Network Design and Architecture Center Books

UNDERSTANDING SUBSCRIBER MANAGEMENT AND BNG

By the writers and editors of the Juniper Networks TechLibrary

UNDERSTANDING
SUBSCRIBER MANAGEMENT AND BNG

By the writers and editors of the Juniper Networks TechLibrary

Industry data indicates that subscribing consumers prefer a bundled voice, video, and data service package and that they prefer getting a single bill from a trusted operator that can supply all of their service demands. But how does that all work? How does an automated process identify individual subscribers, their access rights, their billing information, and keep it all operational while serving millions of customers every day? It's called *subscriber management* and it's handled at the edge between provider and user.

Broadband network operators have evolved from connectivity providers to all-inclusive solution providers that offer IP-based data, video, and voice content to their subscribers on a myriad of devices while making all the content easier to access.

The Evolution of Residential Broadband Multiplay

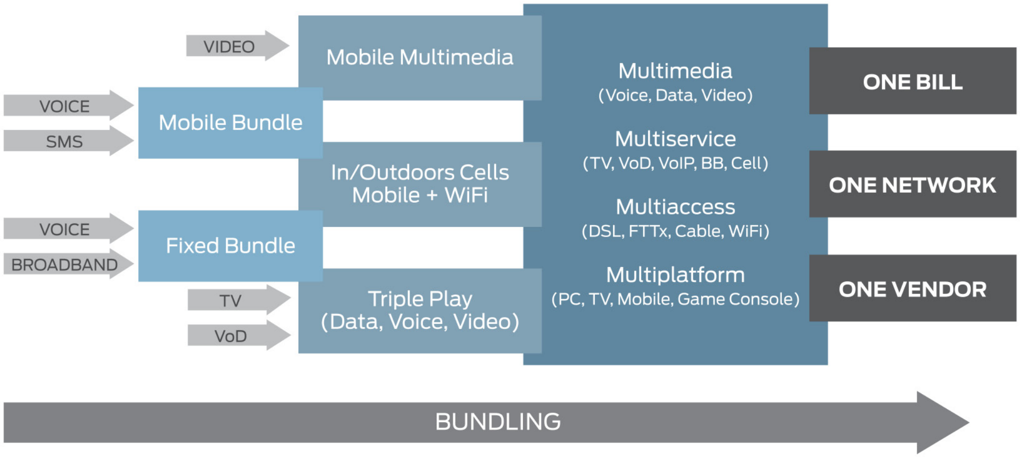

Understanding Subscriber Management and Broadband Network Gateway (BNG) is your introduction to the concepts and technologies of this amazing world. Gleaned from thousands of pages within the Juniper Networks TechLibrary, this book represents clear and lucid coverage of how broadband subscriber management works. As part of the *Juniper Networks Design and Architectural Center* mandate, this book brings readers and engineers together to a common staging area for network enhancement.

Network Design and Architecture Center Books
http://www.juniper.net/documentation/en_US/design-and-architecture/index.html

JUNIPER NETWORKS

UNDERSTANDING
SUBSCRIBER MANAGEMENT AND BNG

By the writers and editors of the Juniper Networks TechLibrary

Published by Juniper Networks Books
Understanding Subscriber Managerment and BNG **is** compiled from various articles and documents from the Juniper Networks TechLibrary listed in the Appendix.

Version History: June 2017

2 3 4 5 6 7 8 9 10

Glossary

AAA

Authentication, authorization, and accounting. Process framework used to standardize the control of access to computer resources, enforcement of policies, audit of usage, and ability to report. Authentication determines who the user is and whether to grant that user access to the network. Authorization determines what the user can do by giving you the ability to limit network services to different users. Accounting tracks the user's activities and provides an audit trail that can be used for billing for connection time or resources used.

AAA method for subscriber authentication

The AAA method that uses authentication (for example, including RADIUS VSAs in the Access-Accept packet) to verify a subscriber and activate a service when the subscriber logs in.

BNG

Broadband network gateway. A broadband remote access server (BRAS, B-RAS, or BBRAS) routes traffic to and from broadband remote access devices such as digital subscriber line access multiplexers (DSLAM) on an Internet service provider's (ISP) network.

Class of service (CoS)

Method of classifying traffic on a packet-by-packet basis using information in the type-of-service (ToS) byte to provide different service levels to different traffic. Enables you to divide traffic into classes and offer various levels of throughput and acceptable packet loss when congestion occurs.

CoA

Change of authorization. RADIUS messages that dynamically modify session authorization attributes, such as data filters.

Customer VLAN (C-VLAN)

Defined by IEEE 802.1ad. A stacked VLAN contains an outer tag corresponding to the S-VLAN, and an inner tag corresponding to the C-VLAN. A C-VLAN often corresponds to CPE. Scheduling and shaping is often used on a C-VLAN to establish minimum and maximum bandwidth limits for a customer. See also S-VLAN.

Demultiplexing (demux) interface

Logical interfaces that share a common, underlying logical interface (in the case of IP demux) or underlying physical interface (in the case of VLAN demux).

DHCP

Dynamic Host Configuration Protocol. Mechanism through which hosts using TCP/IP can obtain protocol configuration parameters automatically from a DHCP server on the network; allocates IP addresses dynamically so that they can be reused when no longer needed.

Domain mapping

Enables you to configure a map that specifies access options and session-specific parameters. The map is based on the domain name of subscriber sessions; the router applies the mapped options and parameters to sessions for subscribers that have the specified domains.

DSL

Digital subscriber line. Most widely deployed and delivered broadband option technology worldwide. Uses existing telephone lines to send broadband information on a different frequency than is used for the existing voice service.

Dynamic profile

A template that defines a set of characteristics that are combined with authorization attributes and are dynamically assigned to static interfaces to provide dynamic subscriber access and services for broadband applications.

Interface set

A logical group of interfaces that describe the characteristics of set of service VLANs, logical interfaces, customer VLANs, or aggregated Ethernet interfaces. Interface sets establish the set and name the traffic control profiles. See also Service VLAN.

L2TP

Layer 2 Tunneling Protocol (L2TP) is a client-server protocol that allows the Point-to-Point Protocol (PPP) to be tunneled across a network. L2TP encapsulates Layer 2 packets, such as PPP, for transmission across a network.

Multilink Point-to-Point Protocol (MLPPP)

Aggregates multiple PPP physical links into a single virtual connection, or logical bundle. More specifically, MLPPP bundles multiple link-layer channels into a single network-layer channel.

Multiservice access node (MSAN)

Refers to a group of commonly used aggregation devices. These devices include digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections.

Passive Optical Networking (PON)

Uses fiber-optic cable to deliver services to the premises. This delivery option provides higher speeds than DSL but lower speeds than Active Ethernet.

Pseudowire

A tunnel that is either an MPLS-based Layer 2 VPN or Layer 2 circuit. The pseudowire tunnel transports Ethernet encapsulated traffic from an access node to the MX Series router that hosts the subscriber management services.

Point-to-Point Protocol (PPP)

PPP support enables you to create and attach dynamic profiles for PPP subscriber interfaces. When the PPP subscriber logs in, the router instantiates the specified dynamic profile and then applies the attributes defined in the profile to the interface.

Point-to-Point Protocol over Ethernet (PPPoE)

Network protocol that encapsulates PPP frames in Ethernet frames and connects multiple hosts over a simple bridging access device to a remote access concentrator. Allows multiple users at a site to share the same digital subscriber line, cable modem, or wireless connection to the Internet.

RADIUS CoA method

The method that uses RADIUS CoA-Request messages and VSAs to activate a service for a subscriber that is already logged in.

Scheduler

Defines the scheduling and queuing characteristics of a queue. Transmit rate, scheduler priority, and buffer size can be specified. In addition, a drop profile may be referenced to describe WRED congestion control aspects of the queue. See also Scheduler map.

Scheduler map

Referenced by traffic control profiles to define queues. The scheduler map establishes the queues that comprise a scheduler node and associates a forwarding class with a scheduler. See also Scheduler.

Service VLAN (S-VLAN)

Defined by IEEE 802.1ad, often corresponds to a network aggregation device such as a DSLAM. Scheduling and shaping is often established for an S-VLAN to provide CoS for downstream devices with little buffering and simple schedulers. See also Customer VLAN.

Stacked VLAN

An encapsulation on an S-VLAN with an outer tag corresponding to the S-VLAN, and an inner tag corresponding to the C-VLAN. See also Service VLAN and Customer VLAN.

Subscriber access technology

The technology used by a subscriber to access services (for example, DHCP).

Traffic control profile

Defines the characteristics of a scheduler node. Traffic control profiles are used at several levels of the CLI, including the physical interface, interface set, and logical interface levels. See also Scheduler and Scheduler map.

VLAN

Virtual LAN, defined on an Ethernet logical interface.

VSA

Vendor-specific attributes. RADIUS server attributes specific to Juniper Networks and are described in RFC 2138, Remote Authentication Dial In User Service (RADIUS).

Understanding Broadband Subscriber Management

NOTE The Appendix contains the links and URLs to all the original articles in this book.

Broadband Subscriber Management dynamically provisions and manages subscriber access in a multiplay or triple play network environment. It uses authentication, authorization, and accounting (AAA) or DHCP with dynamic profiles to provide dynamic, per-subscriber authentication, addressing, and access for a host of broadband services including Internet access, gaming, IPTV, Video on Demand (VoD), and subscriber wholesaling.

Subscriber management provides convenience and flexibility to service providers and subscribers:

- Service providers can separate services, access technology, and eliminate unprofitable flat-rate billing. They gain the ability to efficiently design, manage, and deliver services that subscribers want, and then bill subscribers based on connect time, bandwidth, and the actual service used.
- Subscribers benefit by gaining access to multiple simultaneous services. Depending on the service provider configuration, subscribers can dynamically connect to and disconnect from various services when they want and for however long they want. Subscribers can be billed based on the service level and usage, rather than being charged a set rate regardless of usage.

The Juniper Networks Junos OS subscriber management feature provides subscriber access, authentication, and service creation, activation, and deactivation. You can also collect accounting information and statistics for subscriber service sessions.

The subscriber access feature supports both command-line interface (CLI) and AAA-based configuration (such as RADIUS) for subscribers. Access and services start when the router receives a message from a client (such as a Dynamic Host Configuration Protocol (DHCP) DHCP discover message). For RADIUS clients, RADIUS Access-Accept messages and Change-of-Authorization-Request (CoA-Request) messages can create, modify, and delete subscriber sessions as well as activate and deactivate service sessions. You can use CLI commands to create a dynamic profile to use as a template for user attributes.

A subscriber service is based on the combination of a defined dynamic profile and attributes configured through authentication. Dynamic profiles can include dynamic firewall filters, class-of-service (CoS) settings and protocol (IGMP) settings that define access limits for subscribers and the scope of a service granted to the subscriber after access is obtained.

Broadband Network Gateway (BNG)

Broadband subscriber management is supported on and performed by the Juniper Networks Broadband Network Gateway (BNG).

The BNG dynamically delivers highly customized subscriber services by:

- Managing subscribers, including session and circuit aggregation

- Delivering IP services to subscribers
- Establishing and managing subscriber sessions
- Managing subscriber addressing
- Performing authentication, authorization, accounting (AAA)
- Performing policy and traffic management functions
- Aggregating traffic from various subscriber sessions from an access network, and routes it to the service provider network

The BNG runs on MX Series 3D Universal Edge routers to deliver high subscriber density, automation capabilities, and resiliency, giving service providers flexibility, scalability, and reliability.

AAA Service Framework

The authentication, authorization, and accounting (AAA) Service Framework provides a single point of contact for all the authentication, authorization, accounting, address assignment, and dynamic request services that the router supports for network access. The framework supports authentication and authorization through external servers, such as RADIUS. The framework also supports accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.

The BNG interacts with external servers to determine how individual subscribers access the broadband network. The router also obtains information from external servers for the following:

- Methods used for authentication and accounting
- How accounting statistics are collected and used
- How dynamic requests are handled

When interacting with external back-end RADIUS servers, the AAA Service Framework supports standard RADIUS attributes and Juniper Networks vendor specific attributes (VSAs). The AAA Service Framework also includes an integrated RADIUS client that is compatible with RADIUS servers that conform to RFC-2865, *Remote Authentication Dial In User Service (RADIUS)*, RFC-2866, *RADIUS Accounting*, and RFC-3576, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*, and which can initiate requests.

You create the following types of configurations to manage subscriber access:

Authentication – Authentication parameters defined in the access profile determine the authentication component of the AAA processing. For example, subscribers can be authenticated using an external authentication service such as RADIUS.

Accounting – Accounting parameters in the access profile specify the accounting part of the AAA processing. For example, the parameters determine how the router collects and uses subscriber statistics. You can also configure AAA to enable the router to collect statistics on a per-service session basis for subscribers.

RADIUS-initiated dynamic requests – A list of authentication server IP addresses in the access profile specify the RADIUS servers that can initiate dynamic requests to the router. Dynamic requests include CoA requests, which specify VSA modifications and service changes, and disconnect requests, which terminate subscriber sessions. The list of authentication servers also provides RADIUS-based dynamic service activation and deactivation during subscriber login.

Address assignment – The AAA Service Framework assigns addresses to subscribers based on the configuration of local address-assignment pools. For example, the AAA framework collaborates with RADIUS servers to assign addresses from the specified pools.

Subscriber secure policy – RADIUS VSAs and attributes provide RADIUS-initiated traffic mirroring on a per-subscriber basis.

Class of Service

Class of service (CoS) enables you to divide traffic into classes and offer various levels of throughput and acceptable packet loss when congestion occurs. CoS also provides the option of using differentiated services when best-effort traffic delivery is insufficient. You can also configure the services router to provide hierarchical scheduling for subscribers by dynamically adding or deleting queues when subscribers require services.

By using a dynamic profile, you can provide all subscribers in your network with default CoS parameters when they log in. For example, you can configure an access dynamic profile to specify that all subscribers receive a basic data service. If you use RADIUS variables in the dynamic profile, you can enable the service to be activated for those subscribers at login. You can also use variables to configure a service profile that enables subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages following initial login.

