aaa  service-shutdown-network-logout
  10         no     0           aaa  service-shutdown-remote-reset
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  5          no     0           aaa  service-shutdown-time-limit
  10         no     0           aaa  service-shutdown-volume-limit
  6          no     13          aaa  shutdown-administrative-reset
  4          no     0           aaa  shutdown-idle-timeout
  10         no     0           aaa  shutdown-remote-reset
  5          no     0           aaa  shutdown-session-timeout
```

show network-access aaa terminate-code brief

```
user@host> show network-access aaa terminate-code brief
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  17         no     1           aaa  deny-authentication-denied
  10         no     1           aaa  deny-no-resources
  1200       yes    5           aaa  service-shutdown-subscriber-logout
  6          no     13          aaa  shutdown-administrative-reset
  15         no     7           dhcp nak
  10         no     1           l2tp session-receive-cdn-avp-missing-secret
  10         no     1           ppp  bundle-fail-create
  1          no     1           ppp  lcp-peer-terminate-term-req
  10         no     1           ppp  lcp-tunnel-disconnected
  1          no     10          vlan out-of-band-ancp-port-down
```

show network-access aaa terminate-code summary

```
user@host> show network-access aaa terminate-code summary

Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     1             0           null
  no     12            0           aaa
  no     5             0           dhcp
  no     364           0           l2tp
  no     210           0           ppp
  no     13            10          vlan
```

show network-access aaa terminate-code vlan

```
user@host> show network-access aaa terminate-code vlan
Terminate-code:
  Custom Mapping-Count Usage-Count Type
  no     13            0           vlan
```

show network-access aaa terminate-code vlan detail

```
user@host> show network-access aaa terminate-code vlan detail
Terminate-code:
  RADIUS     Custom Usage-Count Type Code
  6          no     0           vlan admin-logout
  16         no     0           vlan admin-reconnect
  9          no     0           vlan other
  2          no     0           vlan out-of-band-access-interface-down
  6          no     0           vlan out-of-band-admin-access-interface-down
  6          no     0           vlan out-of-band-admin-core-interface-down
  1          no     0           vlan out-of-band-ancp-port-down
  1          no     0           vlan out-of-band-ancp-port-vlan-id-change
  2          no     0           vlan out-of-band-core-interface-down
  15         no     0           vlan out-of-band-l2-wholesale-no-free-vlans
  9          no     0           vlan profile-request-error
  9          no     0           vlan sdb-error
  9          no     0           vlan subscriber-activate-error
```

show network-access aaa subscribers

IN THIS SECTION

Syntax

```
show network-access aaa subscribers
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<statistics>
<username>
<session-id session-id-number detail>
```

Description

Display subscriber-specific AAA statistics.

Options

logical-system *logical-system-name* — (Optional) List subscribers in the specific logical system.

routing-instance *routing-instance-name* — (Optional) List subscribers for the specific routing instance. If you do not specify a routing instance name, the default routing instance is assumed.

statistics — (Optional) Display statistics for the subscriber events.

username — (Optional) Display information for the specified subscriber.

session-id *session-id-number* detail — (Optional) Display information for the specified session ID.

Required Privilege Level

view

Output Fields

"show network-access aaa subscribers" on page 915 lists the output fields for the `show network-access aaa subscribers` command. Output fields are listed in the approximate order in which they appear.

Table 75: show network-access aaa subscribers Output Fields

Field Name	Field Description
`Challenge requests`	Number of authentication requests challenged by the authentication server for this subscriber.
`Challenge responses`	Number of challenge responses sent by the subscriber to the authentication server.
`START sent successfully`	Number of accounting start requests generated by the AAA framework for this subscriber.
`START send failures`	Number of accounting start requests that failed to make it to the accounting server for this subscriber.
`START ack received`	Number of accounting start requests acknowledged by the accounting server for this subscriber.
`INTERIM sent successfully`	Number of accounting interim requests generated by the AAA framework for this subscriber.
`INTERIM send failures`	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
`INTERIM ack received`	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
`Requests received`	Number of reauthentication requests received by the authentication server.
`Successful responses`	Number of successful reauthentication requests granted by the authentication server.

Table 75: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
`Aborts handled`	Number of reauthentication requests terminated by the authentication server.
`Service name`	Name of the subscriber service.
`Creation requests`	Number of requests to create the service.
`Deletion requests`	Number of requests to delete the service.
`Request timeouts`	Number of times the service request was timed out.
`Client type`	Type of client; for example, DHCP, Mobile IP, PPP.
`Session-ID`	ID of the subscriber session.
`Session uptime`	How long the session has been up, in *HH:MM:SS.*
`Accounting`	Status of accounting, and type of accounting if accounting is on.
`Stripped username`	Username of the subscriber session.
`AAA Logical system/ Routing instance`	AAA framework for the subscriber of logical system or routing instance.
`Target Logical system/Routing instance`	Target framework for the subscriber of logical system or routing instance.
`Access-profile`	Profile of the subscriber.

Table 75: show network-access aaa subscribers Output Fields *(Continued)*

Field Name	Field Description
`Accounting Session ID`	ID of the subscriber session for accounting.
`Multi Accounting Session ID`	ID of the subscriber session for multiple accouting.
`IP Address`	IPv4 address of the subscriber.
`IPv6 Address`	IPv6 address of the subscriber.
`IPv6 Prefix`	IPv6 prefix of the subscriber.
`Authentication State`	State of subscriber session authentication.
`Accounting State`	State of subscriber session accounting.
`Provisioning Type`	Type of subscriber provisioning.

Sample Output

show network-access aaa subscribers

```
user@host> show network-access aaa subscribers
Username                 Logical system/Routing instance   Client type    Session-ID
user-example-1           default:default                   pppoe
```

show network-access aaa subscribers (DHCP)

```
user@host> show network-access aaa subscribers
Username          Logical system/Routing instance   Client type    Session-ID
user-example-1     default:default                   dhcp           8
                   default:default                   dhcp           9
```

show network-access aaa subscribers logical-system

```
user@host> show network-access aaa subscribers logical-system
Username            Client type    Logical system/Routing instance
user61@example.net    ppp            default
00010e020304.1231   dhcp           isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp           default:isp-gtown-r3-00
0020df980102.2334   dhcp           isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers logical-system routing-instance

```
user@host> show network-access aaa subscribers logical-system isp-bos-metro-16 routing-instance
isp-cmbrg-12-32
Username            Client type   Logical system/Routing instance
00010e020304.1231   dhcp          isp-bos-metro-12:isp-cmbrg-12
user54@example.com     dhcp          default:isp-gtown-r3-00
0020df980102.2334   dhcp          isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers statistics username

```
user@host> show network-access aaa subscribers statistics username 00010e020304.1231
Authentication statistics
   Challenge requests: 0
   Challenge responses: 0
 Accounting statistics
   START sent successfully: 1
   START send failures: 0
   START ack received: 1
   INTERIM sent successfully: 0
   INTERIM send failures: 0
   INTERIM ack received: 0
 Re-authentication statistics
   Requests received: 0
   Sucessfull responses: 0
   Aborts handled: 0
 Service statistics
   Service name: filter-serv
   Creation requests: 1
   Deletion requests: 0
   Request timeouts: 0
   Service name: filter-serv2
   Creation requests: 144
   Deletion requests: 0
   Request timeouts: 144
```

show network-access aaa subscribers username

```
user@host> show network-access aaa subscribers username user80@example.net
Logical system/Routing instance   Client type  Session-ID  Session uptime  Accounting
isp-bos-metro-16:isp-cmbrg-12     dhcp         7           01:12:56        on/volume
Service name       Service type      Quota          Accounting
I-Cast             volume            1200 Mbps      on/volume+time
Voip                                                on/volume
GamingBurst        time              6000 secs      on/volume
```

show network-access aaa subscribers session-id 26 detail

The following command output is seen when only an IPv4 client is associated with the session:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
```

```
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

The following command output is seen when IPv6 client logs in (after IPv4 association) and is associated with the same session ID:

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
IPv6 Address: 2001:db8:0:8003::2
IPv6 Prefix: 2001:db8:ffff:0:4::/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show network-access aaa subscribers session-id

IN THIS SECTION

Syntax

```
show network-access aaa subscribers session-id session-id
<brief | detail>
```

Description

Display information about the specified subscriber session.

Options

session-id	ID of the subscriber session.
brief | detail	(Optional) Display the specified level of information.

Required Privilege Level

view

Output Fields

Table 76 on page 923 lists the output fields for the `show network-access aaa subscribers session-id` command. Output fields are listed in the approximate order in which they appear.

Table 76: show network-access aaa subscribers session-id Output Fields

Field Name	Field Description	Level of Output
`Type` and `Client type`	Type of client.	All levels

Table 76: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Accounting`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time`, `volume+time`, or `flat-file`. The `time` and `volume+time` types are configured in RADIUS Service-Statistics VSA [26-69].	brief none
`Service type`	Type of accounting: `volume`, `time`, `volume+time`, or `na`.	brief
`Quota`	Quota for service: volume (in Mbps) or time (seconds).	brief
`Username`	Name of the user logged in to the session.	detail
`Stripped username`	Username after the domain has been removed.	detail
`Logical system/ Routing instance` and `AAA Logical system/ Routing instance`	Name of the routing instance, logical system name, or both used for the session.	All levels
`Target Logical system/Routing instance`	Logical system/routing instance to which the session is mapped.	detail
`Access-profile`	Access profile used for AAA services for the session.	detail
`Session ID`	ID of the subscriber session. The session ID value displayed under `Service name` is the service session ID.	detail
`Accounting Session ID`	ID of the accounting session (RADIUS attribute 44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail

Table 76: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Multi Accounting Session ID	Bundle ID for MLPPP sessions. Acct-Multi-Session-Id (RADIUS attribute 50) uses the value of the session database bundle session ID to enable RADIUS to link together multiple related sessions. The value of this field is zero when no MLPPP sessions exist.	detail
IP Address	IP address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
IPv6 Address	IPv6 address of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
IPv6 Prefix	IPv6 prefix of the subscriber. For a single-session dual stack, addresses of both IPv4 and IPv6 clients are displayed.	detail
Authentication State	State of the subscriber authentication session: AuthInit, AuthStart, AuthChallenge, AuthRedirect, AuthClntRespWait, AuthAcctVolStatsAckWait, AuthAcctStopAckWait, AuthServCreateRespWait, AuthLogoutStart, AuthStateActive, AuthClntLogoutRespWait, AuthProfileUpdateWait, AuthProvisionRespWait, AuthProvisionServiceCreationWait	detail
Ocs Subscription-Id-Type	Type of subscriber for an OCS partition. You can define your own or use a predefined value: 0 (END_USER_E164), 1 (END_USER_IMSI), 2 (END_USER_SIP_URI), 3 (END_USER_NAI), 4 (END_USER_PRIVATE).	detail
Ocs Subscription-Id-Data	Subscriber data string concatenated from a list of user-selected data options used to identify the subscriber type for an OCS partition; for example: test-sid	detail
Ocs Interrogation State	State of the OCS interrogation: first, intermediate, final.	detail
Ocs Data State	State of the OCS data: none	detail

Table 76: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Accounting State	State of the subscriber accounting session: Acc-Init, Acc-Start-Sent, Imm-Update-Stats-Pending, Acc-Interim-Sent, Acc-Stop-Stats-Pending, Acc-Stop-Sent, Acc-Stop-On-Fail-Deny-Sent, Acc-Stop-Ackd	detail
Provisioning-type	Provisioning type for this session: • gx-plus—Subscriber service uses Gx-Plus provisioning. • jsrc—Subscriber service uses JSRC provisioning. • none—Provisioning is not enabled.	detail
Service name	Name of the attached service or policy. • For RADIUS-activated and CLI-activated services, displays the full activation string for the service. If the activation string includes service parameters, then both the service name and service parameters are displayed. • For JSRC-activated policies—displays the policy name.	All levels
Service State	State of the service provided in the subscriber session.	detail
Service Family	Network family of the service provided in the subscriber session.	detail
Service Activation Source	Source used to activate the service.	detail
Session uptime	How long the session has been up, in *HH:MM:SS*.	All levels
Service CC-Service-Identifier	Data identification element of the 3GPP Diameter credit control service charging system that uniquely defines the CC-Service-Context.	detail

Table 76: show network-access aaa subscribers session-id Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Service Rating-Group`	Value associated with a charging rule and part of the accounting data stream for the PCRF.	detail
`Ocs Control`	Whether OCS controls the service: `yes` or `no`.	detail
`Accounting status`	Status of the accounting configuration for the service, `on` or `off`, and the type of accounting, `time`, `volume+time`, or `flat-file`. The `time` and `volume+time` types are configured in RADIUS Service-Statistics VSA [26-69].	detail
`Service accounting session ID`	ID of the service accounting session; RADIUS Acct-Session-Id attribute (44). The ID appears in decimal or description format, as specified by the `accounting-session-id-format` statement.	detail
`Service accounting state`	State of the service accounting session: `Acc-Init`, `Acc-Start-Sent`, `Imm-Update-Stats-Pending`, `Acc-Interim-Sent`, `Acc-Stop-Stats-Pending`, `Acc-Stop-Sent`, `Acc-Stop-On-Fail-Deny-Sent`, `Acc-Stop-Ackd`	detail
`Accounting interim interval`	Amount of time between interim accounting updates for this service, in seconds; RADIUS Service-Interim-Acct-Interval VSA [26-140] or Diameter Acct-Interim-Interval AVP (85).	detail

Sample Output

show network-access aaa subscribers session-id brief

```
user@host> show network-access aaa subscribers session-id 6 brief
Logical system/Routing instance    Client type   Session uptime    Accounting
default:default                    dhcp          00:01:29          on/time

Service name                   Service type   Quota           Accounting
```

```
filter-service             -na-           -na-            off
filter-service-2           volume+time    77.00MB/120secs off
1337994190863204450        -na-           -na-            off
```

show network-access aaa subscribers session-id detail

```
user@host> show network-access aaa subscribers session-id 5 detail
Type: dhcp
Username: user23@example.net
Stripped username: user23
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:retail-onlinecompany-ca
Access-profile:retailer-onlinecompany-sjc
Session ID: 5
Accounting Session ID: jnpr ge-1/0/0.101:1
Multi Accounting Session ID: 0
IP Address: 192.168.44.104
Authentication State: AuthStateActive

Ocs Subscription-Id-Type: 15
Ocs Subscription-Id-Data: test-sid
Ocs Interrogation State: intermediate
Ocs Data State: none
Gx-Plus Provisioning State: response-received
Accounting State: Acc-Interim-Sent
Provisioning-type: jsrc
Service name: filter-service-1
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN
 Session ID: 7
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 777
 Service Rating Group: 10
 Ocs Control: yes
Service name: filter-service-2
 Service State: SvcActive
 Service Family: inet
 Service Activation Source: PCRF-LOGIN Session ID: 8
 Session uptime: 00:01:33
 Service CC-Service-Identifier: 778
 Service Rating Group: 11
 Ocs Control: no
Accounting status: on/volume+time
 Service accounting session ID: 1:2-1322506006
 Service accounting state: Acc-Interim-Sent
 Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Service with Multiple Instances)

```
user@host> show network-access aaa subscribers session-id 6 detail
Type: dhcp
Stripped username: user-test-fms2
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: attr_test_profile1
Session ID: 6
Accounting Session ID: 6
Multi Accounting Session ID: 0
IP Address: 198.51.100.10
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
Service name: economy-service(up-filter,down-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius at Reauth
  Session ID: 7
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:7-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
Service name: economy-service(upstrm-filter,dwnstrm-filter)
  Service State: SvcActive
  Service Family: inet
  Service Activation Source: Radius
  Session ID: 8
  Session uptime: 00:04:36
  Accounting status: on/volume+time
  Service accounting session ID: 6:8-1354811427
  Service accounting state: Acc-Start-Sent
  Accounting interim interval: 600
```

show network-access aaa subscribers session-id detail (Single Session Dual Stack with active V4 and V6 subscribers)

```
user@host> show network-access aaa subscribers session-id 26 detail
Type: dhcp
Stripped username: user-test-25
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.10.0.6
IPv6 Address: 00:00:5E:00:53:02
IPv6 Prefix: 00:00:5E:00:53:00/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

show ppp interface

IN THIS SECTION

Syntax

```
show ppp interface interface-name
<extensive | terse>
```

Description

Display information about PPP interfaces.

Options

interface-name Name of a logical interface.

Starting in Junos OS Release 17.3, the * (asterisk) wildcard character is supported for the interface name for debugging purpose. With this support, you can match any string of characters in that position in the interface name. For example, so* matches all SONET/SDH interfaces.

extensive | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 77 on page 932 lists the output fields for the `show ppp interface` command. Output fields are listed in the approximate order in which they appear.

Table 77: show ppp interface Output Fields

Field Name	Field Description	Level of Output
`Session`	Name of the logical interface on which the session is running.	All levels
`Type`	Session type: PPP.	All levels
`Phase`	PPP process phase: `Authenticate`, `Pending`, `Establish`, `LCP`, `Network`, `Disabled`, and `Tunneled`.	All levels
`Session flags`	Special conditions present in the session: `Bundled`, `TCC`, `No-keepalives`, `Looped`, `Monitored`, and `NCP-only`.	All levels
protocol `State`	Protocol state information. See specific protocol state fields for information.	None specified
`AUTHENTICATION`	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the `Authentication` field description for further information.	None specified
`Keepalive settings`	Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS-based PPP sessions are supported only on service interfaces (si). • `Interval`—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the `interval` and `Down-count` values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • `Up-count`—The number of keepalive packets a destination must receive to change a link’s status from down to up. • `Down-count`—The number of keepalive packets a destination must fail to receive before the network takes down a link.	`extensive`

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Magic-Number validation`	Indicates whether the local peer is configured to ignore mismatches between peer magic numbers when the numbers are validated during PPP keepalive (Echo-Request/Echo-Reply) exchanges. • Enable–Mismatch detection sends failed Echo-Reply packets to the Routing Engine. If a valid magic number is not received within the configurable keepalive interval, PPP treats this as a keepalive failure and tears down the PPP sessions. • Disable–The Packet Forwarding Engine does not perform a validation check for magic numbers received from remote peers. A mismatch cannot be detected, so receipt of its own magic number or an unexpected value does not trigger notification to the Routing Engine.	`extensive`
`RE Keepalive statistics`	Keepalive statistics for the packets handled by the Routing Engine. • `LCP echo req Tx`—LCP echo requests sent from the Routing Engine. • `LCP echo req Rx`—LCP echo requests received at the Routing Engine. • `LCP echo rep Tx`—LCP echo responses sent from the Routing Engine. • `LCP echo rep Rx`—LCP echo responses received at the Routing Engine. • `LCP echo req timeout`—Number of keepalive packets where the keepalive aging timer has expired. • `LCP Rx echo req Magic Num Failures`—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • `LCP Rx echo rep Magic Num Failures`—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match.	`extensive`

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
LCP	LCP information: • State—LCP protocol state (all platforms except M120 and M320 routers): • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—LCP protocol state (M120 and M320 routers): • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down). • Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.	extensive

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	- `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. - `Last started`—LCP state start time. - `Last completed`—LCP state completion time. - `Last updated`—Reports the timestamp of the last successful connection update exchange. 1. When LCP negotiation completes, this field has the same value as the `Last completed` field. 2. The field then reports the timestamp of any subsequent successful exchange of Connection-Update-Request and Connection-Update-Ack messages with the peer (such as a home gateway). This field is displayed only when the Connection-Status-Message option is successfully negotiated.	

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	- `Negotiated options:` - `ACFC`—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields. - `Asynchronous map`—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link. - `Authentication protocol`—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required. - `Authentication algorithm`—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported. - `Connection Update Requests`—Number of connection update requests sent by PPP to the remote peer (such as a home gateway). This value does not include retries. This field is displayed even when negotiation fails for the Connection-Status-Message option. This enables you to confirm that an update request was sent. The absence of the Juniper Connection Status Message field indicates the peer does not support the updates. - `Endpoint discriminator class`—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link. - `Juniper Connection Status Message`—The content of the Connection-Status-Message VSA (26-4874–218) most recently received from RADIUS.	

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	This field is displayed only when the Connection-Status-Message option is successfully negotiated. • `Magic number`—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated. • `MRU`—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets. • `MRRU`—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets. • `Multilink header suspendable classes`—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given. • `Multilink header format classes`—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number. • `PFC`—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field. • `short sequence`—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers.	

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Authentication`	CHAP or PAP authentication state information. For CHAP authentication: • `Chap-ans-rcvd`—Packet was sent from the peer, indicating that the peer received the `Chap-resp-sent` packet. • `Chap-ans-sent`—Packet was sent from the authenticator, indicating that the authenticator received the peer's `Chap-resp-rcvd` packet. • `Chap-chal-rcvd`—Challenge packet has been received by the peer. • `Chap-chal-sent`—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • `Chap-resp-rcvd`—CHAP response packet has been received by the authenticator. • `Chap-resp-sent`—CHAP response packet has been sent to the authenticator. • `Closed`—Link is not available for authentication. • `Failure`—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • `Success`—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. For PAP authentication: • `Pap-resp-sent`—PAP response sent to peer (ACK/NACK)t. • `Pap-req-rcvd`—PAP request packet received from peer. • `Pap-resp-rcvd`—PAP response received from the peer (ACK/NACK). • `Pap-req-sent`—PAP request packet sent to the peer. • `Closed`—Link is not available for authentication.	None specified

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	- `Failure`—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. - `Success`—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.	

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPCP	Internet Protocol Control Protocol (IPCP) information. • State—(All platforms except M120 and M320 routers) One of the following values: • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down).	extensive

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	- Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. - Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—IPCP state start time. • Last completed—IPCP state authentication completion time. • Negotiated options: - compression protocol—Negotiate the use of a specific compression protocol. By default, compression is not enabled. - local address—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. - primary DNS server—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link. - primary WINS server—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link. - remote address—IP address of the remote end of the link in dotted quad notation. - secondary DNS server—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link. - secondary WINS server—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: Active or Passive	

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
IPV6CP	Internet Protocol version 6 Control Protocol (IPv6CP) information. • State—(All platforms except M120 and M320 routers) One of the following values: • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is not available for traffic. • Opened—Link is administratively available for traffic. • Req-sent—An attempt has been made to configure the connection. • State—(M120 and M320 routers) One of the following values: • Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. • Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. • Closed—Link is available (up), but no Open has occurred. • Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. • Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. • Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down).	extensive

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	- `Stopped`—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. - `Stopping`—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. - `Last started`—IPV6CP state start time. - `Last completed`—IPV6CP state authentication completion time. - `Negotiated options:` - `local interface identifier`—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address. - `remote interface identifier`—IP address of the remote end of the link in dotted quad notation. - `Negotiation mode`—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: `Active` or `Passive`	
`OSINLCP State`	OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers): - `State:` - `Ack-rcvd`—Configure-Request has been sent and Configure-Ack has been received. - `Ack-sent`—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received. - `Closed`—Link is not available for traffic. - `Opened`—Link is administratively available for traffic. - `Req-sent`—Attempt has been made to configure the connection. - `Last started`—OSINLCP state start time. - `Last completed`—OSINCLP state completion time.	`extensive`

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
TAGCP	TAGCP information. - State—(All platforms except M120 and M320 routers) One of the following values: - Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. - Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. - Closed—Link is not available for traffic. - Opened—Link is administratively available for traffic. - Req-sent—An attempt has been made to configure the connection. - State—(M120 and M320 routers) One of the following values: - Ack-rcvd—A Configure-Request has been sent and a Configure-Ack has been received. - Ack-sent—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received. - Closed—Link is available (up), but no Open has occurred. - Closing—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. - Opened—Link is administratively available for traffic. A Configure-Ack has been both sent and received. - Req-sent—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received. - Starting—An administrative Open has been initiated, but the lower layer is still unavailable (Down).	extensive none

Table 77: show ppp interface Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	◦ Stopped—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack. ◦ Stopping—A Terminate-Request has been sent but a Terminate-Ack has not yet been received. • Last started—TAGCP state start time. • Last completed—TAGCP state authentication completion time.	

Sample Output

show ppp interface

```
user@host> show ppp interface up
Sessions for interface up
  Session up:green-arrow:pp0.3221225473, Type: PPP, Phase: Network
    LCP State: Opened
    Authentication: PAP State: Grant
    IPCP State: Opened
```

show ppp address-pool

IN THIS SECTION

Syntax

```
show ppp address-pool pool-name
<detail>
```

Description

Display PPP address pool information.

Options

pool-name	Address pool name.
detail	(Optional) Display detailed address pool information.

Required Privilege Level

view

Output Fields

Table 78 on page 947 lists the output fields for the `show ppp address-pool` command. Output fields are listed in the approximate order in which they appear.

Table 78: show ppp address-pool Output Fields

Field Name	Field Description	Level of Output
Address pool	Trace address pool code.	All levels
Address range	Range of sequentially ordered IP addresses contained in the address pool.	detail
Number of assigned addresses	Fixed IP address that is to be given to remote users when they dial in. This is a host-only IP address (subnet mask is 255.255.255.255) and is only for single connection receiver profiles.	All levels
Number of addresses configured	Number of IP addresses that are available for allocation and used by PPP sessions.	All levels
Assigned addresses	Addresses assigned to PPP sessions from the address pool.	detail

Sample Output

show ppp address-pool

```
user@host> show ppp address-pool
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
  Number of assigned addresses: 0
  Number of addresses configured: 10
```

show ppp address-pool detail

```
user@host> show ppp address-pool ppp1 detail
Address pool ppp1
  Address range: 203.0.113.221 - 203.0.113.230
  Number of assigned addresses: 2
  Number of addresses configured: 10
  Assigned addresses:
      203.0.113.221
      203.0.113.222
```

show ppp statistics

IN THIS SECTION

Syntax