DHCP-based Subscriber Access Configuration

Figure 1 shows the configuration sequence you perform for DHCP-based subscriber access. It also shows the dynamic configuration performed by the router.

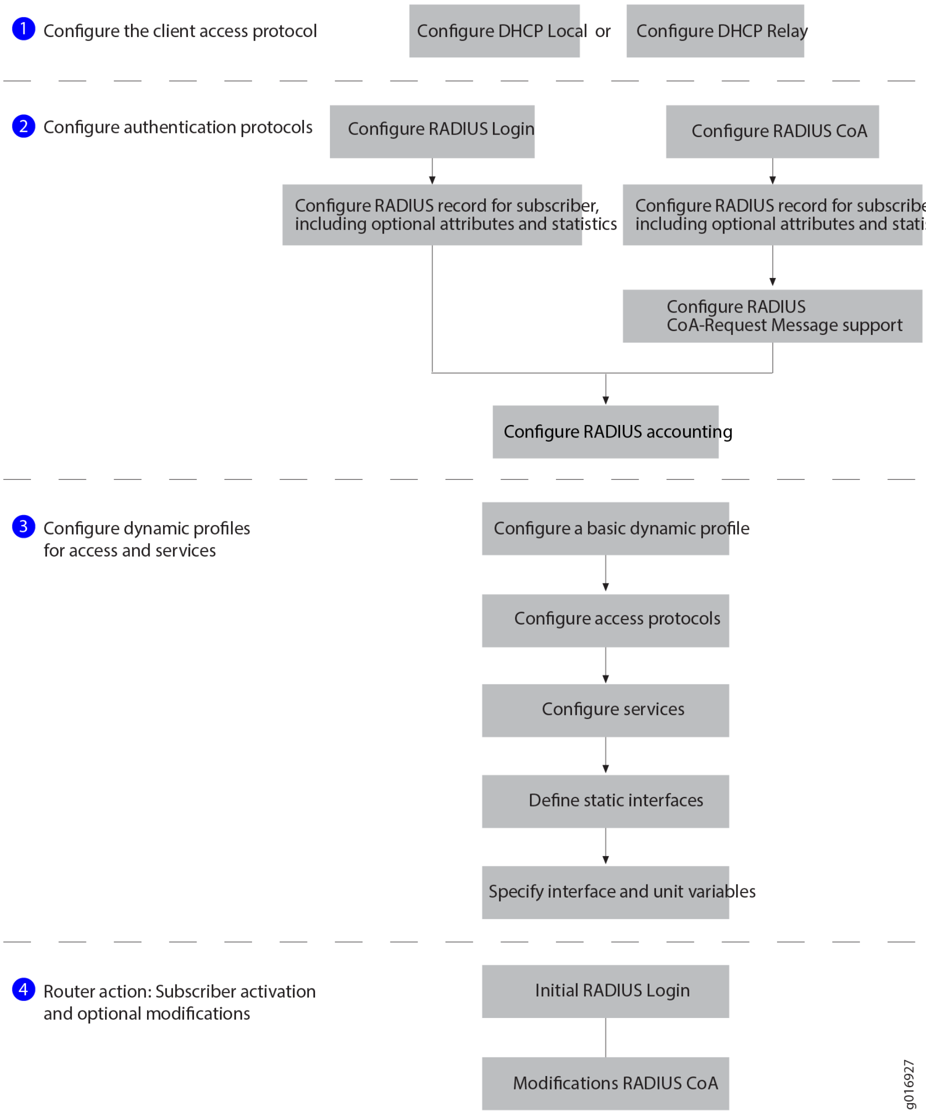
1 Configure the client access protocol
Configure DHCP Local
or
Configure DHCP Relay
2 Configure authentication protocols
Configure RADIUS Login
Configure RADIUS CoA
Configure RADIUS record for subscriber, including optional attributes and statistics
Configure RADIUS record for subscriber, including optional attributes and statistics
Configure RADIUS CoA-Request Message support
Configure RADIUS accounting
3 Configure dynamic profiles for access and services
Configure a basic dynamic profile
Configure access protocols
Configure services
Define static interfaces
Specify interface and unit variables
4 Router action: Subscriber activation and optional modifications
Initial RADIUS Login
Modifications RADIUS CoA
g016927

Figure 1 DHCP-based Subscriber Access Configuration

Components of a Dynamic Profile

The subscriber access feature uses dynamic profiles to activate subscribers and manage services. A dynamic profile is a set of characteristics (defined in a template) that the router uses to provide dynamic subscriber access and services. These services are assigned dynamically to interfaces.

You can use dynamic profiles to define various router components for subscriber access, including the following:

- Dynamic firewall filters – Includes input and output filters to enforce rules that define whether to permit or deny packets that are transmitting an interface on the router. To apply dynamic firewall filters to the subscriber interface, you configure static input and output firewall filters and reference those filters in dynamic profiles.
- Dynamic CoS – Includes CoS values that define a service for a subscriber. For example, you can configure the shaping rate for traffic in a video service by referencing CoS statements in a dynamic profile.
- Dynamic signaling protocol – Includes dynamic IGMP configuration for host to router signaling for IPv4 to support IP multicasting.

Using these profiles enables you to consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes or dynamically created objects simultaneously.

After profiles are created, they reside on the router in a profile library. These profiles can contain various configurations. For example, you can create a client network access configuration, a services activation configuration, or both. When a router interface receives a join message from a client, the router applies the values configured in the specified dynamic profile to that router interface. The profile can contain interface, CoS, and protocol values that are applied directly to the interface. In addition, the dynamic profile can call input or output firewall filters that reside outside of the dynamic profiles hierarchy.

The dynamic-profiles hierarchy appears at the top level of the CLI hierarchy and contains many Juniper Networks configuration statements that you normally define statically.

Dynamic profile statements appear in the following subhierarchies within the [edit dynamic-profiles] hierarchy:

- class-of-service
- firewall
- interfaces
- predefined-variable-defaults
- protocols
- routing-instances
- routing-options
- variables

Understanding Subscriber Access Network

A subscriber access environment includes subscriber access technologies and authentication protocols. The subscriber access technologies include:

- Dynamic Host Configuration Protocol (DHCP) server
- Local DHCP server
- External DHCP server

- Point-to-Point Protocol (PPP)

The subscriber authentication protocols include the RADIUS server. Figure 2 shows an example of a basic subscriber access network.

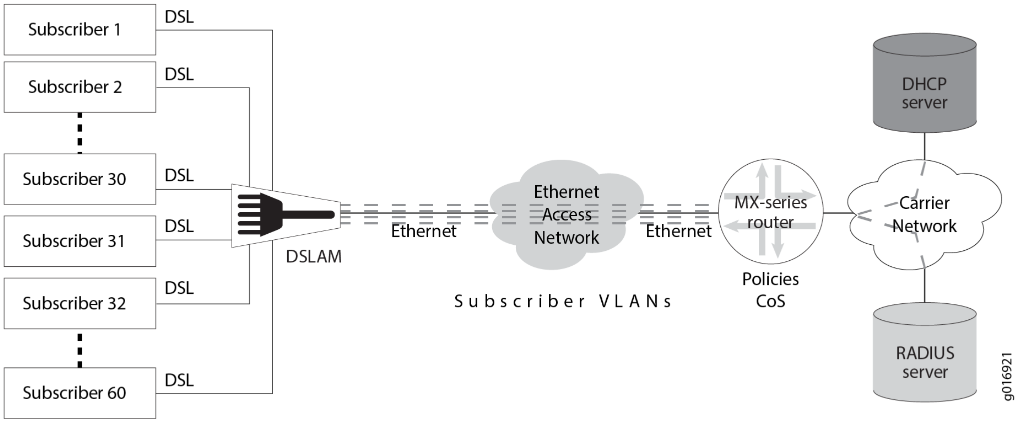

Figure 2 Basic Subscriber Access Network Example

Multiservice Access Node Overview

A *multiservice access node* is a broader term that refers to a group of commonly used aggregation devices. These devices include digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections. Modern MSANs often support all of these connections, as well as providing connections for additional circuits such as plain old telephone service (referred to as POTS) or Digital Signal 1 (DS1 or T1).

The defining function of a multiservice access node is to aggregate traffic from multiple subscribers. At the physical level, the MSAN also converts traffic from the *last mile technology* (for example, ADSL) to Ethernet for delivery to subscribers.

You can broadly categorize MSANs into three types based on how they forward traffic in the network:

- Layer–2 MSAN – This type of MSAN is essentially a Layer 2 switch (though typically not a fully functioning switch) with some relevant enhancements. These MSANs use Ethernet (or ATM) switching to forward traffic. The MSAN forwards all subscriber traffic upstream to an edge router that acts as the centralized control point and prevents direct subscriber-to-subscriber communication. Ethernet Link Aggregation (LAG) provides the resiliency in this type of network.

- Layer 2 DSLAMs cannot interpret IGMP, so they cannot selectively replicate IPTV channels.

- Layer–3 aware MSAN – This IP-aware MSAN can interpret and respond to IGMP requests by locally replicating a multicast stream and forwarding the stream to any subscriber requesting it. Layer 3 awareness is important when supporting IPTV traffic to perform channel changes (sometimes referred to as *channel zaps*). Static IP-aware MSANs always receive all multicast television channels. They do not have the ability to request that specific channels be forwarded to the DSLAM. Dynamic IP-aware DSLAMs, however, can inform the network to begin (or discontinue) sending individual channels to the DSLAM. Configuring IGMP proxy or IGMP snooping on the DSLAM accomplishes this function.

- Layer–3 MSAN – These MSANs use IP routing functionality rather than Layer 2 technologies to forward traffic. The advantage of this forwarding method is the ability to support multiple upstream links going to different upstream routers and improving network resiliency. However, to accomplish this level of resiliency, you must assign a separate IP subnetwork to each MSAN, adding a level of complexity that can be more difficult to maintain or manage.

Ethernet MSAN Aggregation Options

Each MSAN can connect directly to an edge router (broadband services router or video services router), or an intermediate device (for example, an Ethernet switch) can aggregate MSAN traffic before being sent to the services router. Table 1 lists the possible MSAN aggregation methods and under what conditions they are used.

Method	When Used
Direct connection	Each MSAN connects directly to the broadband services router and optional video services router.
Ethernet aggregation switch connection	Each MSAN connects directly to an intermediate Ethernet switch. The switch, in turn, connects to the broadband services router or optional video services router.
Ethernet ring aggregation connection	Each MSAN connects to a ring topology of MSANs. The head-end MSAN (the device closest to the upstream edge router) connects to the broadband services router.

Table 1 Ethernet MSAN Aggregation Methods

You can use different aggregation methods in different portions of the network. You can also create multiple layers of traffic aggregation within the network. For example, an MSAN can connect to a central office terminal (COT), which, in turn, connects to an Ethernet aggregation switch, or you can create multiple levels of Ethernet aggregation switches prior to connecting to the edge router.

Broadband Access Service Delivery Options

Four primary delivery options exist today for delivering broadband network service:

- Digital Subscriber Line
- Active Ethernet
- Passive Optical Networking
- Hybrid Fiber Coaxial

Digital Subscriber Line

Digital subscriber line (DSL) is the most widely deployed broadband technology worldwide. This delivery option uses existing telephone lines to send broadband information on a different frequency than is used for the existing voice service. Many generations of DSL are used for residential service, including Very High Speed Digital Subscriber Line 2 (VDSL2) and versions of Asymmetric Digital Subscriber Line (ADSL, ADSL2, and ADSL2+). These variations of DSL primarily offer asymmetric residential broadband service where different upstream and downstream speeds are implemented. (VDSL2 also supports symmetric operation.) Other DSL variations, like High bit rate Digital Subscriber Line (HDSL) and Symmetric Digital Subscriber Line (SDSL), provide symmetric speeds and are typically used in business applications.

The head-end to a DSL system is the Digital Subscriber Line Access Multiplexer (DSLAM). The demarcation device at the customer premise is a DSL modem. DSL service models are defined by the Broadband Forum (formerly called the DSL Forum).

Active Ethernet

Active Ethernet uses traditional Ethernet technology to deliver broadband service across a fiber-optic network. Active Ethernet does not provide a separate channel for existing voice service, so VoIP (or TDM-to-VoIP) equipment is required. In addition, sending full-speed (10 or 100 Mbps) Ethernet requires significant power, necessitating distribution to Ethernet switches and optical repeaters located in cabinets outside of the central office. Due to these restrictions, early Active Ethernet deployments typically appear in densely populated areas.

Passive Optical Networking

Passive Optical Networking (PON), like Active Ethernet, uses fiber-optic cable to deliver services to the premises. This delivery option provides higher speeds than DSL but lower speeds than Active Ethernet. Though PON provides higher speed to each subscriber, it requires a higher investment in cable and connectivity.

A key advantage of PON is that it does not require any powered equipment outside of the central office. Each fiber leaving the central office is split using a non-powered optical splitter. The split fiber then follows a point-to-point connection to each subscriber.

PON technologies fall into three general categories:

- ATM PON (APON), Broadband PON (BPON), and Gigabit-capable PON (GPON) – PON standards that use the following different delivery options:
- APON – The first passive optical network standard is primarily used for business applications.
- BPON – Based on APON, BPON adds wave division multiplexing (WDM), dynamic and higher upstream bandwidth allocation, and a standard management interface to enable mixed-vendor networks.
- GPON – The most recent PON adaptation, GPON is based on BPON but supports higher rates, enhanced security, and a choice of which Layer 2 protocol to use (ATM, Generic Equipment Model [GEM], or Ethernet).

- Ethernet PON (EPON) – Provides capabilities similar to GPON, BPON, and APON, but uses Ethernet standards. These standards are defined by the IEEE. Gigabit Ethernet PON (GEPON) is the highest speed version.
- Wave Division Multiplexing PON (WDM-PON) – A nonstandard PON which, as the name implies, provides a separate wavelength to each subscriber.

The head-end to a PON system is an Optical Line Terminator (OLT). The demarcation device at the customer premises is an Optical Network Terminator (ONT). The ONT provides subscriber-side ports for connecting Ethernet (RJ-45), telephone wires (RJ-11), or coaxial cable (F-connector).