```
show ppp statistics
<detail>
<memory>
<recovery>
```

Description

Display PPP interface statistics information.

Options

detail (Optional) Display the detailed statistics.

memory (Optional) Display PPP process memory statistics.

recovery (Optional) Display recovery state of PPP after a GRES or restart. It is safe to force another GRES or restart only when the recovery state indicates the recovery is done.

> **NOTE**: When you issue this command option during the recovery process, the command may time out or fail silently rather than display output. Recovery is not complete until the command displays `Recovery state: recovery done`.

Required Privilege Level

view

Output Fields

Table 79 on page 949 lists the output fields for the `show ppp statistics` command. Output fields are listed in the approximate order in which they appear.

Table 79: show ppp statistics Output Fields

Field Name	Field Description	Level of Output
`Total sessions`	Number of PPP sessions on an interface.	none `detail`

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sessions in disabled phase`	Number of PPP sessions disabled. Number of sessions where the link is either administratively or physically down. Once the PPP process learns from the kernel that Layer 2 is ready to send and receive traffic, it will do a phase transition from disabled to established. When LCP and NCP transitions through states, links transition to the establish phase when terminate packets are exchanged or some other failure, such as authentication or expiration of a timer occurs.	none `detail`
`Sessions in establish phase`	Number of PPP sessions in establish phase. In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure and test the data link.	none `detail`
`Sessions in authenticate phase`	Number of PPP sessions in authenticate phase. Each end of the PPP link must first send LCP packets to configure the data link during the link establishment phase. After the link has been established, PPP provides for an optional authentication phase before proceeding to the Network-Layer Protocol (NLP) phase.	none `detail`
`Sessions in network phase`	Number of PPP sessions in the network phase. After a link has been established and optional facilities have been negotiated as needed by the LCP, PPP must send Network Control Protocol (NCP) packets to choose and configure one or more network-layer protocols, such as IP, IPX, or AppleTalk. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.	none `detail`
`Bundles in pending phase`	Number of unique bundles to which PPP links are referring.	none `detail`

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type	Type of structure for which memory is allocated. • Queued rtsock msgs—Queued route socket messages. When a PPP process is unable to send a route socket message to the kernel (typically because of congestion of the route socket interface), the message is queued for deferred processing. • PPP session—Active PPP session. Stores all the information for a PPP session, such as authentication, sequence number, LCP session, and NCP session information. • Interface address—Interface address associated with a PPP connection. Stores the information about the interface address that PPP obtains from the kernel. • Destination profile—Stores the destination profile information associated with an interface address. • ML link settings—Stores information about an MLPPP link, such as the bundle name and compressed real-time transport protocol (CRTP) settings. • IPCP blocked address—When addresses are blocked in an address pool (for example, when the interface address is within the range of an address pool, it will be implicitly blocked), this structure is used to store the address in the pool. • PPP session trace—A PPP session trace is allocated for record keeping for each session listed at the [set protocols ppp monitor-session] hierarchy level. • IFL redundancy state—Stores redundancy state information needed for high availability (HA) operation. • Protocol family—Stores the information about the protocol family that PPP obtains from the kernel.	detail

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Type (continued)	• ML bundle settings—Multilink bundle settings. Stores the context information for a MLPPP bundle. • PPP LCP session—PPP Link Control Protocol session, used for establishing, configuring, and testing the data-link connection. Stores the information for an LCP session, such as negotiated options, current state, and statistics. • PPP NCP session—PPP Network Control Protocol (NCP) phase in the PPP link connection process. Stores the information for an NCP session, such as negotiated options, current state, address family, and statistics. • Physical interface—Stores the information about the physical interface that PPP obtains from the kernel. • Access profile—Stores the information found at the [edit access profile] hierarchy level for each profile. • ML wait entry—Created when there are MLPPP links joining a bundle. before its addition to the PPP process. Links are saved here, and when the bundle is added, are properly assigned to the bundle. • Group profile—Stores information set in the PPP stanza of a group profile, such as the primary and secondary Domain Name System (DNS), primary and secondary NDNS, and address pool name. • Profile client—Stores the per-client information of the access profile (information obtained from the [set access profile *name* client *client-name*] hierarchy level. • PPP Auth session—PPP authentication session. Stores all the session-specific authentication protocol parameters. • Logical interface—Stores the information about the logical interface that PPP obtains from the kernel. • Non-tagged—Generic catch-all for allocations not of a particular structure type.	detail

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Type`	If you specify the `memory` keyword, the following memory statistics are displayed for Ethernet interfaces on M120 and M320 routers. • `authenticate`—Stores information common to all PPP authentication protocols. • `linkInterface`—Stores information about PPP link interfaces. • `pap`—Stores information about PPP PAP authentication protocol. Includes authenticator and authenticate state machines. • `lcp`—PPP Link Control Protocol session. Used for establishing, configuring and testing the data-link connection. Stores information for LCP session, such as negotiated options, state, and statistics. • `chap`—Stores information about PPP CHAP authentication protocol. Includes authenticator and authenticate state machines. • `eapBuffer`—Stores runtime authentication information for EAP. • `eap`—Stores information about PPP EAP authentication protocol. Includes authenticator and authenticate state machines. • `authNone`—Stores information about no PPP authentication. Includes the authenticator state machine. • `networkInterface`—Stores information about NCP portions of PPP protocol. • `ipNcp`—PPP IPCP session information. Used for configuring, negotiating, and establishing IPCP protocol. Stores the current state, and configured and negotiated options. • `ipv6Ncp`—PPP IPv6CP session information. Used for configuring, negotiating, and establishing IPv6CP protocol. Stores the current state, and configured and negotiated options. • `osiNcp`—PPP OSICP session information. Used for configuring, negotiating, and establishing OSICP protocol. Stores the current state, and configured and negotiated options.	`memory`

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• mplsNcp—PPP MPLSCP session information. Used for configuring, negotiating, and establishing MPLSCP protocol. Stores the current state. • trace—Stores information for PPP debugging.	
Total	Total memory allocations.	detail
Size	Size of the structure.	detail
Active	Number of instances of the structure that are used.	detail
Free	Number of instances of the structure that are on the free list. Types with a number in the Free column are pooled structures, and are typically types that are often used.	detail
Limit	Maximum number of instances that can be on the free list. Types with a number in the Limit column are pooled structures, and are typically types that are often used.	detail
Total size	Total amount of memory being used by a type of structure (includes active and free instances).	detail
Requests	Number of allocation requests made by a type of structure.	detail
Failures	Number of failed allocations.	detail
Recovery state	State of PPP recovery after a GRES or restart: • recovery done—All sessions have recovered; it is safe to force another GRES or restart. • recovery cleanup pending—Not all PPP sessions have recovered; it is not safe to force another GRES or restart.	none

Table 79: show ppp statistics Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Subscriber sessions pending retention`	Number of PPP subscriber sessions that are in the process of being recovered.	none
`Subscriber sessions recovered OK`	Number of PPP subscriber sessions that have recovered after a GRES or restart.	none
`Subscriber sessions recovery failed`	Number of PPP subscriber sessions that have failed to recover after a GRES or restart.	none

Sample Output

show ppp statistics

```
user@host> show ppp statistics
Session statistics from PPP universal edge process
  Total subscriber sessions: 1
    Subscriber sessions in disabled phase    : 0
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 1
```

show ppp statistics detail

```
user@host> show ppp statistics detail
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0
Type                   Size  Active  Free  Limit  Total size  Requests Failures
Queued rtsock msgs       28      0      0  65535           0         0
PPP session              60      0                         0         0
Interface address        64      0      0  65535           0         0
Destination profile      65      0                         0         0
ML link settings         68      0                         0         0
IPCP blocked address     68      0                         0         0
PPP session trace        76      0                         0         0
IFL redundancy state     76      0                         0         0
Protocol family          84      0      0  65535           0         0
ML bundle settings      108      0                         0         0
PPP LCP session         120      0                         0         0
PPP NCP session         124      0                         0         0
Physical interface      124    170      0  65535       21080       170
Access profile          132      0                         0         0
ML wait entry           144      0      0     20           0         0
Group profile           164      0                         0         0
Profile client          272      0                         0         0
PPP Auth session        356      0                         0         0
Logical interface       524      0      0  65535           0         0
Non-tagged                                                 8         2
Total                                                  21088       172       0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0
Type                   Size  Active  Free  Limit  Total size Requests Failures
authenticate            224       1    99  16384         224        0        0
linkInterface           152       1    99  16384         152        0        0
pap                     256       1    99  16384         256        0        0
lcp                     272       1    99  16384         272        0        0
chap                    284       0     0  16384           0        0        0
eapBuffer              1464       0     0  16384           0        0        0
eap                     276       0     0  16384           0        0        0
authNone
networkInterface        220       1    99  16384         220        0        0
ipNcp                   256       1    99  16384         256        0        0
ipv6Ncp                 204       0     0  16384           0        0        0
```

```
osiNcp                  192       0     0  16384           0        0        0
mplsNcp                 188       0     0  16384           0        0        0
trace                  2052       0    16      16           0        0        0
Total                                                   1380        0        0
```

show ppp statistics recovery (Safe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery done
    Subscriber sessions recovered OK      : 32001
    Subscriber sessions recovery failed   : 0
```

show ppp statistics recovery (Unsafe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery cleanup pending
    Subscriber sessions pending retention : 32001
    Subscriber sessions recovered OK      : 0
    Subscriber sessions recovery failed   : 0
```

show ppp summary

IN THIS SECTION

Syntax

```
show ppp summary
```

Description

Display PPP session summary information.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 80 on page 958 lists the output fields for the `show ppp summary` command. Output fields are listed in the approximate order in which they appear.

Table 80: show ppp summary Output Fields

Field Name	Field Description
`Interface`	Interface on which the PPP session is running. An interface type of pp0 indicates an Ethernet interface type on a M120 or M320 router.
`Session type`	Type of session: `PPP` or `Cisco-HDLC`.

Table 80: show ppp summary Output Fields *(Continued)*

Field Name	Field Description
Session phase	PPP process phases: Authenticate, Pending, Establish, Network, Disabled.
Session flags	Special conditions present in the session, such as Bundled, TCC, No-keepalives, Looped, Monitored, and NCP-only.

Sample Output

show ppp summary

```
user@host> show ppp summary
Interface                         Session type  Session phase   Session flags
up:green-arrow:pp0.3221225473 PPP           Network
```

show services captive-portal-content-delivery

IN THIS SECTION

Syntax

```
show services captive-portal-content-delivery
<pic pic-name>
<profile profile-name>
<rule rule-name> <term term-name>
<ruleset ruleset-name>
<sset sset-name> <brief> <detail> <summary>
<statistics <interface pic-name>>
```

Description

Display the current operational state of all captive portal interfaces.

Options

brief (Optional) Display brief service set database information.

detail (Optional) Display detailed service set database information.

pic Display the PIC database.

profile Display the profile database.

rule Display the rule database.

ruleset Display the rule set database.

sset Display the service set database.

statistics Display captive portal content delivery statistics about a PIC.

summary (Optional) Display a summary of service set database information.

term (Optional) Display term information for the rule database.

Required Privilege Level

view

Output Fields

Table 81 on page 961 lists the output fields for the `show services captive-portal-content-delivery` command. Output fields are listed in the approximate order in which they appear.

Table 81: show services captive-portal-content-delivery Output Fields

Field Name	Field Description	Level of Output
`Name`	Name of the interface.	none
`Index`		none
`Profile`	Name of the service profile for the HTTP redirect services that contains the rules or rule sets specifying the service.	none
`Rules or Rule Sets`	List of rules or rule sets contained in the HTTP redirect service profile.	none
`Rule Name`	Name of an HTTP redirect service rule.	none
`Term Name`	Name of a rule term.	none
`Rule match direction`	Traffic direction on the interface where the rule match is applied, `input`, `output`, or `input-output`.	none

Table 81: show services captive-portal-content-delivery Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Term action`	Action performed on packets when the rule term is matched: • Accept the packets. • Redirect the packets to a new destination URL. • Rewrite the packets with a new destination address and optionally a new destination port. • Log information about the packet to a system log file.	none
`Term action option`	Additional information related to the term action. • A new destination URL for a redirect action. • A new destination address for a rewrite action. • A new destination port for a rewrite action.	none
`Service Sets`	Name of service sets contained in a profile.	none
`Id`	Identifier number for a service set.	none
`Compiled Rules`		none
`service-set interface`	Interface on which the service set rules are applied.	none
`Packets received`	Number of packets received on the service-set interface.	none
`Packets altered`	Number of packets redirected or rewritten on the service-set interface.	none
`Packets dropped`	Number of packets dropped on the interface.	`detail`
`Received`	Number of packets received for the listed action: `Redirect`, `Rewrite`, or `Insert`.	`detail`

Table 81: show services captive-portal-content-delivery Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Altered`	Number of packets altered by the listed action: `Redirect` or `Rewrite`.	`detail`
`Redirected`	Number of packets redirected by the `Insert` action.	`detail`

Sample Output

show services captive-portal-content-delivery

```
user@host> show services captive-portal-content-delivery pic si-5/0/0
Name         Index
si-5/0/0        20
```

show services captive-portal-content-delivery (Profile)

```
user@host> show services captive-portal-content-delivery profile
Profile         Rules or Rule Sets
http-redirect   1
ipda-rewrite    1
```

show services captive-portal-content-delivery (Profile HTTP Redirect)

```
user@host> show services captive-portal-content-delivery profile http-redirect
Profile         Rules or Rule Sets
http-redirect   1
```

show services captive-portal-content-delivery (Profile IPDA Rewrite)

```
user@host> show services captive-portal-content-delivery profile ipda-rewrite
Profile         Rules or Rule Sets
ipda-rewrite    1
```

show services captive-portal-content-delivery (Rules)

```
user@host> show services captive-portal-content-delivery rule
Rule Name          Term Name
redirect           t2
rewrite            t1
```

show services captive-portal-content-delivery (Rewrite Term)

```
user@host> show services captive-portal-content-delivery rule rewrite term t1
Rule name: rewrite
Rule match direction: input
Term name: t1
Term action: rewrite
Term action option: null
```

show services captive-portal-content-delivery (Redirect Term)

```
user@host> show services captive-portal-content-delivery rule redirect term t2
Rule name: redirect
Rule match direction: input
Term name: t2
Term action: redirect
Term action option: http://www.example.net
```

show services captive-portal-content-delivery (Service Set Detail)

```
user@host> show services captive-portal-content-delivery sset sset1 detail
Service Set        Id       Profile         Compiled Rules
sset1              1        ipda-rewrite    1
```

show services captive-portal-content-delivery (Interface)

```
user@host> show services captive-portal-content-delivery statistics interface sis-5/0/0
service-set interface: si-5/0/0

Packets received   Packets altered
5                     3
```

Release Information

Command introduced in Juniper BNG CUPS Release 23.1R1.

show services l2tp client

IN THIS SECTION

Syntax

```
show services l2tp client
<client-name>
```

Description

Display information about all L2TP clients or a specific L2TP client.

Options

client-name (Optional) Name of a client.

Required Privilege Level

view

Output Fields

Table 82 on page 966 lists the output fields for the `show services l2tp client` command. Output fields are listed in the approximate order in which they appear.

Table 82: show services l2tp client Output Fields

Field Name	Field Description

`Client`
`Client Name`
	Name of the client.

Table 82: show services l2tp client Output Fields *(Continued)*

Field Name	Field Description
`Tunnels`	Number of tunnels in the tunnel group.
`Sessions`	Number of L2TP sessions established for tunnels in the tunnel group.
`Tunnel-group`	Name of a tunnel group to which the client belongs.
`Session-limit-group`	Name of a session-limit group to which the client belongs.

Sample Output

show services l2tp client

```
user@host> show services l2tp client
Client           Tunnels    Sessions      Tunnel-group         Session-limit-group
entA-serviceA    2          20            l2tp-tunnel-group1   enterpriseA
entA-serviceB    3          120           l2tp-tunnel-group2   enterpriseB
```

show services l2tp client (Client Name)

```
user@host> show services l2tp client entA-serviceA
Client Name      Tunnels    Sessions    Tunnel-group          Session-limit-group
entA-serviceA    2          20          l2tp-tunnel-group1    enterpriseA
```

show services l2tp destination

IN THIS SECTION

Syntax

```
show services l2tp destination
<brief | detail | extensive>
<local-gateway gateway-address>
<peer-gateway gateway-address>
<statistics>
```

Description

Display information about L2TP tunnel destinations.

Options

brief | detail | extensive — (Optional) Display the specified level of information.

local-gateway *gateway-address* — (Optional) Display L2TP session information for only the specified local gateway address.

peer-gateway *gateway-address*	(Optional) Display L2TP session information for only the specified peer gateway address.
statistics	(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.

Required Privilege Level

view

Output Fields

Table 83 on page 969 lists the output fields for the `show services l2tp destination` command. Output fields are listed in the approximate order in which they appear.

Table 83: show services l2tp destination Output Fields

Field Name	Field Description	Level of Output
`Local Name`	Name of this destination.	All levels
`Remote IP`	IP address of the remote peer (LNS).	All levels
`Tunnels`	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed

Table 83: show services l2tp destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sessions`	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed
`State`	Administrative state of the L2TP destination: • `Enabled`—No restrictions exist on creation or operation of sessions and tunnels for this destination. • `Disabled`—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the `Disabled` state. • `Drain`—Creation of new sessions and tunnels is disabled for this destination.	All levels
`Local IP`	IP address of the local gateway (LAC).	`detailextensive`
`Transport`	Medium used for tunneling. Only `ipUdp` is supported.	`detailextensive`
`Logical System`	Logical system in which the tunnel is configured.	`detailextensive`
`Router Instance`	Routing instance in which the tunnel is configured.	`detailextensive`
`Lockout State`	Reachability state of the destination: • `not locked`—Destination is considered reachable. • `waiting for lockout timeout`—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	`detailextensive`

Table 83: show services l2tp destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Access Line Information`	State of the LAC per-destination configuration for forwarding subscriber line information to the LNS, `Enabled` or `Disabled`.	`detailextensive`
`Speed Updates`	State of the LAC per-destination configuration for including connection speed updates when it forwards subscriber line information to the LNS, `Enabled` or `Disabled`.	`detailextensive`
`Connections`	Number of total, active, and failed tunnel and session connections for the destination.	`extensive`
`Control Tx`	Amount of control information transmitted, in packets and bytes.	`statistics`
`Control Rx`	Amount of control information received, in packets and bytes.	`statistics`
`Data Tx`	Amount of data transmitted, in packets and bytes.	`statistics`
`Data Rx`	Amount of data received, in packets and bytes.	`statistics`
`Error Tx`	Number of errors transmitted, in packets.	`statistics`
`Error Rx`	Number of errors received, in packets.	`statistics`

Sample Output

show services l2tp destination

```
user@host> show services l2tp destination
  Local Name     Remote IP         Tunnels         Sessions  State
  1              203.0.113.101     1               1         Enabled
```

show services l2tp destination detail

```
user@host> show services l2tp destination detail
 Local name: 1
    Remote IP: 203.0.113.101
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Enabled
 Local name: 1
    Remote IP: 203.0.113.108
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 203.0.113.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: waiting for lockout timeout
    Access Line Information: Enabled, Speed Updates: Enabled
```

show services l2tp destination extensive (LAC)

```
user@host> show services l2tp destination extensive
  Local name: 1
    Remote IP: 203.0.113.101
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Enabled
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination extensive (LNS)

```
user@host> show services l2tp destination extensive
  Local name: 3
    Remote IP: 203.0.113.103
    State: Enabled
    Local IP: 203.0.113.102
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: Enabled, Speed Updates: Disabled
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination statistics (LAC only on MX Series Routers)

```
user@host> show services l2tp destination statistics
 Local name: 2, Tunnels: 1, Sessions: 210
                      Packets        Bytes
    Control Tx            680        63.3k
    Control Rx            283        10.6k
    Data Tx              1129        14.3k
    Data Rx               877        10.9k
    Errors Tx               0
    Errors Rx               0
```

show services l2tp destination lockout

IN THIS SECTION

Syntax

```
show services l2tp destination lockout
```

Description

Display a list of destinations that are currently locked out and the time remaining for each to remain in the lockout state.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 84 on page 974 lists the output fields for the `show services l2tp destination lockout` command. Output fields are listed in the approximate order in which they appear.

Table 84: show services l2tp destination lockout Output Fields

Field Name	Field Description
`Destination`	Name of the destination.
`Time Remaining`	Time remaining for the destination to be locked out.

Table 84: show services l2tp destination lockout Output Fields *(Continued)*

Field Name	Field Description
L2TP lockout destinations found	Total count of lockout destinations.

Sample Output

show services l2tp destination lockout

```
user@host> show services l2tp destination lockout
  Destination    Time Remaining
  4              45
  5              43
  6              8
3 L2TP lockout destinations found
```

show services l2tp session

IN THIS SECTION

Syntax

```
show services l2tp session
<brief | detail | extensive>
<interface interface-name>
<local-gateway gateway-address>
<local-gateway-name gateway-name>
<local-session-id session-id>
<local-tunnel-id tunnel-id>
<peer-gateway gateway-address>
<peer-gateway-name gateway-name>
<statistics>
<tunnel-group group-name>
```

Description

Display information about active L2TP sessions for LAC and LNS.

Options

none

Display standard information about all active L2TP sessions.

brief | detail | extensive

(Optional) Display the specified level of output.

interface *interface-name*

(Optional) Display L2TP session information for only the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- `si-`*fpc*/*pic*/*port*— MPCs on MX Series routers only.

local-gateway *gateway-address*

(Optional) Display L2TP session information for only the specified local gateway address.

local-gateway-name *gateway-name*

(Optional) Display L2TP session information for only the specified local gateway name.

local-session-id *session-id*

(Optional) Display L2TP session information for only the specified local session identifier.

local-tunnel-id *tunnel-id*

(Optional) Display L2TP session information for only the specified local tunnel identifier.

peer-gateway *gateway-address*

(Optional) Display L2TP session information for only the specified peer gateway address.

peer-gateway-name *gateway-name*

(Optional) Display L2TP session information for only the specified peer gateway name.

statistics

(Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.

tunnel-group *group-name*

(Optional) Display L2TP session information for only the specified tunnel group. To display information about L2TP CPU and memory usage, you can include the tunnel group name in the `show services service-sets memory-usage` *`group-name`* and `show services service-sets cpu-usage` *`group-name`* commands. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Output Fields

Table 85 on page 977 lists the output fields for the `show services l2tp session` command. Output fields are listed in the approximate order in which they appear.

Table 85: show services l2tp session Output Fields

Field Name	Field Description	Level of Output
`Interface`	(LNS only) Name of an adaptive services interface.	All levels

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tunnel group`	(LNS only) Name of a tunnel group.	All levels
`Tunnel local ID`	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
`Session local ID`	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	All levels
`Session remote ID`	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
`State`	State of the L2TP session: • `Established`—Session is operating. This is the only state supported for the LAC. • `closed`—Session is being closed. • `destroyed`—Session is being destroyed. • `clean-up`—Session is being cleaned up. • `lns-ic-accept-new`—New session is being accepted. • `lns-ic-idle`—Session has been created and is idle. • `lns-ic-reject-new`—New session is being rejected. • `lns-ic-wait-connect`—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
`Bundle ID`	(LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank `Bundle` field are not participating in the Multilink Protocol. Sessions in a multilink bundle might belong to different L2TP tunnels. For L2TP output organized by bundle ID, issue the `show services l2tp multilink extensive` command.	All levels

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Mode	(LNS) Mode of the interface representing the session: shared or exclusive. (LAC) Mode of the interface representing the session: shared or dedicated. Only dedicated is currently supported for the LAC.	extensive
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	extensive
Remote IP	IP address of remote endpoint of the PPP session.	extensive
Username	(LNS only) Name of the user logged in to the session.	All levels
Assigned IP address	(LNS only) IP address assigned to remote client.	extensive
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	extensive
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	extensive
Local MRU	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	extensive
Remote MRU	(LNS only) MRU setting of the remote device, in bytes.	extensive

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tx speed`	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived. Either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed: • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. When the Tx connect speed method is set to none, the value of zero (0) is displayed.	`extensive`
`Rx speed`	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived. Either the initial (`initial`) line speed or both the initial and current (`update`) line speeds can be displayed : • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. When the Tx connect speed method is set to none, the value of zero (0) is displayed.	`extensive`

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Bearer type`	Type of bearer enabled: • `0`—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • `1`—Digital access requested. • `2`—Analog access requested. • `4`—Asynchronous Transfer Mode (ATM) bearer support.	`extensive`
`Framing type`	Type of framing enabled: • `1`—Synchronous framing • `2`—Asynchronous framing	`extensive`
`LCP renegotiation`	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: `On` or `Off`.	`extensive`
`Authentication`	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	`extensive`
`Interface ID`	(LNS only) Identifier used to look up the logical interface for this session.	`extensive`
`Interface unit`	Logical interface for this session.	All levels
`Call serial number`	Unique serial number assigned to the call.	`extensive`
`Policer bandwidth`	Maximum policer bandwidth configured for this session.	`extensive`

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Policer burst size`	Maximum policer burst size configured for this session.	`extensive`
`Firewall filter`	Configured firewall filter name.	`extensive`
`Session encapsulation overhead`	Overhead allowance configured for this session, in bytes.	`extensive`
`Session cell overhead`	Cell overhead activation (`On` or `Off`).	`extensive`
`Create time`	Date and time when the call was created.	`extensive`
`Up time`	Length of time elapsed since the call became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	`extensive`

Table 85: show services l2tp session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets. • `LCP echo req Tx`—Number of LCP echo requests transmitted, in packets. • `LCP echo req Rx`—Number of LCP echo requests received, in packets. • `LCP echo rep Tx`—Number of LCP echo responses transmitted, in packets. • `LCP echo rep Rx`—Number of LCP echo responses received, in packets. • `LCP echo Req timout`—Number of LCP echo requests that timed out. • `LCP echo Req error`—Number of errors received for LCP echo packets. • `LCP echo Rep error` —Number of errors transmitted for LCP echo packets.	`extensive`

Sample Output

show services l2tp session (LNS on MX Series Routers)

```
user@host> show services l2tp session
Tunnel local ID: 40553
  Local  Remote   State                Interface        Interface
  ID     ID                            unit             Name
  17967  1        Established          1073749824       si-5/2/0
```

show services l2tp session (LAC)

```
user@host> show services l2tp session
Tunnel local ID: 31889
  Local  Remote   State                Interface        Interface
  ID     ID                            unit             Name
   31694      1   Established                311        pp0
```

show services l2tp session detail (LAC)

```
user@host> show services l2tp session detail
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1, Interface unit: 311
    State: Established, Interface: pp0, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns
```

show services l2tp session extensive (LAC)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 0, Rx speed: 0
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LAC on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.102:1701, Remote IP: 203.0.113.101:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 256000, source service-profile
    Rx speed: 128000, source ancp
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LNS on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 40553
  Session local ID: 17967, Session remote ID: 1
    Interface unit: 1073749824
    State: Established
    Interface: si-5/2/0
    Mode: Dedicated
    Local IP: 192.0.2.2:1701, Remote IP: 192.0.2.3:1701
    Local name: lns-mx960, Remote name: testlac
```

```
      Tx speed: initial 64000, Update 256000
      Rx speed: initial 64000, Update 256000
      Bearer type: 2, Framing type: 1
      LCP renegotiation: Off, Authentication: None
      Call serial number: 1
      Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:48
      Idle time: N/A
      Statistics since: Mon Apr 25 20:27:50 2011
                          Packets        Bytes
        Control Tx              4          219
        Control Rx              4          221
        Data Tx                 0            0
        Data Rx                10          228
        Errors Tx               0
        Errors Rx               0
```

show services l2tp session statistics (MX Series Routers)

```
user@host>show services l2tp session statistics local session-id 1
Tunnel local ID: 17185
Session local ID: 1, Session remote ID: 14444, Interface unit: 1073788352
State: Established
Statistics since: Mon Aug 1 13:27:47 2011
                Packets          Bytes
Data Tx         4             51
Data Rx         3             36
```

show services l2tp session-limit-group

IN THIS SECTION

Syntax

```
show services l2tp session-limit-group
<limit-group-name>
```

Description

Display information about all session-limit groups or a specific session limit group.

Options

limit-group-name	(Optional) Name of a session-limit group.

Required Privilege Level

view

Output Fields

Table 86 on page 988 lists the output fields for the `show services l2tp session-limit-group` command. Output fields are listed in the approximate order in which they appear.

Table 86: show services l2tp session-limit-group Output Fields

Field Name	Field Description
`Session-limit-group`	Name of a session-limit group.
`Tunnels`	Number of tunnels associated with the session-limit group in the tunnel group.
`Sessions`	Number of L2TP sessions established for session-limit group.
`Maximum limit`	Maximum number of sessions allowed for the session-limit group.

Sample Output

show services l2tp session-limit-group

```
user@host> show services l2tp session-limit-group

Session-limit-group     Tunnels       Sessions       Maximum limit
enterpriseA             2             10             1000
enterpriseB             10            120            2000
```

show services l2tp session-limit-group (Limit Group Name)

```
user@host> show services l2tp session-limit-group enterpriseA

Session-limit-group     Tunnels       Sessions       Maximum limit
enterpriseA             2             10             1000
```

show services l2tp summary

IN THIS SECTION

Syntax

```
show services l2tp summary
<statistics>
```

Description

Display Layer 2 Tunneling Protocol (L2TP) summary information.

Options

none — Display complete L2TP summary information. For LNS, displays L2TP summary information for all inline services interfaces.

statistics — (Optional) Display a summary of control packets and bytes transmitted and received.

Required Privilege Level

view

Output Fields

Table 87 on page 990 lists the output fields for the `show services l2tp summary` command. Output fields are listed in the approximate order in which they appear.

Table 87: show services l2tp summary Output Fields

Field Name	Field Description
`Administrative state`	Administrative state of the tunnel is drain. In this state you cannot configure new sessions, destinations, or tunnels at the LAC or LNS.
`Failover within a preference level`	State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level.
`Weighted load balancing`	State of this tunnel selection method on the LAC. When enabled, the maximum session limit of a tunnel determines its weight within a preference level. Tunnel selection proceeds from greatest to least weight. When disabled, selection defaults to a round robin method.
`Destination equal load balancing`	State of this tunnel selection method on the LAC. When enabled, the LAC selects tunnels based on the session count for destinations and the tunnel session count.
`Tunnel authentication challenge`	State of tunnel authentication, indicating whether the LAC and LNS exchange an authentication challenge and response during the establishment of the tunnel. The state is `Enabled` when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is `Disabled` when the secret is not present.

Table 87: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Calling number avp`	When the state is `Enabled`, the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is `Disabled`, the attribute is not sent to the LNS.
`Failover Protocol`	When the state is enabled, the LAC operates in the default *failover-protocol-fall-back-to-silent-failover* manner. When the state is disabled, the `disable-failover-protocol` statement has been issued and the LAC operates only in silent failover mode.
`Tx connect speed method`	The connection speed method configured to send the speed values in the L2TP Tx Connect Speed (AVP 24) and L2TP Rx Connect Speed (AVP 38). Possible values are: • `ancp` • `none` • `pppoe-ia-tag` • `service-profile` • `static` This is the default value.
`Rx speed avp when equal`	Indicates if the Rx connect speed when equal configuration is `enabled` or `disabled`.
`Tunnel assignment id`	Format of the tunnel name. Format of the tunnel name, based on RADIUS attributes returned from the AAA server: • `authentication-id`—Name consists of only Tunnel Assignment-Id [82]. This is the default value. • `client-server-id`—Name is a combination of Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82].