Hybrid Fiber Coaxial

Multi-System Operators (MSOs) also known as cable TV operators, offer broadband service through their hybrid fiber-coaxial (HFC) network. The HFC network combines optical fiber and coaxial cable to deliver service directly to the customer. Services leave the central office (CO) using a fiber-optic cable. The service is then converted outside of the CO to a coaxial cable tree using a series of optical nodes and, where necessary, through a trunk radio frequency (RF) amplifier. The coaxial cables then connect to multiple subscribers. The demarcation device is a cable modem or set-top box, which talks to a Cable Modem Termination System (CMTS) at the MSO head-end or master facility that receives television signals for processing and distribution. Broadband traffic is carried using the Data Over Cable Service Interface Specification (DOCSIS) standard defined by CableLabs and many contributing companies.

DHCP Overview

You use DHCP in broadband access networks to provide IP address configuration and service provisioning. DHCP, historically a popular protocol in LANs, works well with Ethernet connectivity and is becoming increasingly popular in broadband networks as a simple, scalable solution for assigning IP addresses to subscriber home PCs, set-top boxes (STBs), and other devices.

Junos OS subscriber management supports the following DHCP allocation models:

- DHCP Local Server
- DHCP Relay
- DHCP Relay Proxy

DHCP uses address assignment pools from which to allocate subscriber addresses. Address-assignment pools support both dynamic and static address assignment:

- Dynamic address assignment – A subscriber is automatically assigned an address from the address-assignment pool.
- Static address assignment – Addresses are reserved and always used by a particular subscriber.

NOTE Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

Extended DHCP Local Server

You can enable the services router to function as an extended DHCP local server. As an extended DHCP local server the services router, and not an external DHCP server, provides an IP address and other configuration information in response to a client request. The extended DHCP local server supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Extended DHCP Relay

You can configure extended DHCP relay options on the router and enable the router 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. You can use DHCP relay in carrier edge applications such as video and IPTV to obtain configuration parameters, including an IP address, for your subscribers. The extended DHCP relay agent supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

DHCP Relay Proxy

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits. Except for the ability to add DHCP relay agent options and the gateway address (GIADDR) to DHCP packets, DHCP relay is transparent to DHCP clients and DHCP servers, and simply forwards messages between DHCP clients and servers. When you configure DHCP relay to operate in proxy mode, the relay is no longer transparent. In proxy mode, DHCP relay conceals DHCP server details from DHCP clients, which interact with a DHCP relay in proxy mode as though it is the DHCP server. For DHCP servers there is no change, because proxy mode has no effect on how the DHCP server interacts with the DHCP relay.

Dynamic Profiles for PPP Subscriber Interfaces

Subscriber management Point-to-Point Protocol (PPP) support enables you to create and attach dynamic profiles for PPP subscriber interfaces. When the PPP subscriber logs in, the router instantiates the specified dynamic profile and then applies the attributes defined in the profile to the interface.

Dynamic profiles are used for both static and dynamic PPP interfaces. For static PPP interfaces, you use the CLI to attach dynamic profiles, which specify PPP options. For dynamic PPP interfaces, the dynamic profile creates the interface, including the PPP options.

NOTE Dynamically created interfaces are supported only on Point-to-Point Protocol over Ethernet (PPPoE) interfaces.

Unlike traditional PPP support, subscriber management does not allow bi-directional PPP authentication – authentication is performed only by the router, never by the remote peer. The router’s AAA process manages authentication and address assignment for subscriber management. When you configure PPP options for a dynamic profile, you can configure either Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP) authentication, and you can control the order in which the router negotiates the CHAP and PAP protocols. In addition, for CHAP authentication, you can modify the default length of the CHAP challenge message. Other PPP options, which are either commonly used or mandatory for a traditional PPP interface configuration, are not supported in subscriber management dynamic profiles.

L2TP for Subscriber Access

The Layer 2 Tunneling Protocol (L2TP) is a client-server protocol that allows the Point-to-Point Protocol (PPP) to be tunneled across a network. L2TP encapsulates Layer 2 packets, such as PPP, for transmission across a network. An L2TP access concentrator (LAC), configured on an access device, receives packets from a remote client and forwards them to an L2TP network server (LNS) on a remote network. The LNS functions as the logical termination point of the PPP session tunneled by the LAC from the remote client. Figure 3 shows a simple L2TP topology.

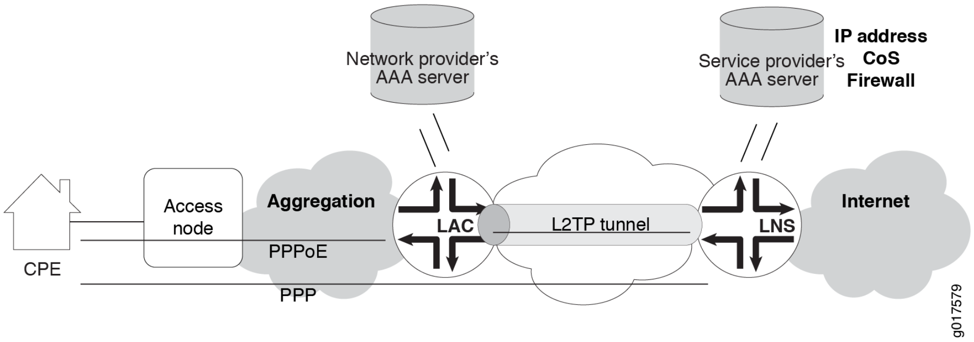

Figure 3 Typical L2TP Topology

L2TP separates the termination of access technologies, such as cable or xDSL, from the termination of PPP and subsequent access to a network. This separation enables public ISPs to outsource their access technologies to competitive local exchange carriers (CLECs). L2TP provides ISPs the capability to supply VPN service; private enterprises can reduce or avoid investment in access technologies for remote workers.

You can configure your router to act as the LAC in PPP pass-through mode in which the LAC receives packets from a remote client and then forwards them at Layer 2 directly to the LNS. The PPP session is terminated on the LNS. This LAC implementation supports only Point-to-Point Protocol over Ethernet (PPPoE) subscribers over dynamic or static logical interfaces.

PPP MPLS Pseudowire Subscriber Logical Interfaces

Subscriber management supports the creation of subscriber interfaces over point-to-point MPLS pseudowires. The pseudowire subscriber interface capability enables service providers to extend an MPLS domain from the access-aggregation network to the service edge, where subscriber management is performed. Service providers can take advantage of MPLS capabilities such as failover, rerouting, and uniform MPLS label provisioning, while using a single pseudowire to service a large number of DHCP and PPPoE subscribers in the service network.

NOTE Pseudowire subscriber logical interfaces are supported on Modular Port Concentrators (MPCs) with Ethernet Modular Interface Cards (MICs) only.

The pseudowire is a tunnel that is either an MPLS-based Layer 2 VPN or Layer 2 circuit. The pseudowire tunnel transports Ethernet encapsulated traffic from an access node (for example, a DSLAM or other aggregation device) to the MX Series router that hosts the subscriber management services. The termination of the pseudowire tunnel on the MX Series router is similar to a physical Ethernet termination, and is the point at which subscriber management functions are performed. A service provider can configure multiple pseudowires on a per-DSLAM basis and then provision support for a large number of subscribers on a specific pseudowire. Figure 4 shows an MPLS network that provides subscriber management support.

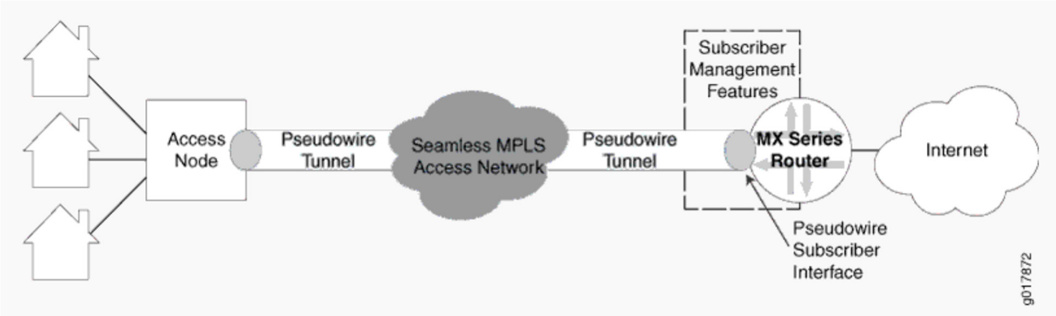

Figure 4 MPLS Access Network with Subscriber Management Support

At the access node end of the pseudowire, the subscriber traffic can be groomed into the pseudowire in a variety of ways, limited only by the number and types of interfaces that can be stacked on the pseudowire. You specify an anchor point, which identifies the logical tunnel interface that terminates the pseudowire tunnel at the access node.

Understanding Subscriber Management VLAN Architecture

The subscriber management logical network architecture is as important as the physical network architecture. You configure the logical portion of the subscriber management network using virtual local area networks (VLANs).

Customer VLANs

Customer VLANs (C-VLANs) provide one-to-one (1:1) subscriber-to-service connectivity; one VLAN carries all traffic to each subscriber on the network. Having a single VLAN per subscriber simplifies operations by providing a 1:1 mapping of technology (VLANs) to subscribers. You can also understand what applications any subscriber is using at any given time. Because you use only one VLAN to carry traffic to each subscriber, this approach is not affected when adding new services. However, using a pure C-VLAN model consumes more bandwidth because a single television channel being viewed by multiple subscribers is carried across the network several times – once on each C-VLAN. This approach requires a more scalable, robust edge router that can support several thousand VLANs.

Configurations that use C-VLANs uniquely identify subscribers by using the VLAN ID and stacked VLAN (S-VLAN) ID. Subscriber packets received from the access node that are either single-tagged with a VLAN ID or double-tagged with both an S-VLAN ID and a VLAN ID are examples of C-VLAN configurations because they provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.

In the C-VLAN architecture, each customer premises equipment (CPE) or subscriber network has its own dedicated Layer 2 path to the router. Each subscriber network is separated by a customer VLAN (C-VLAN) that is dedicated to a particular customer. The services for each customer are transmitted from the router to the access node by means of that customer's C-VLAN.

The ability to uniquely identify subscribers by means of VLAN encapsulation facilitates delivery of services such as authentication, authorization, and accounting (AAA); class of service (CoS); and filters (policers) to subscribers in a C-VLAN configuration.

We recommend using C-VLANs for data and voice traffic to simplify configuration and management when expanding services. However, some MSANs are limited to the number of VLANs they can support, limiting the ability to use C-VLANs.

Service VLANs

Service VLANs (S-VLANs) provide many-to-one (N:1) subscriber-to-service connectivity: The service VLAN carries a service (for example, data, video, or voice) to all subscribers instead of having different services share a VLAN. Adding a new service requires adding a new VLAN and allocating bandwidth to the new service. The service VLAN model enables different groups that are using the broadband network (for example, external application providers) to manage a service. One limitation of service VLANs is the absence of any logical isolation between user sessions at the VLAN level. This lack of isolation requires that the multiservice access node (MSAN) and broadband network gateway (BNG) provide the necessary security filtering.

Service VLANs enable service providers to route different services to different routers to functionally separate network services and reduce network complexity.

Typically, you would use S-VLANs for video and IPTV traffic.

Hybrid VLANs

The hybrid VLAN combines the best of both previous VLANs by using one VLAN per subscriber to carry unicast traffic and one shared multicast VLAN (M-VLAN) for carrying broadcast (multicast) television traffic. You can use both the *pure* and *hybrid* C-VLAN models in different portions of the network, depending upon available bandwidth and MSAN capabilities.

NOTE The term *C-VLAN*, when used casually, often refers to a *hybrid* C-VLAN implementation.

Dynamic 802.1Q VLAN Overview

You can identify VLANs statically or dynamically. You can also configure a mix of static and dynamic VLANs on the same underlying interface.

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces. Many hosts can be connected to the same Gigabit Ethernet switch, but they cannot be in the same routing or bridging domain.

To identify VLANs statically, you can reference a static VLAN interface in a dynamic profile. To identify subscribers dynamically, you use a variable to specify an 802.1Q VLAN that is dynamically created when a subscriber accesses the network.

Dynamic VLAN Configuration

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a particular interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

Dynamic VLAN configuration supports the creation of IPv4 (inet), DHCPv4, IPv6 (inet6), and DHCPv6 VLANs.

Dynamic Mixed VLAN Ranges

Dynamic VLAN and dynamic stacked VLAN configuration supports mixed (or flexible) VLAN ranges. When you configure dynamic mixed VLAN ranges, you must create separate dynamic profiles for VLANs and stacked VLANs. Table 2 lists all valid combinations for the maximum number of dynamic profiles and VLAN and stacked VLAN ranges on a single underlying interface.

VLANs		Stacked VLANs	
Maximum Number of Dynamic Profiles	Maximum Number of VLAN Ranges Per Profile	Maximum Number of Dynamic Profiles	Maximum Number of Stacked VLAN Ranges Per Profile
1	128	1	128
16	32	16	32
1	128	16	32
16	32	1	128

Table 2 Maximum Dynamic Profiles and Ranges for Dynamic Mixed VLAN Configurations

Table 2 shows the valid maximums for the following dynamic mixed VLAN range configuration scenarios, in this order:

- Configurations that require up to 128 VLAN ranges and up to 128 stacked VLAN ranges on a single underlying interface. You must create one VLAN dynamic profile and one stacked VLAN dynamic profile, each with a maximum of 128 ranges per profile.
- Configurations that require up to 32 VLAN ranges and up to 32 stacked VLAN ranges on a single underlying interface. You can configure up to 16 VLAN dynamic profiles and up to 16 stacked VLAN dynamic profiles, each with a maximum of 32 ranges per profile.
- Configurations that consist of one VLAN dynamic profile with a maximum of 128 ranges, and up to 16 stacked VLAN dynamic profiles with 32 ranges each.
- Configurations that consist of up to 16 VLAN dynamic profiles with 32 ranges each, and one stacked VLAN dynamic profile with a maximum of 128 ranges.