Table 87: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Tunnel Tx Address Change`	Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both: • `accept`—Accepts change requests for the IP address or UDP port. This is the default action. • `ignore`—Ignores all change requests. • `ignore-ip-address`—Ignores change requests for the IP address but accepts them for the UDP port. • `ignore-udp-port`—Ignores change requests for the UDP port but accepts them for the IP address.
`Min Retransmission Timeout for control packets`	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
`Min Retransmission Timeout for control packets`	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
`Max Retransmissions for Established Tunnel`	Maximum number of times control messages are retransmitted for established tunnels.
`Max Retransmissions for Not Established Tunnel`	Maximum number of times control messages are retransmitted for tunnels that are not established.
`Tunnel Idle Timeout`	Period that a tunnel can be inactive–that is, carrying no traffic–before it times out and is torn down.
`Destruct Timeout`	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
`Reassembly Service Set`	Indicates active IP reassembly configured for the interface.

Table 87: show services l2tp summary Output Fields *(Continued)*

Field Name	Field Description
`Destination Lockout Timeout`	Timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created.
`Access Line Information`	State of LAC global configuration for forwarding subscriber line information to the LNS, `Enabled` or `Disabled`. Indicates active IP reassembly configured for the interface. This information can also be displayed on the LNS for information it receives from the LAC.
`IPv6 Services for LAC Sessions`	State of LAC IPv6 service configuration for creating the IPv6 (inet6) address family for LAC subscribers, allowing the application of IPv6 firewall filters, `Enabled` or `Disabled`.
`Speed Updates`	State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, `Enabled` or `Disabled`. This information can also be displayed on the LNS for updates it receives from the LAC.
`Destinations`	Number of L2TP destinations for the LAC.
`Tunnels`	Number of L2TP tunnels established on the router.
`Sessions`	Number of L2TP sessions established on the router.
`Switched sessions`	Number of L2TP tunnel-switched sessions established on the router.
`Control`	Count of L2TP control packets and bytes sent and received.
`Data`	Count of L2TP data packets and bytes sent and received.
`Errors`	Count of L2TP error packets and bytes sent and received.

Sample Output

show services l2tp summary (LAC)

```
user@host> show services l2tp summary
Administrative state is Drain
      Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Enabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is enabled
  Tunnel Tx Address Change is Accept
  Min Retransmissions Timeout for control packets is 2 seconds
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
  Reassembly Service Set is ssnr3
  Access Line Information is Enabled, Speed Updates is Enabled
  IPv6 Services For LAC Sessions is Enabled
  Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

show services l2tp summary (LNS)

```
user@host show services l2tp summary
 Administrative state is Drain
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Disabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Enabled
  Tx Connect speed method is static
  reassembly Service Set is ssnr3
  Destinations: 4, Tunnels: 19, Sessions: 65, Switched sessions: 2
  Access Line Information is Enabled, Speed Updates is Enabled
```

show services l2tp summary statistics

```
user@host>show services l2tp summary statistics
Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is advisory
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 4 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 seconds
Destinations: 1, Tunnels: 1, Sessions: 31815, Switched sessions: 0
                    Tx packets    Rx packets  Memory (bytes)
  Control          90.4k           32.0k       245678080
  Data             127.3k          100.8kk                  0
  Errors                 0                      0
```

show services l2tp tunnel

IN THIS SECTION

Syntax

```
show services l2tp tunnel
<brief | detail | extensive>
<local-gateway gateway-address>
<local-gateway-name gateway-name>
<local-tunnel-id tunnel-id>
<peer-gateway gateway-address>
<peer-gateway-name gateway-name>
<statistics>
<tunnel-group group-name>
```

Description

Display information about L2TP tunnels for LAC and LNS; the tunnels may or may not have active sessions.

Options

none
Display standard information about all active L2TP tunnels.

brief | detail | extensive
(Default) Display the specified level of output.

local-gateway *gateway-address*
(Optional) Display L2TP tunnel information for only the specified local gateway address.

local-gateway-name *gateway-name*
(Optional) Display L2TP tunnel information for only the specified local gateway name.

local-tunnel-id *tunnel-id*
(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.

peer-gateway *gateway-address*
(Optional) Display L2TP tunnel information for only the specified peer gateway address.

peer-gateway-name *gateway-name*
(Optional) Display L2TP tunnel information for only the specified peer gateway name.

statistics
(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. The statistics for a tunnel are retained until the tunnel is disconnected, rather than until the last session in the tunnel is cleared. Retaining the statistics enables them to increment in the event a new session subsequently uses the tunnel. You cannot include this option with any of the level options, `brief`, `detail`, or `extensive`.

tunnel-group *group-name*
(Optional) Display L2TP tunnel information for only the specified tunnel group.

Required Privilege Level

view

Output Fields

"show services l2tp tunnel" on page 995 lists the output fields for the `show services l2tp tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 88: show services l2tp tunnel Output Fields

Field Name	Field Description
`Interface`	(LNS only) Name of an adaptive services interface.
`Tunnel group`	(LNS only) Name of a tunnel group.

Table 88: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`Local ID`	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.
`Remote ID`	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.
`Remote IP`	IP address of the peer endpoint of the tunnel.
`Sessions`	Number of L2TP sessions established through the tunnel.

Table 88: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`State`	State of the L2TP tunnel: • `cc_responder_accept_new`—The tunnel has received and accepted the start control connection request (SCCRQ). • `cc_responder_reject_new`—The tunnel has received and rejected the SCCRQ. • `cc_responder_idle`—The tunnel has just been created. • `cc_responder_wait_ctl_conn`—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • `clean-up`—The tunnel is being cleaned up. • `closed`—The tunnel is being closed. • `destroyed`—The tunnel is being destroyed. • Drain—Creation of new sessions and destinations is disabled for this tunnel. • `Established`—The tunnel is operating. This is the only state supported for the LAC. • `Terminate`—The tunnel is terminating. • `Unknown`—The tunnel is not connected to the router.
`Tunnel Name`	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].
`Local IP`	IP address of the local endpoint of the tunnel.
`Local name`	Name used for local tunnel endpoint during tunnel negotiation.
`Remote name`	Name used for remote tunnel endpoint during tunnel negotiation.

Table 88: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`Effective Peer Resync Mechanism`	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.
`Nas Port Method`	NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS: • `cisco-avp`—sends the AVP. • `none`—does not send the AVP.
`Tunnel Logical System`	Logical system in which the L2TP tunnel is brought up.
`Tunnel Routing Instance`	Routing instance in which the L2TP tunnel is brought up.
`Max sessions`	Maximum number of sessions that can be established on this tunnel. The displayed limit for configured sessions is set to the lowest of the following configured session values for either LAC or LNS: • Global (chassis)—`set services l2tp tunnel maximum-sessions`*`number`* • Tunnel profile (individual tunnel)—`set access tunnel-profile` *`profile-name`* `tunnel` *`tunnel-id`*`max-sessions`*`number`*`]` • RADIUS—Value of VSA 26–33, Tunnel-Max-Sessions For LNS only, the following configuration is also considered: • Host profile—`access profile l2tp-profile client default l2tp maximum-sessions-per-tunnel`
`Window size`	Number of control messages that can be sent without receipt of an acknowledgment.

Table 88: show services l2tp tunnel Output Fields *(Continued)*

Field Name	Field Description
`Hello interval`	Interval between the transmission of hello messages, in seconds.
`Create time`	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the when the call was created. If connection to the LAC is severed, the State changes to `Unknown` and the `Create time` value resets.
`Up time`	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.
`Idle time`	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets.

Sample Output

show services l2tp tunnel (LAC)

```
user@host> show services l2tp tunnel
  Local ID  Remote ID  Remote IP              Sessions  State
     17185          1  203.0.113.101:1701            1  Established
```

show services l2tp tunnel detail (LAC)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 203.0.113.101:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 192.0.2.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel detail (LNS)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 198.51.100.15:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/2
    Local IP: 198.51.100.5:1701
    Local name: ce-bras-mx240-e, Remote name: testlac2
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
```

show services l2tp tunnel extensive (LAC)

```
user@host> show services l2tp tunnel extensive
  Tunnel local ID: 17185, Tunnel remote ID:     1
    Remote IP: 203.0.113.101:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 192.0.2.22:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: failover protocol
    Max sessions: 32000, Window size: 4, Hello interval: 60
    Create time: Tue Nov  9 15:23:29 2010, Up time: 00:00:26
    Idle time: 00:00:00
```

show services l2tp tunnel extensive (LNS)

```
user@host> show services l2tp tunnel extensive
Tunnel local ID: 40553, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, State: Established
    Tunnel Name: 3/1838
    Local IP: 203.0.113.2:1701
    Local name: lns-mx960, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
    Max sessions: 60000, Window size: 4, Hello interval: 60
    Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:11
    Idle time: 00:00:00, ToS Reflect: Enabled
    Tunnel Group Name: tg1
    Statistics since: Mon Apr 25 20:27:50 2011
                        Packets        Bytes
      Control Tx              4          219
      Control Rx              4          221
      Data Tx                 0            0
      Data Rx                 6           64
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel statistics

```
user@host>show services l2tp tunnel statistics
Tunnel local ID: 17185, Tunnel remote ID: 1
Sessions: 31.8k, State: Established
Statistics since: Mon Aug 1 13:21:38 2011
                    Packets         Bytes
Control Tx           90.3k             9.0M
Control Rx           32.0k             1296.9k
Data Tx          127.3k         1591.6k
Data Rx             100.8k         1273.4k
Errors Tx            0
Errors Rx               0
```

show services l2tp tunnel-group

IN THIS SECTION

Syntax

```
show services l2tp tunnel-group
<group-name>
```

Description

Display information about all L2TP tunnel groups or a specific L2TP tunnel group.

Options

group-name	(Optional) Name of a tunnel group.

Required Privilege Level

view

Output Fields

Table 89 on page 1005 lists the output fields for the `show services l2tp tunnel-group` command. Output fields are listed in the approximate order in which they appear.

Table 89: show services l2tp tunnel-group Output Fields

Field Name	Field Description
`Tunnel-group`	Name of a tunnel group.
`Tunnels`	Number of tunnels in the tunnel group.
`Sessions`	Number of L2TP sessions established for tunnels in the tunnel group.

Sample Output

show services l2tp tunnel-group

```
user@host> show services l2tp tunnel-group
Tunnel-group            Tunnels        Sessions
l2tp-tunnel-group1      2              20
l2tp-tunnel-group2      3              120
```

show services l2tp tunnel-group (Group Name)

```
user@host> show services l2tp tunnel-group l2tp-tunnel-group1
Tunnel-group            Tunnels        Sessions
l2tp-tunnel-group1      2              20
```

show services l2tp tunnel-switch destination

IN THIS SECTION

Syntax

```
show services l2tp tunnel-switch destination
< detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnel destinations.

Options

none Display standard information for all L2TP switched tunnel destinations.

detail | extensive (Optional) Display the specified level of information.

statistics (Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

Table 90 on page 1007 lists the output fields for the `show services l2tp tunnel-switch destination` command. Output fields are listed in the approximate order in which they appear.

Table 90: show services l2tp tunnel-switch destination Output Fields

Field Name	Field Description	Level of Output
`Local Name`	Name of this destination.	All levels
`Remote IP`	IP address of the remote peer (LNS).	All levels
`Tunnels`	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed

Table 90: show services l2tp tunnel-switch destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sessions`	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total `extensive` for active and failed
`Switched-sessions`	Number of L2TP sessions established by tunnel switching.	All levels
`State`	Administrative state of the L2TP destination: • `Enabled`—No restrictions exist on creation or operation of sessions and tunnels for this destination. • `Disabled`—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the `Disabled` state.	All levels
`Local IP`	IP address of the local gateway (LAC).	`detailextensive`
`Transport`	Medium used for tunneling. Only `ipUdp` is supported.	`detailextensive`
`Logical System`	Logical system in which the tunnel is configured.	`detailextensive`
`Router Instance`	Routing instance in which the tunnel is configured.	`detailextensive`
`Lockout State`	Reachability state of the destination: • `not locked`—Destination is considered reachable. • `waiting for lockout timeout`—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	`detailextensive`

Table 90: show services l2tp tunnel-switch destination Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Connections`	Number of total, active, and failed tunnel and session connections for the destination.	`extensive`
`Control Tx`	Amount of control information transmitted, in packets and bytes.	`extensivestatistics`
`Control Rx`	Amount of control information received, in packets and bytes.	`extensivestatistics`
`Data Tx`	Amount of data transmitted, in packets and bytes.	`extensivestatistics`
`Data Rx`	Amount of data received, in packets and bytes.	`extensivestatistics`
`Error Tx`	Number of errors transmitted, in packets.	`extensivestatistics`
`Error Rx`	Number of errors received, in packets.	`extensivestatistics`

Sample Output

show services l2tp tunnel-switch destination

```
user@host> show services l2tp tunnel-switch destination
  Local Name    Remote IP       Tunnels   Sessions  Switched-sessions  State
  1             192.0.2.3       1         1         1                  Enabled
  2             203.0.113.10    1         1         1                  Enabled
```

show services l2tp tunnel-switch destination detail

```
user@host> show services l2tp tunnel-switch destination detail
  Local name: 1
```

```
    Remote IP: 192.0.2.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.51
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
  Local name: 2
    Remote IP: 198.51.100.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.31
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
```

show services l2tp tunnel-switch destination extensive

```
user@host> show services l2tp tunnel-switch destination extensive
Waiting for statistics...
  Local name: 1
    Remote IP: 192.0.2.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 203.0.113.51
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              6          239
      Control Rx              6          267
      Data Tx                67          815
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
  Local name: 2
    Remote IP: 198.51.100.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP:203.0.113.31
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              7          462
      Control Rx              6          171
      Data Tx                 0            0
      Data Rx                66          798
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel-switch destination statistics

```
user@host> show services l2tp tunnel-switch destination statistics
Waiting for statistics...
  Local name: 2, Tunnels: 1, Sessions: 1
                        Packets        Bytes
      Control Tx              5          452
      Control Rx              4          147
      Data Tx                 0            0
      Data Rx                 4           54
      Errors Tx               0
      Errors Rx               0
  Local name: 1, Tunnels: 1, Sessions: 1
                        Packets        Bytes
      Control Tx              4          184
      Control Rx              4          243
      Data Tx                 5           71
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel-switch session

IN THIS SECTION

Syntax

```
show services l2tp tunnel-switch session
<detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnel sessions.

Options

none	Display standard information about all active L2TP switched tunnel sessions.
detail | extensive	(Optional) Display the specified level of output.

statistics (Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

"show services l2tp tunnel-switch session" on page 1012 lists the output fields for the `show services l2tp tunnel-switch session` command. Output fields are listed in the approximate order in which they appear.

Table 91: show services l2tp tunnel-switch session Output Fields

Field Name	Field Description	Level of Output
`Tunnel local ID`	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
`Local ID`	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	none
`Remote ID`	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	none

Table 91: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
State	State of the L2TP session: • Established—Session is operating. This is the only state supported for the LAC. • closed—Session is being closed. • destroyed—Session is being destroyed. • clean-up—Session is being cleaned up. • lns-ic-accept-new—New session is being accepted. • lns-ic-idle—Session has been created and is idle. • lns-ic-reject-new—New session is being rejected. • lns-ic-wait-connect—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
Interface unit	Logical interface for this session.	All levels
Interface Name	(LNS only) Name of an adaptive services interface.	none
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	detailextensive
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	detailextensive
Tunnel switch profile name	Name of a tunnel switch profile.	detailextensive
Mode	(LNS) Mode of the interface representing the session: shared or exclusive. (LAC) Mode of the interface representing the session: shared or dedicated. Only dedicated is currently supported for the LAC.	detailextensive

Table 91: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Local IP`	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	`detailextensive`
`Remote IP`	IP address of remote endpoint of the PPP session.	`detailextensive`
`Local name`	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	`detailextensive`
`Remote name`	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	`detailextensive`
`Bearer type`	Type of bearer enabled: • `0`—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • `1`—Digital access requested. • `2`—Analog access requested. • `4`—Asynchronous Transfer Mode (ATM) bearer support.	`extensive`
`Framing type`	Type of framing enabled: • `1`—Synchronous framing • `2`—Asynchronous framing	`extensive`
`LCP renegotiation`	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: `On` or `Off`.	`extensive`
`Authentication`	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	`extensive`

Table 91: show services l2tp tunnel-switch session Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Interface ID`	(LNS only) Identifier used to look up the logical interface for this session.	`extensive`
`Call serial number`	Unique serial number assigned to the call.	`extensive`
`Tx speed`	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	`extensive`
`Rx speed`	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	`extensive`
`Create time`	Day, date, and time when the call was created.	`extensive`
`Up time`	Length of time elapsed since the call became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	`extensive`
`ToS Reflect`	Status of IP ToS value reflection, `Disabled` or `Enabled`.	`extensive`
`Statistics since`	Date and time when collection of the following statistics began: • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes.	`extensive`

Sample Output

show services l2tp tunnel-switch session

```
user@host> show services l2tp tunnel-switch session
Tunnel local ID: 37602
  Local  Remote   State               Interface         Interface
  ID     ID                           unit              Name
  13545  1        Established         1073741842        si-2/1/0

Tunnel local ID: 37060
  Local  Remote   State               Interface         Interface
  ID     ID                           unit              Name
  58296  1        Established         1073741843        si-2/1/0
```

show services l2tp tunnel-switch session detail

```
user@host> show services l2tp tunnel-switch session detail
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1, Interface unit: 1073741842
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.0.2.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1, Interface unit: 1073741843
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
```

show services l2tp tunnel-switch session extensive

```
user@host> show services l2tp tunnel-switch session extensive
Tunnel local ID: 37602
```

```
  Session local ID: 13545, Session remote ID: 1
    Interface unit: 1073741842
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.51:1701, Remote IP: 192.0.2.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Bearer type: 2, Framing type: 1
    LCP renegotiation: On, Authentication: None, Interface ID: si-2/1/0
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:06:50
    Idle time: N/A, ToS Reflect: Disabled
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets        Bytes
      Data Tx                85         1031
      Data Rx                 0            0

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1
    Interface unit: 1073741843
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 203.0.113.31:1701, Remote IP: 198.51.100.10:1701
    Local name: lns, Remote name: lns
    Bearer type: 2, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:06:48
    Idle time: N/A
    Statistics since: Fri Jan 18 03:01:14 2013
                        Packets        Bytes
      Data Tx                 0            0
      Data Rx                84         1014
```

show services l2tp tunnel-switch summary

IN THIS SECTION

Syntax

```
show services l2tp tunnel-switch summary
<statistics>
```

Description

Display L2TP tunnel switch summary information.

Options

none Display complete L2TP switched tunnel summary information.

statistics (Optional) Display the number of control packets and bytes transmitted and received for all switched tunnels and sessions.

Required Privilege Level

view

Output Fields

"show services l2tp tunnel-switch summary" on page 1019 lists the output fields for the `show services l2tp tunnel-switch summary` command. Output fields are listed in the approximate order in which they appear.

Table 92: show services l2tp tunnel-switch summary Output Fields

Field Name	Field Description
`Tunnel switch profile name`	Name of a tunnel switch profile.
`LNS local session id`	Identifier assigned by the LNS function on the LTS to the local endpoint of the L2TP session originating on a remote LAC (the first session)
`LAC local session id`	Identifier assigned by the LAC function on the LTS to the local endpoint of the L2TP session originating on the LTS (the second session).
`LNS state`	State of the L2TP session (the first session) between a remote LAC and the LNS function on the LTS.
`LAC state`	State of the L2TP session (the second session) between the LAC function on the LTS and a remote LNS.

Sample Output

show services l2tp tunnel-switch summary

```
user@host> show services l2tp tunnel-switch summary
Tunnel switch profile name: ce-lts-profile
```

```
LNS local   LAC local   LNS state     LAC state     Interface
session ID  session ID                              name
13545       58296       established   established   si-2/1/0
```

show services l2tp tunnel-switch tunnel

IN THIS SECTION

Syntax

```
show services l2tp tunnel-switch tunnel
<detail | extensive>
<statistics>
```

Description

Display information about L2TP switched tunnels.

Options

none	Display standard information about all active L2TP tunnels.
detail | extensive	(Default) Display the specified level of output.
statistics	(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with either of the level options, `detail` or `extensive`.

Required Privilege Level

view

Output Fields

Table 93 on page 1022 lists the output fields for the `show services l2tp tunnel-switch tunnel` command. Output fields are listed in the approximate order in which they appear.

Table 93: show services l2tp tunnel-switch tunnel Output Fields

Field Name	Field Description	Level of Output
`Local ID`	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	none
`Remote ID`	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	none
`Remote IP`	IP address of the peer endpoint of the tunnel.	All levels

Table 93: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Sessions`	Number of L2TP sessions established through the tunnel.	All levels
`Switched-sessions`	Number of L2TP sessions established by tunnel switching.	All levels
`State`	State of the L2TP tunnel: • `cc_responder_accept_new`—The tunnel has received and accepted the start control connection request (SCCRQ). • `cc_responder_reject_new`—The tunnel has received and rejected the SCCRQ. • `cc_responder_idle`—The tunnel has just been created. • `cc_responder_wait_ctl_conn`—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • `clean-up`—The tunnel is being cleaned up. • `closed`—The tunnel is being closed. • `destroyed`—The tunnel is being destroyed. • `Established`—The tunnel is operating. This is the only state supported for the LAC. • `Terminate`—The tunnel is terminating. • `Unknown`—The tunnel is not connected to the router.	All levels
`Tunnel local ID`	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	`detailextensive`

Table 93: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Tunnel remote ID`	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	`detailextensive`
`Tunnel Name`	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].	`detailextensive`
`Local IP`	IP address of the local endpoint of the tunnel.	`detailextensive`
`Local name`	Name used for local tunnel endpoint during tunnel negotiation.	`detailextensive`
`Remote name`	Name used for remote tunnel endpoint during tunnel negotiation.	`detailextensive`
`Effective Peer Resync Mechanism`	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.	`detailextensive`
`NAS Port Method`	(LAC only) Status of interoperation with Cisco LNS devices: • none—NAS port method is not enabled for interoperation. • cisco-avp—NAS port method is enabled for interoperation.	`detailextensive`
`Tunnel Logical System`	Logical system in which the L2TP tunnel is brought up.	`detailextensive`
`Tunnel Routing Instance`	Routing instance in which the L2TP tunnel is brought up.	`detailextensive`

Table 93: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Max sessions`	Maximum number of sessions that can be established on this tunnel.	`extensive`
`Window size`	Number of control messages that can be sent without receipt of an acknowledgment.	`extensive`
`Hello interval`	Interval between the transmission of hello messages, in seconds.	`extensive`
`Create time`	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to `Unknown` and the `Create time` value resets.	`extensive`
`Up time`	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.	`extensive`
`Idle time`	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.	`extensive`
`ToS Reflect`	Status of IP ToS value reflection, `Disabled` or `Enabled`.	`extensive`
`Interface Name`	(LNS only) Name of an adaptive services interface.	`extensive`
`Tunnel Group Name`	(LNS only) Name of a tunnel group.	`extensive`

Table 93: show services l2tp tunnel-switch tunnel Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Statistics since`	Date and time when collection of the following statistics began: • `Control Tx`—Amount of control information transmitted, in packets and bytes. • `Control Rx`—Amount of control information received, in packets and bytes. • `Data Tx`—Amount of data transmitted, in packets and bytes. • `Data Rx`—Amount of data received, in packets and bytes. • `Errors Tx`—Number of errors transmitted, in packets. • `Errors Rx`—Number of errors received, in packets.	`extensive`

Sample Output

show services l2tp tunnel-switch tunnel

```
user@host> show services l2tp tunnel-switch tunnel
  Local ID Remote ID Remote IP              Sessions Switched-sessions State
  37602    1         192.0.2.3:1701                1                 1 Established
  37060    1         198.51.100.10:1701            1                 1 Established
```

show services l2tp tunnel-switch tunnel detail

```
user@host> show services l2tp tunnel-switch tunnel detail
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 1/1
    Local IP: 203.0.113.51:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 198.51.100.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 203.0.113.31:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel-switch tunnel extensive

```
user@host> show services l2tp tunnel-switch tunnel extensive
Waiting for statistics...
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.0.2.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 1/1
    Local IP: 203.0.113.51:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:07:49
    Idle time: 00:00:00, ToS Reflect: Disabled
    Interface Name: si-2/1/0, Tunnel Group Name: ce-l2tp-tunnel-group
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets       Bytes
      Control Tx              7         259
      Control Rx              7         279
      Data Tx                97        1175
      Data Rx                 0           0
      Errors Tx               0
      Errors Rx               0
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 198.51.100.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 203.0.113.31:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:07:46
    Idle time: 00:00:00
    Statistics since: Fri Jan 18 03:01:14 2013
                        Packets        Bytes
      Control Tx              8          482
      Control Rx              7          183
      Data Tx                 0            0
      Data Rx                96         1158
      Errors Tx               0
      Errors Rx               0
```

show system license (View)

IN THIS SECTION

Syntax

```
show system license
<detail>
```

Description

Display licenses and information about how licenses are used.

Options

none	Display all license information.
installed	(Optional) Display installed licenses only.
keys	(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.
status	(Optional) Display license status for a specified logical system or for all logical systems.
usage	(Optional) Display the state of licensed features.

Required Privilege Level

view

Output Fields

Table 94 on page 1029 lists the output fields for the `show system license` command. Output fields are listed in the approximate order in which they appear.

Table 94: show system license Output Fields

Field Name	Field Description
`Feature name`	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.

Table 94: show system license Output Fields *(Continued)*

Field Name	Field Description
`Licenses used`	Number of licenses used by the device. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.
`Licenses installed`	Information about the installed license key: • `License identifier`—Identifier associated with a license key. • `License version`—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. • `Valid for device`—Device that can use a license key. • `Features`—Feature associated with a license.
`Licenses needed`	Number of licenses required for features being used but not yet properly licensed.
`Expiry`	Time remaining in the grace period before a license is required for a feature being used.
`Logical system license status`	Displays whether a license is enabled for a logical system.

Sample Output

show system license

```
user@host> show system license

License usage:
                         Licenses     Licenses     Licenses           Expiry
Feature name                 used    installed       needed
cBNG-SSWLUP-Tier                0       100000            0           permanent
```

show system license detail

```
user@host> show system license detail

Licenses         Licenses   Licenses      Licenses     Licenses      Expiry
Feature name      used       installed      needed       available
scale-subscriber   0          10             0            10         permanent
cBNG-SSWLUP-Tier   0          20000          0            20000      2021-06-29 15:56:00 IST
Licenses installed:
License identifier: RMS123000002
License version: 1
Order Type: commercial
Software Serial Number: AID000000121
Customer ID: SAM7709
License count: 20000
Features:
cBNG-SSWLUP-Tier - Subscriber Services Wireline User Plane Feature
date-based, 2020-06-29 15:56:00 IST - 2021-06-29 15:56:00 IST
```

show system resource-monitor fpc

IN THIS SECTION

Syntax

```
show system resource-monitor fpc
<slot slot-number>
```

Description

Display the utilization of memory resources on the Packet Forwarding Engines for all FPCs or a specific FPC. The filter memory denotes the filter counter memory used for firewall filter counters. The asterisk (*) displayed next to each of the memory regions denotes the ones for which the configured threshold is being currently exceeded.

Options

`slot` *`slot-number`* Display the Junos OS utilization information of memory resources for the specified slot number in which the FPC (or MPC) is installed.

Additional Information

The filter memory denotes the filter counter memory used for firewall filter counters. From the Ukern perspective, MPC5E contains only one Packet Forwarding Engine instance. The `show chassis fabric plane` command output displays the state of fabric plane connections to the Packet Forwarding Engine. Because two Packet Forwarding Engines exist, you notice PFE-0 and PFE-1 in the output.

Required Privilege Level

view

Output Fields

Table 95 on page 1033 lists the output fields for the `show system resource-monitor fpc` command. Output fields are listed in the approximate order in which they appear.