The following guidelines apply to the limits in Table 2 when you configure VLAN ranges and S-VLAN ranges for use with dynamic profiles:

- These limits apply to both single-tagged and double-tagged dynamic VLAN ranges.
- These limits apply only to MX Series routers with MPCs. For MX Series routers with Enhanced Queuing IP Services DPCs (DPCE-R-Q model numbers) or Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q model numbers), the maximum number of VLAN ranges for a dynamic profile on an underlying interface remains unchanged at 32 VLAN ranges and 32 S-VLAN ranges.
- These limits have no effect on the maximum number of VLAN IDs on a given underlying interface. The valid range of ID values for a dynamic VLAN range or dynamic S-VLAN range remains unchanged at 1 through 4094.

Understanding Subscriber Interfaces

DHCP Subscriber Interfaces

You can identify subscribers statically or dynamically.

Statically Identifying Subscribers

To identify subscribers statically, you can reference a static VLAN interface in a dynamic profile. Before you can configure static subscriber interfaces in a dynamic profile, you must first configure the logical interfaces on the router to which you expect clients to connect. After you have created the static interfaces, you can modify them by using dynamic profiles to apply configuration parameters.

You can also configure subscribers by creating sets of static IP demux interfaces that are not referenced in a dynamic profile.

When configuring the interfaces stanza within a dynamic profile, you use variables to specify the interface name and the logical unit value. When a DHCP subscriber sends a DHCP request to the interface, the dynamic profile replaces the interface-name and unit variables with the actual interface name and logical unit number of the interface that received the DHCP request. After this association is made, the router configures the interface with any CoS or protocol (that is, IGMP) configuration within the dynamic profile, or applies any input or output filter configuration that you have associated with that dynamic profile.

Dynamically Identifying Subscribers

To identify subscribers dynamically, you create variables for demux interfaces that are dynamically created by DHCP when subscribers log in. You can configure demux interfaces to represent a subscriber interface in a dynamic profile. When a subscriber logs in using a DHCP access method, the demux interface is dynamically created.

You specify variables for the unit number, the name of the underlying interface, and the IP address in the dynamic profile. These variables are replaced with the values that are supplied by DHCP when the subscriber logs in.

Demultiplexing Interfaces

You can create logical subscriber interfaces using static or dynamic demultiplexing interfaces. In addition, you can use either IP demultiplexing interfaces or VLAN demultiplexing interfaces when creating logical subscriber interfaces.

Demultiplexing (demux) interfaces are logical interfaces that share a common, underlying logical interface (in the case of IP demux) or underlying physical interface (in the case of VLAN demux). You can use these interfaces to identify specific subscribers or to separate individual circuits by IP address (IP demux) or VLAN ID (VLAN demux).

The subscriber interfaces can provide different levels of services for individual subscribers in an access network. For example, you can apply CoS parameters for each subscriber.

Interface Sets of Static Demux Interfaces

You can group static demux interfaces to create individual subscriber interfaces using interface sets. Interface sets enable you to provide the same level of service for a group of subscribers; for example, all residential subscribers who receive the basic data service.

Figure 5 shows a subscriber interface configured using a set of IP demux interfaces with an underlying VLAN interface.

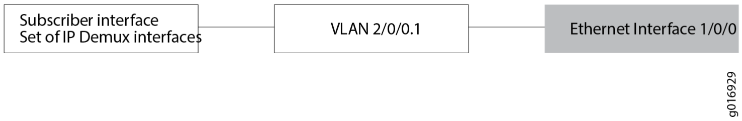

Figure 5 IP Demux Subscriber Interface

Dynamic Demultiplexing Interfaces

You can configure demux interfaces to represent a dynamic subscriber interface in a dynamic profile.

Demux interfaces are dynamically created by a DHCP access method when the underlying interface for the demux interface is configured for the access method. The DHCP access model creates the demux interface with the subscriber's assigned IP address (for IP demux interfaces) or VLAN ID (for VLAN demux interfaces).

To configure an IP demux interface in the dynamic profile, you specify variables for the unit number, the name of the underlying interface, and the IP address. To configure a VLAN demux interface in the dynamic profile, you specify variables for the unit number, the name of the underlying interface, and the VLAN ID. These variables are replaced with the values that are supplied by DHCP when the subscriber logs in.

Guidelines for Configuring Demux Interfaces for Subscriber Access

When you configure static or dynamic demux interfaces for subscriber access, consider the following guidelines:

- Hierarchical and per-unit scheduling is supported for dynamically created demux interfaces on the EQ DPC.
- IP demux interfaces support IPv4 (family inet) and IPv6 (family inet6)).
- IP demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only MPCs installed. If the router has other cards in addition to MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.
- You can configure IPv4 and IPv6 addressing for static and dynamic demux interfaces.
- You can configure only one demux0 interface per chassis.
- For IP demux interfaces, you can define logical demux interfaces on top of the demux0 interface (for example, demux0.1, demux0.2, and so on).
- Demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.
- You must associate IP demux interfaces with an underlying logical interface.
- You must associate VLAN demux interfaces with an underlying device (physical interface).
- You cannot use a dynamic demux interface to represent multiple subscribers in a dynamic profile attached to an interface. One dynamic demux interface represents one subscriber. Do not configure the aggregate-clients option when attaching a dynamic profile to a demux interface for DHCP.

PPPoE Subscriber Interfaces

You can configure the router to dynamically create Point-to-Point Protocol over Ethernet (PPPoE) logical interfaces on statically created underlying Ethernet interfaces. The router creates the dynamic interface in response to the receipt of a PPPoE Active Discovery Request (PADR) control packet on the underlying interface. Because the router creates a dynamic PPPoE logical interface on demand when a subscriber logs in to the network, dynamic PPPoE logical interfaces are also referred to as *dynamic PPPoE subscriber interfaces*.

Configuration of dynamic PPPoE subscriber interfaces over static underlying Ethernet interfaces is supported on MPC/MIC interfaces on MX Series 3D Universal Edge Routers.

Configuring and using dynamic PPPoE subscriber interfaces offers the following benefits:

- On-demand dynamic interface creation
- Dynamic PPPoE subscriber interfaces provides the flexibility of dynamically creating the PPPoE subscriber interface only when needed; that is, when a subscriber logs in on the associated underlying Ethernet interface. By contrast, statically created interfaces allocate and consume system resources when the interface is created. Configuring and using dynamically created interfaces helps you effectively and conveniently manage edge or access networks in which large numbers of subscribers are constantly logging in to and logging out from the network on a transient basis.
- Dynamic removal of PPPoE subscriber interfaces without manual intervention
- When the PPPoE subscriber logs out or the PPPoE session is terminated, the router dynamically deletes the associated PPPoE subscriber interface without your intervention, thereby restoring any consumed resources to the router.
- Use of dynamic profiles to efficiently manage multiple subscriber interfaces
- By using a profile, you reduce the management of a large number of interfaces by applying a set of common characteristics to multiple interfaces. When you configure a dynamic profile for PPPoE, you use predefined dynamic variables in the profile to represent information that varies from subscriber to subscriber, such as the logical unit number and underlying interface name. These variables are dynamically replaced with the values supplied by the network when the subscriber logs in.
- Denial of service (DoS) protection
- You can configure the underlying Ethernet interface with certain PPPoE-specific attributes that can reduce the potential for DoS attacks. Duplicate protection, which is disabled by default, prevents activation of another dynamic PPPoE logical interface on the underlying interface when a PPPoE logical interface for the same client is already active on the underlying interface. You can also specify the maximum number of PPPoE sessions that the router can activate on the underlying interface. By enabling duplicate protection and restricting the maximum number of PPPoE sessions on the underlying interface, you can ensure that a single toxic PPPoE client cannot monopolize allocation of the PPPoE session.

- Support for dynamic PPPoE subscriber interface creation from PPPoE service name tables
- You can assign a previously configured PPPoE dynamic profile to a named, empty, or any service entry in a PPPoE service name table, or to an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services. The router uses the attributes defined in the profile to instantiate a dynamic PPPoE subscriber interface based on the service name, ACI, and ARI information provided by the PPPoE client during PPPoE negotiation. To specify the routing instance in which to instantiate the dynamic PPPoE subscriber interface, you can assign a previously configured routing instance to a named, empty, or any service, or to an ACI/ARI pair defined for these services. The dynamic profile and routing instance configured for the PPPoE service name table overrides the dynamic profile and routing instance assigned to the PPPoE underlying interface on which the dynamic subscriber interface is created.

MLPPP Support for LNS and PPPoE Subscribers

Multilink Point-to-Point Protocol (MLPPP) aggregates multiple PPP physical links into a single virtual connection, or logical bundle. More specifically, MLPPP bundles multiple link-layer channels into a single network-layer channel. Peers negotiate MLPPP during the initial phase of Link Control Protocol (LCP) option negotiation. Each router indicates that it is multilink capable by sending the multilink option as part of its initial LCP configuration request.

An MLPPP bundle can consist of multiple physical links of the same type – such as multiple asynchronous lines – or can consist of physical links of different types – such as leased synchronous lines and dial-up asynchronous lines.

Packets received with an MLPPP header are subject to fragmentation, reassembly, and sequencing. Packets received without the MLPPP header cannot be sequenced and can be delivered only on a first-come, first-served basis.

MLPPP is used to bundle multiple low speed links to create a higher bandwidth pipe such that the combined bandwidth is available to traffics from all links, and to support link fragmentation and interleaving (LFI) support on the bundle to reduce the transmission delay of high priority packets. LFI interleaves voice packets with fragmented data packets to ensure timely delivery of voice packets. Figure 6 shows how incoming packets are distributed and aggregated into an MLPPP bundle.

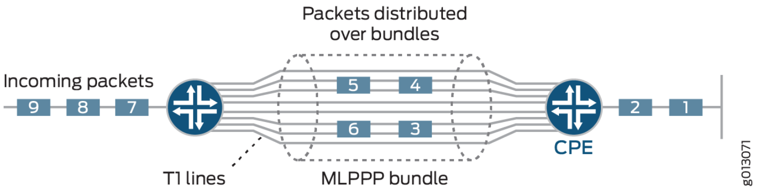

Figure 6 *MLPPP Aggregation of Traffic Into Single Bundle*

Starting in Junos OS Release 14.1, multilink PPP (MLPPP) support is provided to LNS (L2TP network server) and PPPoE (Point-to-Point Protocol over Ethernet) terminated and tunneled subscribers running on MX Series with access-facing MPC2s.

For customers with both MLPPP and single link PPP clients, the router needs to determine client capability during link control protocol (LCP) negotiation and support either multilink or single link access modules accordingly (mixed mode support).

Member Link and Bundle Configuration

An MLPPP subscriber consists of two IFLs (logical interfaces), a member link, and a bundle. For MLPPP subscribers, you can configure the member link and bundle statically, or dynamically using dynamic profiles.

Static MLPPP Subscribers – You must configure both member link and bundle IFLs manually before the member link IFL can start LCP (link control protocol) negotiation either for an LNS session or for a PPPoE session.

Dynamic MLPPP Subscribers – You configure dynamic member IFLs using dynamic profiles. The member link dynamic profile includes the family mlppp statement containing the bundle dynamic profile and the service interface (si), or a pool of service interfaces. This information is then used to create the dynamic bundle IFL.

Each bundle accepts only one member link. If more than one member link attempts to join the same bundle, the system fails the new member session. Dual-stack is supported for the bundle.

LNS Subscribers and MX Series

Figure 7 shows a network diagram with the MX Series functioning as the LNS. Both PPP and MLPPP bundles are terminated at the LNS. The following three domains are shown passing traffic through the LNS network:

- PPP domain – Contains data and voice traffic
- MLPPP domain – Contains data traffic only
- L2TP domain – Contains all types of traffic

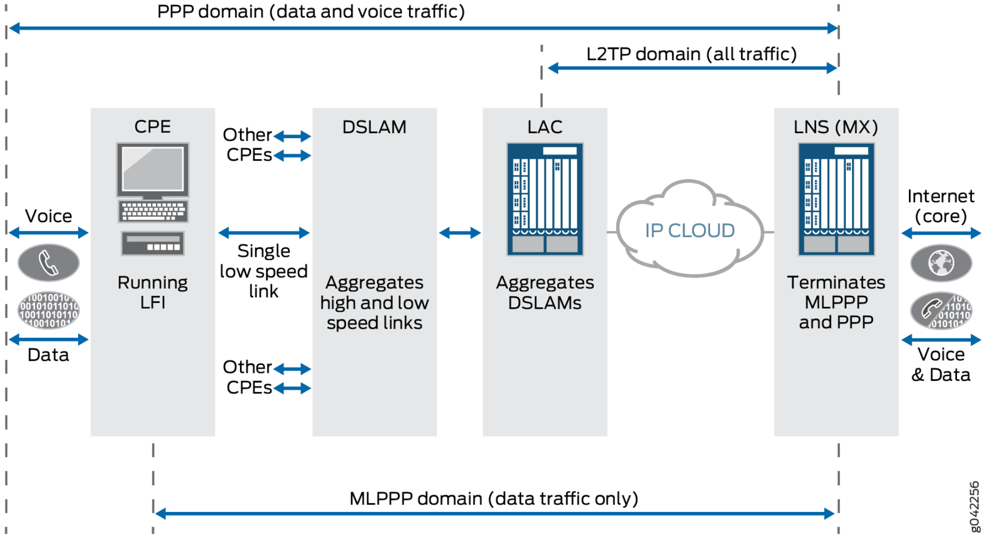

Figure 7 MLPPP Bundles Terminated at MX Series as the LNS Network

ATM for Subscriber Access

By using the ATM Modular Interface Card (MIC) with small form-factor pluggable transceiver (SFP) and a supported Modular Port Concentrator (MPC), you can configure the MX Series router to support configurations that enable subscribers to access the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs). Using these configurations enables the delivery of subscriber-based services, such as class of service (CoS) and firewall filters, for subscribers accessing the router over an ATM network.

On MX Series routers with MPC/MIC interfaces that use the ATM MIC with SFP (Model Number MIC-3D-8OC3-2OC12-ATM), you can create the following configurations to enable subscribers to access the router over an ATM network using ATM Adaptation Layer 5 (AAL5) permanent virtual connections (PVCs):

- PPP-over-Ethernet-over-ATM
- Routed IP-over-ATM
- Bridged IP-over-Ethernet-over-ATM
- PPP-over-ATM
- Concurrent PPP-over-Ethernet-over-ATM interfaces and IP-over-Ethernet-over-ATM interfaces on a single ATM PVC

Understanding Subscriber Sessions

RADIUS Server Options for Subscriber Access

You can specify options that the router uses when communicating with RADIUS authentication and accounting servers for subscriber access.

The following list describes the RADIUS options you can configure:

- access-loop-id-local – The Agent-Remote-Id and Agent-Circuit-Id generated locally when these values are not present in the client database. The interface description of the logical interface is used as the Agent-Remote-Id and the interface description portion of the NAS-Port-Id using the format <underlying-interface-name>:<outer-tag>-<inner-tag> is used as the Agent-Circuit-Id.

NOTE The NAS-Port-Id format changes (established by [set access profile *profile-name* radius options interface-description-format]) are applied before generating the Agent-Circuit-Id.

The NAS-Port-Id format (established by [set access profile *profile-name* radius options interface-description-format]) leverages the locally generated Agent-Remote-Id and Agent-Circuit-Id.

- accounting-session-id-format – The format the router uses to identify the accounting session. The identifier can be in one of the following formats:
- decimal – The default format. For example, 435264
- description – In the format, *jnpr interface-specifier*:*subscriber-session-id*. For example, jnpr fastEthernet 3/2.6:1010101010101
- calling-station-id-delimiter – The character that the router uses as the separator between concatenated values in the Calling-Station-Id string (RADIUS attribute 31).
- calling-station-id-format – Optional information that the router includes in the Calling-Station-Id (RADIUS attribute 31).
- client-accounting-algorithm and client-authentication-algorithm – The method the router uses to access RADIUS accounting and RADIUS authentication servers. You can specify the following methods:
- direct – The default method, in which there is no load balancing. For example, in the direct method, the router always accesses server1 (the primary server) first, and uses server2 and server3 as backup servers.
- round-robin – The method that provides load balancing by rotating router requests among the list of configured RADIUS servers. For example, if three RADIUS servers are configured to support the router, the router sends the first request to server1, and uses server2 and server3 as backup servers. The router then sends the second request to server2, and uses server3 and server1 as backups.

NOTE When a RADIUS server in the round-robin list becomes unreachable, the next reachable server in the round-robin list is used for the current request. That same server is also used for the next request because it is at the top of the list of available servers. As a result, after a server failure, the server that is used takes up the load of two servers.

- coa-dynamic-variable-validation – The optional method that the router uses when processing CoA requests that include changes to a client profile dynamic variable that cannot be applied. The optional configuration specifies that when a CoA operation is unable to apply a requested change to a client profile dynamic variable, subscriber management does not apply any changes to client profile dynamic variables in the CoA request and then responds with a NACK. In the default method, subscriber management does not apply the incorrect update but does apply the other changes to the client profile dynamic variables, and then responds with an ACK message.
- ethernet-port-type-virtual – The physical port type of virtual that the router uses to authenticate clients. The port type is passed in RADIUS attribute 61 (NAS-Port-Type). By default the router passes a port type of ethernet in RADIUS attribute 61.
- interface-description-format – The information that is excluded from the interface description that the router passes to RADIUS for inclusion in the RADIUS attribute 87 (NAS-Port-Id). By default, the router includes both the subinterface and the adapter in the interface description. You can specify:
- exclude-adapter – Exclude the adapter.
- exclude-subinterface – Exclude the subinterface.
- nas-identifier – The value for the client RADIUS attribute 32 (NAS-Identifier), which is used for authentication and accounting requests. You can specify a string in the range 1 through 64 characters.
- nas-port-extended-format – The extended format for RADIUS attribute 5 (NAS-Port) and for the width of the fields in the NAS-Port attribute that the RADIUS client uses. You can specify:
- adapter-width *width* – Number of bits in the adapter field.
- port-width *width* – Number of bits in the port field.
- pw-width – Number of bits in the pseudowire field.
- slot-width *width* – Number of bits in the slot field.
- stacked-vlan-width *width* – Number of bits in the SVLAN ID field.
- vlan-width *width* – Number of bits in the VLAN ID field.

NOTE The total of the widths must not exceed 32 bits, or the configuration fails.

You can configure an extended format for the NAS-Port attribute for both Ethernet subscribers and ATM subscribers. For ATM subscribers, you can specify:

- adapter-width – Number of bits in the ATM adapter field, in the range 0 through 32

- port-width – Number of bits in the ATM port field, in the range 0 through 32
- slot-width – Number of bits in the ATM slot field, in the range 0 through 32
- vci-width – Number of bits in the ATM virtual circuit identifier (VCI) field, in the range 0 through 32
- vpi-width – Number of bits in the ATM virtual path identifier (VPI) field, in the range 0 through 32

NOTE For ATM subscribers, the combined total of the widths of all fields must not exceed 32 bits, or the configuration fails. The router may truncate the values of individual fields depending on the bit width you specify.

- nas-port-id-delimiter – The character used as the separator between values in the NAS-Port-Id string.
- nas-port-id-format – Optional information included in RADIUS attribute 87 (NAS-Port-Id).
- nas-port-type – The port type used to authenticate subscribers.
- revert-interval – The number of seconds that the router waits after a server has become unreachable. The router rechecks the connection to the server when the revert-interval expires. If the server is then reachable, it is used in accordance with the order of the server list. You can configure from 0 (off) through 604800 seconds. The default is 60 seconds.
- service-activation – Setting that determines whether newly authenticated subscriber can successfully log in when service activation failures related to configuration errors occur during authd processing of the activation request for the subscriber’s address family. You can specify this behavior for services configured in dynamic profiles (dynamic-profile) or in Extensible Subscriber Services Manager (ESSM) operation scripts (extensible-service):
- optional-at-login – Service activation is optional. Activation failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Service activation failures due to causes other than configuration errors cause network family activation to fail. The login attempt is terminated unless another address family is already active for the subscriber. This is the default behavior for the extensible-service service type.
- required-at-login – Service activation is required. Activation failure for any reason causes network family activation to fail. The login attempt is terminated unless another address family is already active for the subscriber. This is the default value for the dynamic-profile service type.
- vlan-nas-port-stacked-format – The format that turns off RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

Global RADIUS Options for Subscriber Access

You can specify options that the router uses when communicating with all configured RADIUS servers for subscriber access.

The following list describes the global RADIUS options you can configure:

- revert-interval – The number of seconds that the router waits after a server has become unreachable. The router rechecks the connection to the server when the revert-interval expires. If the server is then reachable, it is used in accordance with the order of the server list. You can configure from 0 (off) through 604800 seconds. The default is 60 seconds.
- request-rate – The number of requests per second that the router can send to all configured RADIUS servers collectively. By limiting the flow of requests from the router to the RADIUS servers, you can prevent the RADIUS servers from being flooded with requests. You can configure from 100 through 4000 requests per second. The default is 500 requests per second.

RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework

The AAA Service Framework supports RADIUS attributes and vendor-specific attributes (VSAs). This support provides tunable parameters that the subscriber access management feature uses when creating subscribers and services.

RADIUS attributes are carried as part of standard RADIUS request and reply messages. The subscriber management access feature uses the RADIUS attributes to exchange specific authentication, authorization, and accounting information. VSAs allow the subscriber access management feature to pass implementation-specific information that provide extended capabilities, such as service activation or deactivation, and enabling and disabling filters.

When you use dynamic profiles, the AAA Service Framework supports the use of Junos OS predefined variables to specify the RADIUS attribute or VSA for the information obtained from the RADIUS server.

Specifying the Authentication and Accounting Methods for Subscriber Access

You can specify the authentication and accounting methods that subscriber access management uses.

You can configure multiple authentication and accounting methods – the authentication-order and accounting order statements specify the order in which the subscriber access management feature uses the methods. For example, an authentication entry of RADIUS password specifies that RADIUS authentication is performed first and, if it fails, local authentication (password) is done.

You can specify the following authentication methods:

NOTE You must always specify the RADIUS authentication method. Subscriber access management does not support the password keyword (the default), and authentication fails when no method is specified.

- password – Local authentication
- radius – RADIUS-based authentication

You can specify the following accounting methods:

- radius – RADIUS-based accounting

To configure the authentication and accounting methods for subscriber access management:

- Specify the authentication methods and the order in which they are used. Only radius is supported.

```
[edit access profile isp-bos-metro-fiber-basic]
user@host# set authentication-order radius
```

- Specify the accounting method.

```
[edit access profile isp-bos-metro-fiber-basic]
user@host# set accounting order radius
```

RADIUS Accounting Statistics for Subscriber Access

The AAA Service Framework enables you to configure how the router collects and uses accounting statistics for subscriber management.

For example, you can specify when statistics collection is terminated, the order in which different accounting methods are used, the types of statistics collected, and how often statistics are collected. You can also configure the router to request that the RADIUS server immediately update the accounting statistics when certain events occur, such as when a subscriber logs in or when a change of authorization (CoA) occurs.

Subscriber management provides two levels of subscriber accounting – subscriber session and service session. In subscriber session accounting, the router collects statistics for the entire subscriber session. In service session accounting, the router collects statistics for specific service sessions for the subscriber.

NOTE Subscriber management counts forwarded packets only. Dropped traffic (for example as a result of a filter action) and control traffic are not included in the accounting statistics.

The router uses the RADIUS attributes and Juniper Networks VSAs listed in Table 3 to provide the accounting statistics for subscriber and service sessions. If the session has both IPv4 and IPv6 families enabled, the router reports statistics for both families.

NOTE RADIUS reports subscriber statistics as an aggregate of both IPv4 statistics and IPv6 statistics.

For an IPv4-only configuration, the standard RADIUS attributes report the IPv4 statistics and the IPv6 VSA results are all reported as 0.

For an IPv6-only configuration, the standard RADIUS attributes and the IPv6 VSA statistics are identical, both reporting the IPv6 statistics.

When both IPv4 and IPv6 are configured, the standard RADIUS attributes report the combined IPv4 and IPv6 statistics. The IPv6 VSAs report IPv6 statistics.

Attribute Number	Attribute Name	Type of Statistics
26-151	IPv6-Acct-Input-Octets	IPv6
26-152	IPv6-Acct-Output-Octets	IPv6
26-153	IPv6-Acct-Input-Packets	IPv6
26-154	IPv6-Acct-Output-Packets	IPv6
26-155	IPv6-Acct-Input-Gigawords	IPv6
26-156	IPv6-Acct-Output-Gigawords	IPv6
47	Acct-Input-Packets	IPv4 and IPv6 aggregation
48	Acct-Output-Packets	IPv4 and IPv6 aggregation
52	Acct-Input-Gigawords	IPv4 and IPv6 aggregation
53	Acct-Output-Gigawords	IPv4 and IPv6 aggregation

Table 3 RADIUS Attributes and VSAs Used for Per-Subscriber Session Accounting

Understanding Session Options for Subscriber Access

You can configure several characteristics of the sessions that are created for DHCP, L2TP, and terminated PPP subscribers. You can place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both. You can also set parameters that modify a subscriber's username at login based on the subscriber's access profile.

Subscriber Session Timeouts

You can limit subscriber access by configuring a session timeout or an idle timeout. Use a session timeout to specify a fixed period of time that the subscriber is permitted to have access. Use an idle timeout to specify a maximum period of time that the subscriber can be idle. You can use these timeouts separately or together. By default, neither timeout is present.

NOTE For all subscriber types other than DHCP (such as L2TP-tunneled and PPP-terminated subscribers), the session timeout value limits the subscriber session. For DHCP subscribers, the session timeout value is used to limit the lease when no other lease time configuration is present. The lease expires when the timeout value expires. If this value is not supplied by either the CLI or RADIUS, the DHCP lease does not expire.

The idle timeout is based on accounting statistics for the subscriber. The router determines subscriber inactivity by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction.

Optionally, you can specify that only subscriber ingress traffic is monitored; egress traffic is ignored. This configuration is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore cannot detect that the peer is not up. In this situation, because by default the LAC monitors both ingress and egress traffic, it detects the egress traffic from the LNS and either does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored, the LAC can detect that the peer is inactive and then initiate logout.

When either timeout period expires, the non-DHCP subscribers are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout. DHCP subscribers are disconnected. The Acct-Terminate-Cause [RADIUS attribute 49] value includes a reason code of 5 for a session timeout and a code of 4 for an idle timeout.

You can configure these limitations to subscriber access on a per-subscriber basis by using the RADIUS attributes Session-Timeout [27] and Idle-Timeout [28]. RADIUS returns these attributes in Access-Accept messages in response to Access-Request messages from the access server.

Service providers often choose to apply the same limitations to large numbers of subscribers. You can reduce the RADIUS provisioning effort for this scenario by defining the limitations for subscribers in an access profile on a per-routing-instance basis. If you do so, RADIUS attributes subsequently returned for a particular subscriber logged in with the profile override the per-routing-instance values.

BEST PRACTICE We recommend that you do not configure a session timeout for subscribers receiving voice services. Because the session timeout is based only on time and not user activity, it is likely to interrupt subscribers actively using a voice service and terminate their calls unexpectedly (from the subscriber viewpoint). This result is a particular concern for emergency services calls.

BEST PRACTICE We recommend that you do not configure an idle timeout for DHCP subscribers. When the timeout expires with no activity and the connection is terminated, the protocol has no means to inform the client. Consequently, these subscribers are forced to reboot their CPE device the next time they attempt to access the Internet. Contrast the behavior when an idle timeout is configured for PPP subscribers. In this case, timeout expiration causes PPP to terminate the link with the peer. Depending on the CPE device, this termination enables the peer to automatically retry the connection either on demand or immediately. In either case, no subscriber intervention is required.

The available range for setting a timeout is the same whether you configure it in the CLI or through the RADIUS attributes:

- Session timeouts can be set for 1 minute through 527,040 minutes in the CLI and the corresponding number of seconds (60 through 31,622,400) in the Session-Timeout attribute [27].
- Idle timeouts can be set for 10 minutes through 1440 minutes in the CLI and the corresponding number of seconds (600 through 86,400) in the Idle-Timeout attribute [28].