Table 95: show system resource-monitor fpc Output Fields

Field Name	Field Description
`Free Heap Memory Watermark`	Configured watermark value for the percentage of free memory space used for ukernel or heap memory to be monitored
`Free FW Memory Watermark`	Configured watermark value for the percentage of free memory space used for firewall or filter memory to be monitored
`Free NH Memory Watermark`	Configured watermark value for the percentage of free memory space used for next-hop memory to be monitored
`* - watermark reached`	An asterisk (*) displayed beside any of the memory regions denotes the memory types for which the configured threshold is being currently exceeded.
`Slot #`	Slot number in which the line card is installed
`PFE #`	Number or identifier of the Packet Forwarding Engine in the specified line card slot
`Heap % free`	Percentage of free space associated with heap or ukernel memory
`Encap mem % free`	Percentage of free space associated with encapsulation memory
`NH mem % free`	Percentage of free space associated with next-hop memory

Table 95: show system resource-monitor fpc Output Fields *(Continued)*

Field Name	Field Description
`Filter / FW mem % free`	Percentage of free space associated with firewall or filter memory

Sample Output

show system resource-monitor fpc (All Slots)

```
show system resource-monitor fpc
FPC Resource Usage Summary

Free Heap Mem Watermark         : 20  %
Free NH Mem Watermark           : 20  %
Free Filter Mem Watermark       : 20  %

* - Watermark reached


Slot #         % Heap Free        RTT      Average RTT
     0             89                103     102(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              NA                 78                 99

Slot #         % Heap Free        RTT      Average RTT
     2             88                103     103(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              NA                 80                 99
                    1              NA                 80                 99

Slot #         % Heap Free        RTT      Average RTT
     3             91                  0      --(--)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
```

```
                    0              99                 82                 72
                    1              99                 82                 72
```

show system resource-monitor fpc (Specific Slot)

```
show system resource-monitor fpc slot 2
FPC Resource Usage Summary

Free Heap Mem Watermark         : 20  %
Free NH Mem Watermark           : 20  %
Free Filter Mem Watermark       : 20  %

* - Watermark reached


Slot #         % Heap Free        RTT      Average RTT
     2             88                 103     103(30)

                 PFE #      % ENCAP mem Free      % NH mem Free      % FW mem Free
                    0              NA                 80                 99
                    1              NA                 80                 99
```

show system resource-monitor subscribers-limit

IN THIS SECTION

Syntax

```
show system resource-monitor subscribers-limit
<chassis>
<fpc slot-number>
<pic number>
<port number>
<extensive>
<terse>
```

Description

Display information about subscriber limits for the specified hardware element, chassis, FPC, PIC, or port by client type. Shows the configured limit, the number of subscribers of the type currently logged in, and the number of subscribers that have been denied login because the limit has been reached. Use the `extensive` option to display information for the specified element and all subordinate elements that have a configured subscriber limit.

Options

extensive (Optional) Display information for the specified hardware element and all subordinate elements that have a configured subscriber limit.

chassis (Optional) Subscriber limit statistics for the chassis.

fpc *slot-number* (Optional) Subscriber limit statistics for FPC in the specified slot.

pic *number* (Optional) Subscriber limit statistics for the specified PIC.

port *number* (Optional) Subscriber limit statistics for the specified port.

Required Privilege Level

view

Output Fields

Table 96 on page 1037 lists the output fields for the `show system resource-monitor subscribers-limit` command. Output fields are listed in the approximate order in which they appear.

Table 96: show system resource-monitor subscribers-limit Output Fields

Field Name	Field Description	Level of Output
`fpc`, `pic`, `port`	Hardware element on which a maximum subscriber limit is configured.	All levels
`Client-type`	Type of client for which a maximum subscriber limit is configured on the specified hardware element: `ANY`, `DHCP`, `L2TP`, or `PPPoE`.	All levels
`Configured limit`	Maximum number of subscribers that can be logged in for the client type.	All levels
`Current count`	Current number of subscribers that can log in for the client type.	All levels
`Denied count`	Number of subscribers for the client type that have been denied login because the maximum subscriber limit has been reached.	All levels

Sample Output

show system resource-monitor subscribers-limit (Chassis)

```
user@host> show system resource-monitor subscribers-limit chassis
Client-type : pppoe
    Configured limit     : 0
    Current count        : 1
    Denied count         : 0
```

```
Client-type : DHCP
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

Client-type : L2TP
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (Chassis Extensive)

```
user@host> show system resource-monitor subscribers-limit chassis extensive
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

    fpc : 1
    Client-type : pppoe
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

     pic : 0
        Client-type : pppoe
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0

          port : 2
            Client-type : pppoe
                Configured limit    : 0
                Current count       : 1
                Denied count        : 0
```

show system resource-monitor subscribers-limit (FPC)

```
user@host> show system resource-monitor subscribers-limit fpc 1
Client-type : pppoe
```

```
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (FPC Extensive)

```
user@host> show system resource-monitor subscribers-limit fpc 1 extensive
FPC : 1
  Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
  pic : 0
    Client-type : pppoe
        Configured limit    : 0
        Current count       : 1
        Denied count        : 0

      port : 2
        Client-type : pppoe
            Configured limit    : 0
            Current count       : 1
            Denied count        : 0
```

show system resource-monitor subscribers-limit (PIC)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 0
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (PIC Extensive)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 0 extensive
Client-type : pppoe
```

```
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0

  port : 0
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor subscribers-limit (Port)

```
user@host> show system resource-monitor subscribers-limit fpc 1 pic 0 port 2
Client-type : pppoe
    Configured limit    : 0
    Current count       : 1
    Denied count        : 0
```

show system resource-monitor-summary

IN THIS SECTION

Syntax

```
show system resource-monitor summary
```

Description

Display information about round-trip time load throttling for all line cards in the chassis.

Required Privilege Level

view

Output Fields

Table 97 on page 1041 lists the output fields for the `show system resource-monitor summary` command. Output fields are listed in the approximate order in which they appear.

Table 97: show system resource-monitor summary Output Fields

Field Name	Field Description
`Throttle`	Status of throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels: **Enabled** or **Disabled**.
`Heap Mem Threshold`	Percentage of heap memory in use that represents the threshold for throttling subscribers and services.
`Round Trip Delay Threshold`	Internal threshold value against which calculated delay times are evaluated for throttling subscribers and services.

Table 97: show system resource-monitor summary Output Fields *(Continued)*

Field Name	Field Description
`IFL Counter Threshold`	Percentage of filter counter memory in use that represents the threshold for throttling subscribers.
`Filter Counter Threshold`	Percentage of filter counter memory in use that represents the threshold for throttling subscribers.
`Expansion Threshold`	Percentage of expansion memory in use that represents the threshold for throttling subscribers.

`MFS Threshold`
`Used`
	Percentage of main file system memory in use that represents the threshold for throttling services. The Used values is how much of this memory is in use.

`Slot #`	Slot number in which the line card is installed.
`Client allowed`	Whether clients are currently allowed to connect: • Yes—Clients are allowed to connect. • No—Clients are not allowed to connect.
`Service allowed`	Whether services are currently allowed to be created for the subscriber: • Yes—Services are allowed to be created. • No—Services are not allowed to be created.
`Client denied`	Number of new subscribers denied login.
`Service denied`	Number of new services denied completion because the throttle has been exceeded

`Heap memory used`
`In %`
	Number of bytes and percentage of heap memory in use on the line card.

Table 97: show system resource-monitor summary Output Fields *(Continued)*

Field Name	Field Description
`Average Round-trip Delay`	Average calculated round-trip delay for the last 30 round-trip delays.
`Round-trip Delay`	Current calculated round-trip delay. An asterisk indicates that the `Max session rate allowed (%)` is less than 100%. This means that subscriber throttling is active.
`MAX session rate allowed(%)`	Percentage of new subscriber sessions allowed per unit time. When this value is less than 100%, the Round-trip Delay field displays an asterisk.
`Performance Denial Client`	Number of load-based client sessions denied completion because the throttle has been exceeded
`Performance Denial Service`	Number of service sessions denied completion because the throttle has been exceeded.

`Filter memory`
`used | %`
	Number of bytes or total percentage of memory in use for the firewall filter counter on the Packet Forwarding Engine.

`IFL memory`
`used | %`
	Number of bytes or total percentage of memory in use for the logical interface counter on the Packet Forwarding Engine.

`Expansion memory`
`used | %`
	Number of bytes or total percentage of memory in use for expansion memory on the Packet Forwarding Engine. Expansion memory is used when the memory allocated for next-hop and firewall filters is fully consumed.

`PFE #`	Number that identifies the Packet Forwarding Engine for which statistics are displayed.

Sample Output

show system resource-monitor summary

```
show system resource-monitor summary
Resource Usage Summary
Throttle                         : Enabled
Heap Mem Threshold                : 80  %
Round Trip Delay Threshold        :120 ms
IFL Counter Threshold             : 80  %
Filter Counter Threshold          : 80  %
Expansion Threshold               : 95  %
MFS threshold                     : 80  %        Used : 1
Slot # 1
     Client allowed           : Yes
     Service allowed          : Yes
     Client denied            : 0         Performance Denial Client : 0
     Service denied           : 0         Performance Denial Service :0
     Heap memory used         : 183985808       In % : 9
     Average Round-trip Delay    : 127 ms
     Round Trip Delay : 130 ms  *        MAX session rate allowed(%) : 80
                 Filter memory        IFL memory        Expansion memory
      PFE #       used  |   %         used  |   %          used  |   %
          0      29696      0         5056      0             0      0
          1      29536      0         4896      0             0      0
Slot # 2
     Client allowed           : Yes
     Service allowed          : Yes
     Client denied            : 0         Performance Denial Client : 0
     Service denied           : 0         Performance Denial Service :0
     Heap memory used         : 183982960       In % : 9
     Average Round-trip Delay        : 98 ms
     Round Trip Delay         : 100 ms           MAX session rate allowed(%) : 100
                 Filter memory        IFL memory        Expansion memory
      PFE #       used  |   %         used  |   %          used  |   %
          0      29856      0         5216      0             0      0
          1      29376      0         4736      0             0      0
```

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 98 on page 1046 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 98: 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 98: 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 98: 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
```

test aaa authd-lite user

IN THIS SECTION

Syntax

```
test aaa authd-lite user username password password profile access-profile-name
<port nas-port>
<zero-stats>
```

Description

Verify authd-lite subscriber access authentication, accounting, and address allocation configuration.

The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see No Link Title. For information about Juniper Networks VSAs, see No Link Title.

Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag.

Options

username	Specify the subscriber username to test.
password *password*	Specify the password associated with the username.
profile *access-profile-name*	Specify the access profile associated with the subscriber.
port *nas-port*	(Optional) Specify the NAS port used for the test.
zero-stats	(Optional) Specify that no accounting statistics are set for this test.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa authd-lite user

The following example tests the configuration for authd-lite subscriber user1bt with a password of $ABC123 and an access profile of employee12, and displays the resulting output:

```
user@host> test aaa authd-lite user user1bt password $ABC123 profile employee12
    Authentication Grant
    ************User Attributes***********
         User Name -                              user1bt
         Framed IPv6 Prefix -                     ::/0
         Framed IPv6 Pool -                       NULL
         Nas IPv6 Address -                       ::
         NDRA IPv6 Prefix -                       NULL
         Login IPv6 Host -                        ::
         Framed Interface Id -                    0:0:0:0
         Delegated IPv6 Prefix -                  ::/0
         NDRA IPv6 Pool -                         NULL
         User Password -                          $ABC123
         Nas Ip Address -                         0.0.0.0
         NAS Port -                               0
         Service Type-                            0
         Framed IP Address -                      0.0.0.0
         Framed IP Netmask -                      0.0.0.0
         Filter Id -                              NULL
         Framed MTU -                             0
         Reply Message -                          NULL
         Framed Route-                           <not set>
         Framed MTU -                             0
         Class -                                  SBR2CL
         Virtual Router Name (LS:RI)              default:default
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
         Secondary WINS IP Address -              0.0.0.0
         Ingress Statistics -                     disabled
         Egress Statistics -                      disabled
         Ingress Policy Name -                   <not set>
         Engress Policy Name -                   <not set>
         IGMP Enable -                             disabled
```

```
     Redirect VR Name (LS:RI)                 default:default
     Service Bundle                          <not set>
     Framed Ip Route Tag                     <not set>
     LI Action                                0
     LI Interception Identifier               0
     LI Mediation Device IP Address           0.0.0.0
     LI_Mediation_Device_Port_Number          0
     Activate Service                         NULL
     Deactivate Service                       NULL
     Service Statistics                       0
     Ignore_DF_Bit -                          disabled
     IGMP Access Group Name                  <not set>
     IGMP Access Source Group_Name -         <not set>
     MLD Access Group Name                   <not set>
     MLD Access Source Group Name            <not set>
     MLD Version -                            MLD Version not set
     IGMP Version                             IGMP Version not set
     IGMP Immediate Leave -                  <not set>
     MLD Immediate Leave -                   <not set>
     IPv6_Ingress_Policy_Name -              <not set>
     IPv6_Egress_Policy_Name -               <not set>
     Cos_Parameter_Type -                    <not set>
     Service Interim Acct Interval            0
     Max Clients Per Interface                0
     Cos_Scheduler_Pmt_Type                  <not set>
     Session Timeout                          599999940
     NAS Port Type                            0
     Framed Pool                              NULL
     Idle Timeout                             0
Acct-start sent
Acct-start succeeded
Pausing 10 seconds
Interim-Acct sent
Acct-interim succeeded
Pausing 10 seconds
Acct-stop sent
Acct-stop succeeded
Logging out subscriber
Test complete. Exiting
```

test aaa authd-lite user (XML Output)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa authd-lite user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Framed IPv6 Prefix -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Framed IPv6 Pool -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>NDRA IPv6 Prefix -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa dhcp user

IN THIS SECTION

Syntax

```
test aaa dhcp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<mac-address mac-address>
<no-address-request>
<option-82 option-82>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
<service-type service-type>
<source-address source-address>
<terminate-code code-value>
```

Description

Verify Dynamic Host Configuration Protocol (DHCP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

NOTE: The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see No Link Title. For information about Juniper Networks VSAs, see No Link Title.

Options

username

Subscriber username to test.

agent-remote-id *ari*

(Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).

logical-system *logical-system-name*

(Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).

mac-address *mac-address*

(Optional) MAC address of the DHCP client.

no-address-request

(Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed.

> **NOTE**: The `test aaa dhcp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.

option-82 *option-82*

(Optional) DHCP relay agent information option (option-82) value.

password *password*

(Optional) Password associated with the username.

profile *access-profile-name*

(Optional) Access profile associated with the subscriber.

routing-instance *routing-instance-name*

(Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.

service-type *service-type*

(Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)
callback-login (3)	–

source-address *source-address*

(Optional) IP address of the outgoing interface.

terminate-code *code-value*

(Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa dhcp user

The following example tests the configuration for DHCP subscriber user1DB and password $ABC123, and displays the resulting output:

```
user@host> test aaa dhcp user user1DB@test.net password $ABC123
Authentication Grant
    ************User Attributes***********
         User Name -                              user1DB@test.net
         Client IP Address -                      192.168.1.1
         Client IP Netmask -                      255.255.0.0
         Virtual Router Name (LS:RI)-             default:default