The router interprets the values in the attributes to conform to the supported ranges. For example, for Session-Timeout [27]:

- A value of zero is treated as no timeout.
- A value in the range 1 through 59 is raised to 60 seconds.
- A value that exceeds 31,622,400 is reduced to 31,622,400 seconds.
- For Idle-Timeout [28]:
- A value of zero is treated as no timeout.
- A value in the range 1 through 599 is raised to 600 seconds.
- A value that exceeds 86,400 is reduced to 86,400 seconds.

In configurations using dynamically created subscriber VLANs, the idle timeout also deletes the inactive subscriber VLANs when the inactivity threshold has been reached. In addition to deleting inactive dynamic subscriber VLANs, the idle timeout also removes dynamic VLANs when no client sessions were ever created (for example, in the event no client sessions are created on the dynamic VLAN or following the occurrence of an error during session creation or client authentication where no client sessions are created on the dynamic VLAN).

Session and idle timeouts for deleting dynamic subscriber VLANs are useful only in very limited use cases; typically neither timeout is configured for this purpose.

A possible circumstance when they might be useful is when the dynamic VLANs have no upper layer protocol that helps determine when the VLAN is removed with the remove-when-no-subscribers statement; for example, when the VLAN is supporting IP over Ethernet without DHCP in a business access model with fixed addresses. However, business access is generally a higher-tier service than residential access and as such typically is not subject to timeouts due to inactivity as might be desired for residential subscribers.

An idle timeout might be appropriate in certain Layer 2 wholesale situations, where the connection can be regenerated when any packet is received from the CPE.

When using the idle timeout for dynamic VLAN removal, keep the following in mind:

- The idle timeout period begins after a dynamic subscriber VLAN interface is created or traffic activity stops on a dynamic subscriber VLAN interface.
- If a new client session is created or a client session is reactivated successfully, the client idle timeout resets.
- The removal of inactive subscriber VLANs functions only with VLANs that have been authenticated.

Subscriber Username Modification

For Layer 2 wholesale applications, some network service providers employ username modification to direct subscribers to the appropriate retail enterprise network. This modification is also called username *stripping*, because some of the characters in the username are stripped away and discarded. The remainder of the string becomes the new, modified username. The modified username is used by an external AAA server for session authentication and accounting. The modification parameters are applied according to a subscriber access profile that also determines the subscriber and session context; that is, the logical system:routing instance (LS:RI) used by the subscriber. Only the default (master) logical system is supported. Because the wholesaler differentiates between multiple retailers by placing each in a different LS:RI, the usernames are appropriately modified for each retailer.

You can select up to eight characters as delimiters to mark the boundary between the discarded and retained portions of the original username; there is no default delimiter. The portion of the name to the right of the selected delimiter is discarded along with the delimiter. By configuring multiple delimiters, a given username structure can result in different modified usernames. You can configure the direction in which the original name is parsed to determine which delimiter marks the boundary. By default, the parse direction is from left to right.

Domain Mapping Overview

Domain mapping enables you to configure a map that specifies access options and session-specific parameters. The map is based on the domain name of subscriber sessions – the router applies the mapped options and parameters to sessions for subscribers that have the specified domains. For example, you might configure a domain map that is based on the domain name example.com. The options and parameters in that domain map are then applied when subscribers with the specified domain name (for example, bob@example.com, raj@example.com, and juan@example.com) request an AAA service.

NOTE A subscriber's username is typically made up of two parts – the user's name followed by the user's domain name, which are separated by a delimiter character. The domain name is always to the right of the domain delimiter. For example, in the username, juan@example.com, the user's name, juan is followed by the domain name example.com, and the two are separated by the @ delimiter character.

However, some systems use a username format in which the domain name precedes the user's name. To avoid confusion with the typical domain name usage, this type of preceding domain name is referred to as a realm name, and the realm name is to the left of the realm delimiter. For example, in the username, top321-example.com/mary, the top321-example.com part is the realm name, mary is the user's name, and the / character is the delimiter character.

The domain map provides efficiency, and enables you to make changes for a large number of subscribers in one operation. For example, if an address assignment pool becomes exhausted due to the number of subscribers obtaining addresses from the pool, you can create a domain map that specifies that subscribers in a particular domain obtain addresses from a different pool. In another use of the domain map, you might create a new dynamic profile and then configure the domain map to specify which subscribers (by their domain) use that dynamic profile.

NOTE Subscriber management is supported in the default logical system only. The documentation for the subscriber management domain mapping feature describes using the aaa-logical-system and target-logical-system statements to configure mapping to a non-default logical system. These statements are for future extensions of subscriber management.

Table 4 describes the access options and parameters you can configure in the domain map.

Options	Description
AAA logical system/ routing instance	Logical system/routing instance in which AAA sends authentication and accounting requests for the subscriber sessions. Subscriber management is supported in the default logical system only.
Access profile	Access profile applied to subscriber sessions.
Address pool	Address pool used to allocate addresses to subscribers.
Domain and realm name rules	Rules for domain and realm name usage, including domain name stripping, supported delimiters, and parse direction (delimiters and the parse direction are configured globally).
Dynamic profile	Dynamic profile applied to subscriber sessions.
PADN parameters	PPPoE route information for subscriber sessions.
Target logical system/ routing instance	Logical system/routing instance to which the subscriber interface is attached. Subscriber management is supported in the default logical system only.
Tunnel profile	Tunnel profile applied to subscriber sessions.

Table 4 Domain Map Options and Parameters

Types of Domain Maps and Order of Precedence

Starting in Junos OS Release 16.1, subscriber management uses a specific order when searching for a domain map that matches the subscriber domain name. The following list shows that order:

- Exact match domain map – The subscriber domain name is an exact match to a configured domain map.
- Wildcard domain map – The subscriber domain name is a partial match to a wildcard domain map.
- default domain map – The subscriber domain name is neither an exact match nor a partial wildcard match to a domain map.

NOTE If the subscriber username does not have a domain name, then no search is performed and the subscriber is associated with the none domain map, if configured.

Using RADIUS Dynamic Requests for Subscriber Access Management

RADIUS dynamic requests provide an efficient way to centrally manage subscriber sessions. The AAA Service Framework's RADIUS dynamic request support allows RADIUS servers to initiate user-related operations, such as a termination operation, by sending unsolicited request messages to the router. Without the RADIUS dynamic request feature, the only way to disconnect a RADIUS user is from the router, which can be cumbersome and time-consuming in large networks.

In a typical client-server RADIUS environment, the router functions as the client and initiates requests sent to the remote RADIUS server. However, when using RADIUS dynamic requests, the roles are reversed. For example, during a disconnect operation, the remote RADIUS server performs as the client and initiates the request (the disconnect action) – the router functions as the server in the relationship.

You create an access profile to configure the router to support RADIUS dynamic requests. This configuration enables the router to receive and act on the following types of messages from remote RADIUS servers:

- Access-Accept messages – Dynamically activate services based on attributes in RADIUS Access-Accept messages received when a subscriber logs in.
- Change-of-Authorization (CoA) messages – Dynamically modify active sessions based on attributes in CoA messages. CoA messages can include service creation requests, deletion requests, RADIUS attributes, and Juniper Networks VSAs.
- Disconnect messages – Immediately terminate specific subscriber sessions.

Extended DHCP Local Server Overview

Junos OS includes an extended DHCP local server that enhances traditional DHCP server operation by providing additional address assignment and client configuration functionality and flexibility in a subscriber-aware environment. The extended DHCP local server enables service providers to take advantage of external address-assignment pools and integrated RADIUS-based configuration capabilities in addition to the continued support of traditional local address pools. The address-assignment pools are considered external because they are external to the DHCP local server. The pools are managed independently of the DHCP local server, and can be shared by different client applications, such as DHCP or PPPoE access. Table 5 provides a comparison of the extended DHCP local server and a traditional DHCP local server.

The extended DHCP local server provides an IP address and other configuration information in response to a client request. The server 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 DHCP client authentication. You can configure the dynamic profile and authentication support on a global basis or for a specific group of interfaces.

Feature	Extended DHCP Local Server	Traditional DHCP Local Server
Local address pools	X	X
External, centrally-managed address pools	X	–
Local configuration	X	X

External configuration using information from address-assignment pools or RADIUS servers	X	–
Dynamic-profile attachment	X	–
RADIUS-based subscriber authentication, and configuration using RADIUS attributes and Juniper Networks VSAs	X	–
IPv6 client support	X	–
Default minimum client configuration	X	X

Table 5 Comparing the Extended DHCP Local Server to the Traditional DHCP Local Server

You can also configure the extended DHCP local server to support IPv6 clients. Both DHCP local server and DHCPv6 local server support the specific address request feature, which enables you to assign a particular address to a client.

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.

Extended DHCP Relay Agent Overview

You can configure extended DHCP relay options on the router or on the switch and 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 DHCP client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

NOTE The PTX Series Packet Transport Routers do not support authentication for DHCP relay agents.

On the routers, you can use DHCP relay in carrier edge applications such as video/IPTV to obtain configuration parameters, including an IP address, for your subscribers.

On the switches, you can use DHCP relay to obtain configuration parameters including an IP address for DHCP clients.

NOTE The extended DHCP relay agent options configured with the dhcp-relay statement are incompatible with the DHCP/BOOTP relay agent options configured with the bootp statement. As a result, you cannot enable both the extended DHCP relay agent and the DHCP/BOOTP relay agent on the router at the same time.

You can also configure the extended DHCP relay agent to support IPv6 clients.

To configure the extended DHCP relay agent on the router (or switch), include the dhcp-relay statement at the [edit forwarding-options] hierarchy level.

You can also include the dhcp-relay statement at the following hierarchy levels:

- [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]

DHCP Relay Proxy Overview

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits.

Normally, extended DHCP relay operates as a helper application for DHCP operations. Except for the ability to add DHCP relay agent options and the gateway address (giaddr) to DHCP packets, DHCP relay is transparent to DHCP clients and DHCP servers, and simply forwards messages between DHCP clients and servers.

When you configure DHCP relay to operate in proxy mode, the relay is no longer transparent. In proxy mode, DHCP relay conceals DHCP server details from DHCP clients, which interact with a DHCP relay in proxy mode as though it is the DHCP server. For DHCP servers there is no change, because proxy mode has no effect on how the DHCP server interacts with the DHCP relay.

DHCP relay proxy provides the following benefits:

- DHCP server isolation and DoS protection – DHCP clients are unable to detect the DHCP servers, learn DHCP server addresses, or determine the number of servers that are providing DHCP support. Server isolation also provides denial-of-service (DoS) protection for the DHCP servers.
- Multiple lease offer selection – DHCP relay proxy receives lease offers from multiple DHCP servers and selects a single offer to send to the DHCP client, thereby reducing traffic in the network. Currently, the DHCP relay proxy selects the first offer received.
- Support for both numbered and unnumbered Ethernet interfaces – For DHCP clients connected through Ethernet interfaces, when the DHCP client obtains an address, the DHCP relay proxy adds an access internal host route specifying that interface as the outbound interface. The route is automatically removed when the lease time expires or when the client releases the address.
- Logical system support – DHCP relay proxy can be configured in a logical system, whereas a non-proxy mode DHCP relay cannot.

NOTE You cannot configure both DHCP relay proxy and extended DHCP local server on the same interface.

Default Subscriber Service Overview

Subscriber management enables you to specify a default subscriber service for DHCP subscribers. The default service (dynamic profile) is applied to subscribers when the subscriber logs in. By configuring a default service, you can apply a particular service (for example, a basic service) to subscribers who are not explicitly assigned a service.

When a subscriber logs in, the configured default service is always activated, even when remote service provisioning or RADIUS service activation is configured for the subscriber. The default service is deactivated only when the subscriber is successfully provisioned by the PCRF by means of the GX-Plus application. (Remote provisioning is configured by the provisioning-order statement at the [edit access profile] hierarchy level.)

In all other cases, the default service remains active. For example, if RADIUS authentication is configured but service activation is not, the default subscriber service remains activated. Likewise, if RADIUS authentication is not configured, the default subscriber service remains activated.

Default services can also be deactivated either with a RADIUS CoA deactivate request or with the `request network-access aaa subscriber delete session-id` command.

To create and assign a default subscriber service, you must complete the following operations:

- Create the service – Ensure that the service you want to use has been configured in a dynamic profile. The actual service is no different than any other service used for subscriber management.
- Specify the default service – Use the Junos OS CLI to specify the service that is used as the default service.
- Specify the interfaces on which the default service is assigned – Use the Junos OS CLI to specify that the default service is used globally, for a group of interfaces, or for a specific interface.

Junos OS Enhanced Subscriber Management Overview

Junos OS enhanced subscriber management is a next-generation broadband edge software architecture for wireline subscriber management. Enhanced subscriber management enables you to take advantage of increased scaling and performance for configuring and managing dynamic interfaces and services for subscriber management.

Enhanced subscriber management delivers optimized scaling and performance for the existing dynamic subscriber management feature set. Enhanced subscriber management provides feature parity with the legacy Junos OS subscriber management feature set, with certain exceptions. For a list of these feature exceptions, see the latest *Junos OS Release Notes for MX Series 3D Universal Edge Routers* for your Junos OS software.

In order to use dynamic profiles to create and manage dynamic subscriber interfaces and services, you *must* explicitly configure and enable enhanced subscriber management. When enhanced subscriber management is enabled, it handles all subscriber-management control protocol traffic (DHCP, PPP, PPPoE, L2TP, and dynamic VLAN creation) to direct the creation of subscriber sessions and their associated dynamic interfaces.

If you are using only static network configurations and static services in a business edge environment, you do not need to enable enhanced subscriber management to configure these static topologies. When enhanced subscriber management is *not* enabled, the following client applications do not support the use of dynamic profiles, the creation of dynamic interfaces, or dynamic authentication services:

- Dynamic VLANs
- PPPoE
- PPP
- L2TP
- DHCP

From an operational perspective, enhanced subscriber management introduces only minimal changes to existing subscriber management configuration and verification procedures. For example, enhanced subscriber management consolidates several subscriber management components previously distributed across multiple processes into a single process. As a result, enhanced subscriber management can display consolidated information for subscriber management in a single show command.

Address-Assignment Pools Overview

The address-assignment pool feature supports subscriber management and DHCP management functionality by enabling you to create centralized IPv4 and IPv6 address pools independently of the client applications that use the pools. The process manages the pools and the address allocation, whether the addresses come from local pools or from a RADIUS server. For example, multiple client applications, such as DHCP, can use the same address-assignment pool to provide addresses for their particular clients. Client applications can acquire addresses for either authenticated or unauthenticated clients.

Address-assignment pools support both dynamic and static address assignment. In dynamic address assignment, a client is automatically assigned an address from the address-assignment pool. In static address assignment, which is supported for IPv4 pools only, you reserve an address that is then always used by a particular client. Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

You can configure named address ranges within an address-assignment pool. A named range is a subset of the overall address range. A client application can use named ranges to manage address assignment based on client-specific criteria. For example, for IPv4 address-assignment pools, you might create a named range that is based on a specific DHCP option 82 value. Then, when a DHCP client request matches the specified option 82 value, an address from the specified range is assigned to the client.

You can link address-assignment pools together to provide backup pools for address assignment. When no addresses are found to be available in the specified (primary) address pool, the router or switch automatically proceeds to the linked (secondary) address pool to search for an available address to allocate.

The address-assignment pool hold-down feature enables you to specify that no additional address are allocated from an existing active address-assignment pool. This configuration gracefully transforms the active pool to an inactive state as the previously allocated addresses are returned to the pool. When the pool is inactive, you can safely perform maintenance on the pool without affecting any active subscribers.

You can also explicitly identify that an address-assignment pool is used for ND/RA.

DNS Name Server Address Overview

When a client attempts to access a domain – for example, www.example.com – a request is sent to a Domain Name System (DNS) name server. The name server stores information that correlates domain names with IP addresses; the IP address is used to reach the requested domain. In response to the client request, the name server looks up the IP address for the domain – 192.0.2.10 for www.example.com – and returns it to the client.

In your network configuration, you must configure the address of one or more name servers locally on the router or on your RADIUS server. The local configuration supports the following subscriber types:

- DHCPv4 or DHCPv6
- IP over Ethernet (VLAN)
- Terminated PPPoE (IPv4 or IPv6)
- Tunneled PPPoE (IPv4 or IPv6)

You can configure the name server addresses globally (per routing instance), per access profile, or, for DHCP only, per address pool. You can configure more than one name server in a routing instance or access profile by repeating the statement for each address.

Because you can configure name server addresses at more than one level, the address returned to the client is determined by the order of preference among the levels. The preference depends on the client type.

For DHCP subscribers, the preference in descending order is RADIUS > DHCP address pool > access profile > global.

For non-DHCP subscribers, the preference in descending order is RADIUS > access profile > global.

According to the preference order, a name server address configured in RADIUS is preferred by all subscriber types over all other configuration levels. For all subscriber types, the global name server address is used only when no other name server addresses are configured. When a name server address is configured only in a DHCP address pool, then no address is available to non-DHCP subscribers.

When you configure multiple addresses for a name server, the order in which you configure them determines the preference within that configuration. The preference according to configuration level supersedes this ordering.

There is no restriction on the number of DNS name server addresses that you can configure. For DHCP subscribers, all the addresses are sent in DHCP messages. However, only two addresses – determined by preference order – are sent to PPP subscribers.

All changes in these locally configured DNS name servers affect only new subscribers that subsequently log in. Existing subscribers are not affected by the changes.

CLI-Activated Subscriber Services

Subscriber management enables you to use the Junos OS CLI to locally activate and deactivate dynamic subscriber services. CLI-based activation and deactivation provides local control for dynamic subscriber services that is similar to subscriber management's change of authorization (CoA) feature. CoA is considered a remote activation method because the commands, or triggers, are received from a remote server, such as a RADIUS or provisioning server. Both the CoA and CLI-based methods enable you to manage services for subscribers who are currently logged in to the network – you can activate a new service for the subscriber or deactivate a current service.

The CLI-based feature activates the specified service – you cannot use it to modify a subscriber's dynamic profile instantiation or to modify user-defined variables in a dynamic profile. You can, however, include variables that are defined for the service in the dynamic profile.

Subscriber management does not support accounting for CLI-activated subscriber services. Accounting for any service is disabled by default. Therefore when you use the CLI to activate a service, it is activated with accounting disabled, and there is no way to explicitly enable accounting for the service. CLI deactivation of a service previously activated (such as by RADIUS) has no effect on accounting for that service.

CLI-based activation and deactivation is useful in service provider networks that do not use provisioning servers or RADIUS servers to activate and deactivate subscriber services. The local control provided by the CLI-based operations enables service providers to add and remove services for existing subscribers without requiring that the subscriber log out and than log in again to complete the change. For example, a service provider might allow subscribers to log in and initially use the default service, which provides basic features. After the default service is established, the provider might then use CLI-activation to upgrade qualified subscribers to an advanced service, in addition to retaining the initial service. Later, the provider can use CLI-deactivation to terminate the subscriber's advanced service session. The subscriber retains the initial service until the service is deactivated.

CLI-based activation or deactivation of a subscriber service fails if any of the following conditions exist:

- A RADIUS CoA operation or a previous CLI-based activation or deactivation is currently in progress for the subscriber. Only one dynamic request can be active for the subscriber.

- A unified in-service software upgrade (unified ISSU) operation is active.
- The specified service could not be activated or deactivated.
- A CLI-based activation or deactivation of a subscriber service also fails if a PCRF has successfully activated any services for the subscriber. You must override the PCRF provisioning to be able to activate or deactivate services for such a subscriber.

Subscriber Services with Multiple Instances Overview

Services are activated for subscribers either at login, or by using Change of Authorization (CoA) RADIUS messages or command-line interface (CLI) requests. A subscriber can have multiple instances of the same named service, provided that each instance of the subscriber service has a different set of parameters. Support for multiple instances of a subscriber service enables you to use service parameters to customize the same service to meet different needs for a particular subscriber.

In a subscriber access network, each subscriber has its own set of services. You can configure a specific *service instance* for a particular subscriber by specifying a *service name*, also referred to as a *service profile*, and unique service parameters for that service instance. *Service parameters* can include a combination of policy lists, filters, rate-limit profiles, class of service (CoS) profiles, and interface profiles.

For example, filter-service(up-filter,down-filter) and filter-service(upstream-filter,downstream-filter) are considered two different instances of the same service (filter-service) because their parameters, enclosed in parentheses after the service name, are different.

Each service instance is uniquely identified by the combination of its service name and service parameters. In CoA messages, the router identifies a subscriber service by its complete activation string, which consists of the service name and, if configured, one or more service parameters in the order specified.

Diameter Base Protocol Overview

The Diameter protocol is defined in *RFC 3588, Diameter Base Protocol*, and provides an alternative to RADIUS that is more flexible and extensible. The Diameter base protocol provides basic services to one or more applications (also called functions) that each runs in a different Diameter instance. The individual application provides the extended AAA functionality. Applications that use Diameter include Gx-Plus, JSRC, NASREQ, and PTSP.

Diameter peers communicate over a reliable TCP transport layer connection by exchanging Diameter messages that convey status, requests, and acknowledgments by means of standard Diameter AVPs and application-specific AVPs. The Diameter transport layer configuration is based on Diameter network elements (DNEs); multiple DNEs per Diameter instance are supported. Currently only the predefined *master* Diameter instance is supported, but you can configure alternative values for many of the master Diameter instance values.

Each DNE consists of a prioritized list of peers and a set of routes that define how traffic is forwarded. Each route associates a destination with a function, a function partition, and a metric. When an application sends a message to a routed destination, all routes within the Diameter instance are examined for a match. When the best route to the destination has been selected, the message is forwarded by means of the DNE that includes that route.

Multiple routes to the same destination can exist within a given DNE and in different DNEs. In the case of multiple routes that match a request for forwarding, Diameter selects the best route as follows:

- Diameter compares the metric of the routes and selects the route with the lowest metric.
- If multiple routes have the same lowest metric, then Diameter selects the most-qualified route. Diameter evaluates multiple attributes of the route to determine a score that reflects how specifically each route matches the request. By default, the score of a route is 0. Points are added to the score as follows:
- If the destination realm matches the request, add 1.
- If the destination host matches the request, add 2.
- If the function matches the request, add 3.
- If the function partition matches the request, add 4.
- If multiple routes are equally qualified, then Diameter compares the names of the DNEs in lexicographical order and selects the route in the DNE that has the lowest value. For example, dne-austin has a lower value than dne-boston.
- If the routes are tied within the same DNE, then Diameter compares the route names in lexicographical order and selects the route with the lowest value.
- When the state of any DNE changes, the route lookup for all destinations is reevaluated. All outstanding messages to routed destinations are rerouted as needed, or discarded.

To configure a Diameter network element, include the network-element statement at the [edit diameter] hierarchy level, then include the route statement at the [edit diameter network-element *element-name* forwarding] hierarchy level.

To configure a route for the DNE, include the destination (optional), function (optional), and metric statements at the [edit diameter network-element *element-name* forwarding route *dne-route-name*] hierarchy level.

Specify the Diameter peers associated with the DNE by including one or more peer statements at the [edit diameter network-element *element-name*] hierarchy level.

Specify the Diameter peers associated with the DNE by including one or more peer statements at the [edit diameter network-element *element-name*] hierarchy level.

Set the priority for each peer with the priority statement at the [edit diameter network-element *element-name* peer *peer-name*] hierarchy level.

Diameter requires you to configure information about the origin node; this is the endpoint node that originates Diameter for the Diameter instance. Include the host and realm statements at the [edit diameter] hierarchy level to configure the Diameter origin.

You can optionally configure one or more *transports* to specify the source (local) address of the transport layer connection. To configure a Diameter transport, include the transport statement at the [edit diameter] hierarchy level. Then include the address statement at the [edit diameter transport *transport-name*] hierarchy level.

You can optionally specify a logical system and routing instance for the connection by including the logical-system and routing-instance statements at the [edit diameter transport *transport-name*] hierarchy level. By default, Diameter uses the *default* logical system and *master* routing instance. The logical system and routing instance for the transport connection must match that for the peer, or a configuration error is reported.

Each Diameter peer is specified by a name. Peer attributes include address and the destination TCP port used by active connections to this peer. To configure a Diameter peer, include the peer statement at the [edit diameter] hierarchy level, and then include the address and connect-actively statements at the [edit diameter peer *peer-name*] hierarchy level.

To configure the active connection, include the port and transport statements at the [edit diameter peer *peer-name* connect-actively] hierarchy level. The assigned transport identifies the transport layer source address used to establish active connections to the peers.

Understanding CoS for Subscriber Access

Junos class-of-service (CoS) enables you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This functionality allows packet loss to happen according to rules that you configure. The Junos CoS features provide a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient.

In a subscriber access environment, service providers want to provide video, voice, and data services over the same network for subscribers. Subscriber traffic is delivered from the access network, through a router, through a switched Ethernet network, to an Ethernet digital subscriber line access multiplexer (DSLAM). The DSLAM forwards the subscriber's traffic to the residential gateway over a digital subscriber line (DSL). An MX Series router that is installed in a subscriber access network as an edge router can perform subscriber management functions that include subscriber identification and per-subscriber CoS.

In a subscriber access network, a subscriber is an authenticated user – a user that has logged in to the access network at a subscriber interface and then been verified by the configured authentication server and subsequently granted initial CoS services. Subscribers can be identified statically or dynamically. In this network, subscribers are mapped to VLANs, demux, or PPPoE interfaces.

You can configure the router to provide *hierarchical scheduling* or *per-unit scheduling* for subscribers:

- Hierarchical CoS enables you to apply traffic scheduling and queuing parameters (which can include a delay-buffer bandwidth) and packet transmission scheduling parameters (which can include buffer management parameters) to an individual subscriber interface rather than to all interfaces configured on the port.

Hierarchical CoS enables you to dynamically modify queues when subscribers require services.

Per-unit scheduling enables one set of output queues for each logical interface configured under the physical interface. In per-unit scheduling configurations, each Layer 3 scheduler node is allocated a dedicated set of queues.

Because the interface sets corresponding to VLANs using agent-circuit-identifier information are created dynamically, you can apply CoS attributes, such as shaping, at the household level. You must set and define the CoS policy for the agent-circuit-identifier virtual VLAN interface set using the dynamic profile for the agent-circuit-identifier interface set (not the subscriber profile). CoS on dynamic VLANs includes support for Level 3 or Level 2 scheduler nodes for a dynamic interface set. You can also configure a traffic-control profile and a remaining traffic-control profile for a dynamic interface set. CoS on dynamic VLANs enables you to configure a dynamic scheduler map for a traffic-control profile that is used by a dynamic interface set. In this case, the dynamic scheduler map must use the unique ID format.

CoS for Aggregated Ethernet Subscriber Interfaces Overview

You can apply static or dynamic hierarchical CoS to a scheduler node at the aggregated Ethernet logical interface, its underlying physical interface, or an interface set.

When you configure CoS for aggregated Ethernet interfaces, consider the following guidelines:

- Configure the aggregated Ethernet logical interface over two physical interfaces capable of performing hierarchical scheduling.
- For VLAN subscriber interfaces over aggregated Ethernet, you must enable link protection on the aggregated Ethernet interface for hierarchical CoS to operate.
- Link protection is not required for IP or demux subscriber interfaces over aggregated Ethernet. We recommend that you enable targeted distribution on the demux interface to provide accurate hierarchical scheduling for these links.

Keep the following guidelines in mind when configuring interface sets of aggregated Ethernet interfaces:

- Sets of aggregated Ethernet interfaces are supported on MPC/MIC interfaces on MX Series routers only.
- The supported logical interfaces for aggregated Ethernet in an interface set include VLAN demux interfaces, IP demux interfaces, and PPPoE logical interfaces over VLAN demux interfaces.
- The link membership list and scheduler mode of the interface set are inherited from the underlying aggregated Ethernet interface over which the interface set is configured.
- When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.

- If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links and applied to each of the aggregated interface member links.