         Agent Remote Id -                        NULL
         Reply Message -                          NULL
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
         Secondary WINS IP Address -              0.0.0.0
         Primary DNS IPv6 Address  -              ::
         Secondary DNS IPv6 Address  -            ::
         Framed Pool -                           <not set>
         Service Type -                           0
         DHCP Guided Relay Server -               0
         Class Attribute -                        TEST
         Client IPv6 Address -                    ::
         Client IPv6 Mask -                       null
         Framed IPv6 Prefix -                     ::/0
         Framed IPv6 Pool -                      <not-set>
         NDRA IPv6 Prefix -                      <not-set>
```

```
          Login IPv6 Host -                        ::
          Framed Interface Id -                    0:0:0:0
          Delegated IPv6 Prefix -                  ::/0
          Delegated IPv6 Pool -                   <not-set>
          User Password -                          $ABC123
          CHAP Password -                          NULL
          Mac Address -                            00:00:5E:00:53:ab
          Idle Timeout -                           600
          Session Timeout -                        6000
          Service Name (1) -                       cos-service(video_sch, nc_sch)
          Service Statistics (1) -                 1
          Service Acct Interim (1) -               600
          Service Activation Type (1) -            1
          Service Name (2) -                       filter-service(in_filter, out_filter)
          Service Statistics (2) -                 2
          Service Acct Interim (2) -               900
          Service Activation Type (2) -            1
          Cos shaping rate -                       100m
          Filter Id -                             <not set>
          Framed MTU -                             (null)
          Framed Route -                          <not set>
          Ingress Policy Name -                   <not set>
          Egress Policy Name -                    <not set>
          IGMP Enable -                            disabled
          Redirect VR Name (LS:RI)-                default:default
          Service Bundle -                         Null
          Framed Ip Route Tag -                   <not set>
          Ignore DF Bit -                          disabled
          IGMP Access Group Name -                <not set>
          IGMP Access Source Group Name -         <not set>
          MLD Access Group Name -                 <not set>
          MLD Access Source Group Name -          <not set>
          IGMP Version -                          <not set>
          MLD Version -                           <not set>
          IGMP Immediate Leave -                  <not set>
          MLD Immediate Leave -                   <not set>
          IPv6 Ingress Policy Name -              <not set>
          IPv6 Egress Policy Name -               <not set>
          Dynamic Profile -                       <not set>
          Acct Session ID -                        1
          Acct Interim Interval -                  750
          Acct Type -                              1
          Ingress Statistics -                     disabled
```

```
         Egress Statistics -                      disabled
         Chargeable user identity -               0
         NAS Port Id -                            -0/0/0.0
         NAS Port -                               4095
         NAS Port Type -                          15
         Framed Protocol -                        1
         IPv4 ADF Rule -                          010100
         IPv4 ADF Rule -                          010101
         IPv6 ADF Rule -                          030100
         IPv6 ADF Rule -                          030101
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                          <not set>
    Test complete. Exiting
```

test aaa dhcp user (XML Output)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa dhcp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
            <radius-server-attribute-value>default:default</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Client IP Address -</radius-server-attribute-name>
            <radius-server-attribute-value>198.51.100.7</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Client IP Netmask -</radius-server-attribute-name>
            <radius-server-attribute-value>255.255.255.255</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>


...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```

test aaa ppp user

IN THIS SECTION

Syntax

```
test aaa ppp user username
<agent-remote-id ari>
<logical-system logical-system-name>
<no-address-request>
<password password>
<profile access-profile-name>
<routing-instance routing-instance-name>
```

```
<service-type service-type>
<terminate-code code-value>
```

Description

Verify Point-to-Point Protocol (PPP) subscriber access authentication, accounting, and address allocation configuration by creating a test pseudo session.

> **NOTE**: The `test aaa` command supports all RADIUS-sourced attributes, both IETF standard attributes and Juniper Networks VSAs. Received attributes are displayed in the output. For information about standard RADIUS attributes, see No Link Title. For information about Juniper Networks VSAs, see No Link Title.

Options

username Subscriber username to test.

agent-remote-id *ari* (Optional) Value of the DSL Forum Agent-Remote-Id (VSA 26–2).

logical-system *logical-system-name* (Optional) Logical system in which the subscriber is authenticated. This is the logical system in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25).

no-address-request (Optional) Request is sent for authentication without address allocation. Use for Layer 2-only scenarios where no address allocation request is needed.

> **NOTE**: The `test aaa ppp user` command tries to allocate an IPv4 address even when the subscriber is supposed to get only an IPv6 address. If that behavior is undesirable, include the `no-address-request option` when you issue the command.

password *password* (Optional) Password associated with the username.

profile *access-profile-name* (Optional) Access profile associated with the subscriber.

NOTE: The system logically treats this profile as a client-level configuration. An access profile configured in a domain map takes precedence over client-level configurations. If you have configured one or more domain maps, the username for the user under test is evaluated against the domain maps the same as any other subscriber.

For example, the username can exactly match a domain map or partially match a wildcard domain map. If it matches neither of those, then it matches the `default` domain map if it is configured. If the username has no domain or realm ,then it matches the `none` domain map, if it is configured.

The consequence is that if the test user matches any configured domain map, then an access profile configured in that map is used for the test in preference to an access profile that you specify with the `test` command.

See Specifying an Access Profile in a Domain Map for more information about domain maps and access profiles.

routing-instance ***routing-instance-name***

(Optional) Routing instance in which the subscriber is authenticated. This is the routing instance in the AAA LS:RI context for the subscriber. This context differs from the subscriber context, which is the LS:RI in which the subscriber is placed, by either the Virtual-Router VSA (26-1) or the Redirect-VRouter-Name VSA (26–25). In the case of VSA 26-25, the subscriber is re-authenticated in the subscriber context.

service-type ***service-type***

(Optional) Value of the Service Type RADIUS attribute [6] that is associated with the test user; either a number in the range 1 through 255 or one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)

callback-login (3)	–

terminate-code *code-value*

(Optional) Code associated with the subscriber termination.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request. For information about output fields related to authentication, accounting, and subscriber-specific information, see the **show network-access aaa statistics**, **show network-access aaa statistics authentication**, **show network-access aaa subscribers**, and **show subscribers** commands.

The **test** command does not support volume-time accounting. If volume-time accounting is configured for the test subscriber, the **test** command replaces the statistics with time-only accounting statistics.

This command displays only attributes that are supported by Junos OS; these attributes appear even when their values are not set. The Virtual Router Name (LS:RI) field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays default:default. The displayed value for all other attributes that are not received is `<not set>`.

Sample Output

test aaa ppp user

The following example tests the configuration for PPP subscriber user98BEDC and password $ABC123, and displays the resulting output:

```
user@host> test aaa ppp user user98BEDC@test.net password $ABC123
Authentication Grant
    ************User Attributes***********
```

```
          User Name -                              user98BEDC@test.net
          Client IP Address -                      192.168.1.1
          Client IP Netmask -                      255.255.0.0
          Virtual Router Name (LS:RI) -            default:default
          Agent Remote Id -                        NULL
          Reply Message -                          NULL
          Primary DNS IP Address -                 0.0.0.0
          Secondary DNS IP Address -               0.0.0.0
          Primary WINS IP Address -                0.0.0.0
          Secondary WINS IP Address -              0.0.0.0
          Primary DNS IPv6 Address  -              ::
          Secondary DNS IPv6 Address  -            ::
          Framed Pool -                           <not set>
          Class Attribute -                        TEST
          Service Type -                           0
          Client IPv6 Address -                    ::
          Client IPv6 Mask -                       null
          Framed IPv6 Prefix -                     ::/0
          Framed IPv6 Pool -                      <not-set>
          NDRA IPv6 Prefix -                      <not-set
          Login IPv6 Host -                        ::
          Framed Interface Id -                    0:0:0:0
          Delegated IPv6 Prefix -                  ::/0
          Delegated IPv6 Pool -                   <not-set>
          User Password -                          $ABC123
          CHAP Password -                          NULL
          Mac Address -                            00:00:5E:00:53:ab
          Idle Timeout -                           600
          Session Timeout -                        6000
          Service Name (1) -                       cos-service(video_sch, nc_sch)
          Service Statistics (1) -                 1
          Service Acct Interim (1) -               600
          Service Activation Type (1) -            1
          Service Name (2) -                       filter-service(in_filter, out_filter)
          Service Statistics (2) -                 2
          Service Acct Interim (2) -               900
          Service Activation Type (2) -            1
          Cos shaping rate -                       100m
          Filter Id -                             <not set>
          Framed MTU -                             (null)
          Framed Route -                          <not set>
          Ingress Policy Name -                   <not set>
          Egress Policy Name -                    <not set>
```

```
         IGMP Enable -                            disabled
         Redirect VR Name (LS:RI) -               default
         Service Bundle -                         Null
         Framed Ip Route Tag -                   <not set>
         Ignore DF Bit -                          disabled
         IGMP Access Group Name -                <not set>
         IGMP Access Source Group Name -         <not set>
         MLD Access Group Name -                 <not set>
         MLD Access Source Group Name -          <not set>
         IGMP Version -                          <not set>
         MLD Version -                           <not set>
         IGMP Immediate Leave -                  <not set>
         MLD Immediate Leave -                   <not set>
         IPv6 Ingress Policy Name -              <not set>
         IPv6 Egress Policy Name -               <not set>
         Dynamic Profile -                       <not set>
         Acct Session ID -                        1
         Acct Interim Interval -                  750
         Acct Type -                              1
         Chargeable user identity -               0
         NAS Port Id -                            -0/0/0.0
         NAS Port -                               4095
         NAS Port Type -                          15
         Framed Protocol -                        1
         IPv4 ADF Rule -                          010100
         IPv4 ADF Rule -                          010101
         IPv6 ADF Rule -                          030100
         IPv6 ADF Rule -                          030101
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                          <not set>
    Test complete. Exiting
```

test aaa ppp user (tunneled user)

The following example tests the configuration for PPP tunneled subscriber accounting14, with password $ABC123 and access profile finance-b, and displays the resulting output:

```
user@host> test aaa ppp user accounting14 password $ABC123 14 profile finance-b
    Authentication Grant with Tunnel Attributes
    ************Tunnel Attributes***********
```

```
         ****Tunnel Definiton -                   1
             Tunnel Medium           -            1
             Tunnel Type             -            3
             Tunnel Max Sessions     -            100
             Tunnel Server Endpoint  -            192.0.2.4
             Tunnel Client Endpoint  -            198.51.100.5
             Tunnel Server AuthId    -            rt1
             Tunnel Client AuthId    -            ts1
             Tunnel Password         -            radius
             Tunnel Assignment Id    -            til
             Tunnel Logical System   -
             Tunnel Routing Instance -
    ****Pausing 10 seconds before disconnecting the test user*********
    Logging out subscriber
         Terminate Id -                           l2tp session-receive-cdn-avp-bad-hidden
    Test complete. Exiting
```

test aaa ppp user (authentication failure)

The following example shows sample output when the authentication grant fails due to an invalid password:

```
user@host>test aaa ppp user user45@test.net password $ABC123123
 Authentication Deny
    Reason : Access Denied
    Received Attributes :
         User Name -                              user45@test.net
         Client IP Address -                      0.0.0.0
         Client IP Netmask -                      0.0.0.0
         Virtual Router Name (LS:RI)-             default
         Agent Remote Id -                        NULL
         Reply Message -                          NULL
         Primary DNS IP Address -                 0.0.0.0
         Secondary DNS IP Address -               0.0.0.0
         Primary WINS IP Address -                0.0.0.0
         Secondary WINS IP Address -              0.0.0.0
         Primary DNS IPv6 Address  -              ::
         Secondary DNS IPv6 Address  -            ::
         Framed Pool -                            not set
         Class Attribute -                        not set
         Service Type -                           0
         Client IPv6 Address -                    ::
         Client IPv6 Mask -                       null
         Framed IPv6 Prefix -                     ::/0
         Framed IPv6 Pool -                       not-set
         NDRA IPv6 Prefix -                       not-set
         Login IPv6 Host -                        ::
         Framed Interface Id -                    0:0:0:0
         Delegated IPv6 Prefix -                  ::/0
         Delegated IPv6 Pool -                    not-set
         User Password -                          $ABC123123
         CHAP Password -                          NULL
         Mac Address -                            00:00:5E:00:53:ab
         Filter Id -                              not set
         Framed MTU -                             (null)
         Framed Route -                           not set
         Ingress Policy Name -                    not set
         Egress Policy Name -                     not set
         IGMP Enable-                             disabled
         Redirect VR Name (LS:RI)-                default
         Service Bundle -                         Null
         Framed Ip Route Tag -                    not set
         Ignore DF Bit -                          disabled
         IGMP Access Group Name -                 not set
         IGMP Access Source Group Name -          not set
         MLD Access Group Name -                  not set
         MLD Access Source Group Name -           not set
         IGMP Version -                           not set
         MLD Version -                            not set
         IGMP Immediate Leave -                   not set
         MLD Immediate Leave -                    not set
         IPv6 Ingress Policy Name -               not set
         IPv6 Egress Policy Name -                not set
         Acct Session ID -                        12
         Acct Interim Interval -                  0
         Acct Type -                              0                         Chargeable user
identity -               0
         NAS Port Id -                            -0/0/0.0
         NAS Port -                               4095
         NAS Port Type -                          15
         Framed Protocol -                        0
    Test complete. Exiting
```

test aaa ppp user (XML Output)

The following example shows an excerpt of sample XML output in the new format:

```
user@host>test aaa ppp user user45@test.net password $ABC123 | display xml

<rpc-reply xmlns:junos="namespace-URL">
    <aaa-test-result>
        <aaa-test-status>Authentication Grant</aaa-test-status>
        <aaa-test-status>************User Attributes***********</aaa-test-status>
        <radius-server-data>
            <radius-server-attribute-name>User Name -</radius-server-attribute-name>
            <radius-server-attribute-value>user45@test.net</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Virtual Router Name (LS:RI) -</radius-server-attribute-
name>
            <radius-server-attribute-value>default:default</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Service Type -</radius-server-attribute-name>
            <radius-server-attribute-value>Framed</radius-server-attribute-value>
        </radius-server-data>
        <radius-server-data>
            <radius-server-attribute-name>Agent Remote Id -</radius-server-attribute-name>
            <radius-server-attribute-value>&lt;not set&gt;</radius-server-attribute-value>
        </radius-server-data>
...
        <aaa-test-status>Test complete. Exiting</aaa-test-status>
    </aaa-test-result>
    <cli>
        <banner></banner>
    </cli>
</rpc-reply>
```