CoS for PPPoE Subscriber Interfaces Overview

For all supported hardware platforms, you can attach an output traffic-control profile that contains basic shaping and scheduling properties directly to a static or dynamic PPPoE interface. In this type of scenario, you can use each PPPoE interface to represent a household and shape all of the household traffic to an aggregate rate. Each forwarding class is mapped to a queue, and represents one type of services provided to a household customer.

For MPCs that support hierarchical scheduling, you can shape subscriber or access node traffic at different levels of the PPPoE interface hierarchy by attaching traffic-control profiles to interface sets that contain PPPoE members.

MPCs support subscriber interfaces with PPPoE encapsulation over aggregated Ethernet interfaces. These PPPoE subscriber interfaces are configured over VLAN demux interfaces, which are also configured over Aggregated Ethernet interfaces.

You can configure 802.3ad link aggregation group (LAG) stateful port and dense port concentrator (DPC) redundancy. This provides targeted distribution of non-replicated (stacked) PPPoE or IP demux links over VLAN demux links, which in turn are over an aggregated Ethernet (AE) logical interface. Service providers with PPPoE or IP demux interfaces for CoS configurations can provide DPC and port redundancy to subscribers.

NOTE For static PPPoE underlying logical interfaces, use PPPoE interface sets.

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview

Queuing Ethernet Modular Port Concentrators (MPCs) provide a set of dedicated queues for subscriber interfaces configured with hierarchical scheduling or per-unit scheduling.

The dedicated queues offered on these MPCs enable service providers to reduce costs through different scaling configurations. These queuing MPCs enable service providers to reduce the cost per subscriber by allowing many subscriber interfaces to be created with four or eight queues.

This topic describes the overall queue, scheduler node, and logical interface scaling for subscriber interfaces created on these MIC and MPC combinations.

Queue Scaling for MPCs

Beginning with Junos OS Release 15.1, MPC2E-3D-NG-Q, MPC3E-3D-NG-Q, MPC5EQ-40G10G, and MPC5EQ-100G10G MPCs support up to five levels of heirarchical queuing. Table 6 lists the number of dedicated queues and nodes supported per MPC.

MPC	Dedicated Queues	Level 4 Nodes	Level 3 Nodes	Level 2 Nodes	Level 1 Nodes (Ports)
MPC2E-3D-NG-Q MPC3E-3D-NG-Q	512,000	64,000	16,000	4000	384
MPC5EQ-40G10G MPC5EQ-100G10G	1 million	128,000	32,000	4000	384

Table 6 Dedicated Queues for MPCs

CAUTION The maximum scaling targets provided in Table 6 are based on system level design specifications. Actual realized subscriber or session scale is highly dependent upon the configuration and can be influenced by configuration variables including: the number of routes, the number of enabled services, the number of policy and firewall filters, policers, counters, statistics and access model type. Once you define a configuration, your Juniper account team can help characterize the expected system level scale or scale range for your live deployment.

MPCs vary in the number of Packet Forwarding Engines on board. MPC2E-3D-NG-Q and MPC3E-3D-NG-Q MPCs each have one Packet Forwarding Engine, allowing all 64,000 Level 4 (subscriber) nodes to be allocated to a single MIC. MPC5EQ MPCs have two Packet Forwarding Engines, one for each possible MIC, each supporting 64,000 Level 4 (subscriber) nodes.

NOTE The nonqueuing MPCs MPC2E-3D-NG, MPC3E-3D-NG, MPC5E-40G10G, and MPC5E-100G10G provide up to eight queues per port in standard configuration. However, each of these MPCs can be configured to provide limited-scale hierarchical class of service (HCoS) and up to 32,000 queues.

Managing Remaining Queues

When the number of available dedicated queues on the MPC drops below 10 percent, an SNMP trap is generated to notify you.

When the maximum number of dedicated queues on the MPCs is reached, a system log message, COSD_OUT_OF_DEDICATED_QUEUES, is generated. The system does not provide subsequent subscriber interfaces with a dedicated set of queues. For per-unit scheduling configurations, there are no configurable queues remaining on the MPC.

For hierarchical scheduling configurations, remaining queues are available when the maximum number of dedicated queues is reached on the MPC. Traffic from these logical interfaces is considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces. These common queues are the default port queues that are created for every port. You can configure a traffic-control profile and attach that to the interface to provide CoS parameters for the remaining queues.

These subscriber interfaces remain with this traffic-control profile, even if dedicated queues become available.

Bandwidth Management for Downstream Traffic in Edge Networks Overview

In a subscriber access network, traffic with different encapsulations can be passed downstream to other customer premise equipment (CPE) through the MX Series router. Managing the bandwidth of downstream ATM traffic to Ethernet interfaces can be especially difficult because of the different Layer 2 encapsulations.

The downstream network is not necessarily the directly attached network. In typical broadband network gateway (BNG) configurations, the directly attached network is an Ethernet access network, which provides access to either another frame-based network, or a cell-based network.

The *overhead accounting* feature enables you to shape traffic based on whether the downstream network is a frame-based network, like Ethernet, or a cell-based network, like ATM. It assigns a byte adjustment value to account for different encapsulations.

This feature is available on MIC and MPC interfaces.

Effective Shaping Rate

The shaping-rate, also known as peak information rate (PIR), is the maximum rate for a scheduler node or queue.

The true rate of a subscriber at the access-loop/CPE is a function of:

- The shaping-rate in effect for the subscriber's household, in bits per second.
- Whether the subscriber is connected to a frame-based or cell-based network.
- Number of bytes in each frame that are accounted for by the shaper.

NOTE Chassis egress-shaping-overhead is not included in the effective rate. Egress-shaping-overhead accounts for the physical interface overhead (ISO OSI Layer 1). Effective shaping-rate is a Layer 2 (ISO OSI) rate.

Shaping Modes

There are two modes used for adjusting downstream traffic:

- *Frame shaping mode* is useful for adjusting downstream traffic with different encapsulations. Shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead. Frame is the default shaping mode on the router.
- *Cell shaping mode* is useful for adjusting downstream cell-based traffic. In cell shaping mode, shaping is based on the number of bytes in cells, and accounts for the cell encapsulation and padding overhead.

When you specify cell mode, the resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.

To account for ATM segmentation, the router adjusts all of the rates by 48/53 to account for 5-byte ATM AAL5 encapsulation. In addition, the router accounts for cell padding, and internally adjusts each frame by 8 bytes to account for the ATM trailer.

Byte Adjustments

When the downstream traffic has different byte sizes per encapsulation, it is useful to configure a *byte adjustment* value to adjust the number of bytes per packet to be included in or excluded from the shaping mechanism. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment. For example, to properly account for a 4-byte header stripped by the downstream network, set the overhead-accounting bytes to -4. To properly account for a 12-byte header added by the downstream network, set the overhead-accounting bytes to 12.

We recommend that you specify a byte adjustment value that represents the difference between the CPE protocol overhead and B-RAS protocol overhead.

The system rounds up the byte adjustment value to the nearest multiple of 4. For example, a value of 6 is rounded to 8, and a value of –10 is rounded to –8.

You do not need to configure a byte adjustment value to account for the downstream ATM network. However, you can specify the byte value to account for additional encapsulations or decapsulations in the downstream network.

Relationship with Other CoS Features

Enabling the overhead accounting feature affects the resulting shaping rates, guaranteed rate, and excess rate parameters, if they are configured.

The overhead accounting feature also affects the egress shaping overhead feature that you can configure at the chassis level. We recommend that you use the egress shaping-overhead feature to account for the Layer 2 overhead of the outgoing interface, and use the overhead-accounting feature to account for downstream traffic with different encapsulations and cell-based networks.

When both features are configured, the total byte adjustment value is equal to the adjusted value of the overhead-accounting feature plus the value of the egress-shaping-overhead feature. For example, if the configured byte adjustment value is 40, and the router internally adjusts the size of each frame by 8, the adjusted overhead accounting value is 48. That value is added to the egress shaping overhead of 24 for a total byte adjustment value of 72.

Changing CoS Services Overview

This topic describes how to provide CoS when subscribers dynamically upgrade or downgrade services in an access environment.

You can configure your network with an *access profile* that provides all subscribers with default CoS parameters when they log in. For example, all subscribers can receive a basic data service. By configuring the access profile with Junos OS predefined variables for RADIUS-provided CoS parameters, you also enable the service to be activated for those subscribers at login.

To enable subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages after login, configure a *service profile* that includes user-defined variables.

Types of CoS Variables Used in a Service Profile

You can configure variables for the following CoS parameters in a service profile:

- Shaping rate
- Delay buffer rate
- Guaranteed rate
- Scheduler map

For each CoS parameter, you must associate a RADIUS vendor ID. For each vendor ID, you must assign an attribute number and a tag. The tag is used to differentiate between values for different CoS variables when you specify the same attribute number for those variables. These values are matched with the values supplied by RADIUS during subscriber authentication. All of the values in the dynamic profile must be defined in RADIUS or none of the values are passed.

Optionally, you can configure default values for each parameter. Configuring default values is beneficial if you do not configure RADIUS to enable service changes. During service changes, RADIUS takes precedence over the default value that is configured.

Static and Dynamic CoS Configurations

Depending on how you configure CoS parameters in the access and service profiles, certain CoS parameters are replaced or merged when subscribers change or activate new services.

Static configuration is when you configure the scheduler map and schedulers in the static [edit class-of-service] hierarchy and reference the scheduler map in the dynamic profile. Dynamic configuration is when you configure the scheduler map and schedulers within the dynamic profile.

The CoS configuration also depends on whether you have enabled multiple subscribers on the same logical interface using the aggregate-clients statements in the dynamic profile referenced by DHCP. When you specify the aggregate-clients replace statement, the scheduler map names are replaced. In both cases, if the length of the scheduler map name exceeds 128 characters, subscribers cannot log in. When you specify the aggregate-clients merge statement, the scheduler map names specified in the dynamic profile are appended.

NOTE To improve CoS performance in IPv4, IPv6, and dual-stack networks, we recommend that you use the aggregate-clients replace statement rather than the aggregate-clients merge statement.

CoS for Interface Sets of Subscribers Overview

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

Guidelines for Configuring Dynamic Interface Sets in a Subscriber Access Network

Interface sets enable service providers to group logical interfaces so they can apply CoS parameters to all of the traffic in the group.

Interface sets are beneficial for various scenarios in a subscriber access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

When configuring interface sets for subscriber access, keep the following guidelines in mind:

- You can configure interface sets of VLAN demux, PPPoE, or demux interfaces over aggregated Ethernet interfaces.
- An interface can only belong to one interface set. If you try to add the same interface to different interface sets, the commit operation fails.
- You configure the interface set and the traffic scheduling and shaping parameters in a dynamic profile. However, you must apply the traffic-control profile to the interface set in the static [edit class-of-service] hierarchy.

NOTE This rule applies to all interface sets except ACI sets.

The $junos-interface-set-name predefined variable is available only for RADIUS Accept messages; change of authorization (CoA) requests are not supported.

The $junos-svlan-interface-set-name predefined variable locally generates an 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*. For example, an aggregated Ethernet interface “ae0,” with a dual-tagged VLAN interface that has an outer tag of “111,” results in a $junos-svlan-interface-set-name dynamic variable of “ae0-111”. Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of “111,” results in a $junos-svlan-interface-set-name dynamic variable of “ge-1/1/0-111”.

The $junos-phy-ifd-interface-set-name predefined variable locally generates an 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.

The $junos-tagged-vlan-interface-set-name predefined variable locally generates an 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 as follows:

- Dual-tagged (client) VLAN – *physical_interface_name* - *outer_VLAN_tag* - *inner_VLAN_tag*. For example, an aggregated Ethernet interface "ae0," with a dual-tagged VLAN interface that has an outer tag of "111" and an inner tag of "200," results in a $junos-tagged-vlan-interface-set-name dynamic variable of "ae0-200-111". Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same dual-tagged VLAN interface that has an outer tag of "111" and an inner tag of "200," results in a $junos-tagged-vlan-interface-set-name dynamic variable of "ge-1/1/0-200-111".
- Single tagged (service) VLAN – *physical_interface_name* - *VLAN_tag*. For example, an aggregated Ethernet interface "ae0," with an N:1 VLAN using the single tag of "200," results in a $junos-tagged-vlan-interface-set-name dynamic variable of "ae0-200". Similarly, a non-aggregated Ethernet interface of ge-1/1/0, with the same N:1 VLAN using the single tag of "200," results in a $junos-tagged-vlan-interface-set-name dynamic variable of "ge-1/1/0-200".

All dynamic demux, dual-tagged VLAN logical interfaces with the same outer VLAN tag and physical interface are assigned to the same interface set and all CoS values provisioned with the dynamic profile are applied to the interfaces that are part of the set.

The interface set name must be explicitly referenced in the CoS configuration as part of the static configuration outside of the dynamic profile. The CoS configuration is static and the interface set name must be statically referenced.

NOTE This rule applies to all interface sets except ACI sets.

RADIUS can return an *access-accept* message under certain conditions. A configured RADIUS VSA for the interface set name takes precedence over the locally generated variable on the router. This means that if the interface-set-name VSA is configured on RADIUS, the router continues to use this variable instead of the locally generated value from the dynamic variable.

Sets of aggregated Ethernet interfaces are supported on MPC/MIC interfaces on MX Series routers only.

The supported interface stacks for aggregated Ethernet in an interface set include VLAN demux interfaces, IP demux interfaces, and PPPoE logical interfaces over VLAN demux interfaces.

The link membership list and scheduler mode of the interface set are inherited from the underlying aggregated Ethernet interface over which the interface set is configured.

When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.

If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links and applied to each of the aggregated interface member links.

Understanding Firewall Filters

Understanding Dynamic Firewall Filters

Firewall filters provide rules that define whether to accept or reject packets that are transiting an interface on a router. The subscriber management feature supports four categories of firewall filters:

Classic filters are static filters that are applied to an interface dynamically. They are compiled at commit time and then, when a service is activated, an interface-specific filter is created and attached to a logical interface. This dynamic application is performed by associating input or output filters with a dynamic profile. When triggered, a dynamic profile applies the filter to an interface. Because classic filters are static, they cannot contain subscriber-specific terms (also called rules).

Parameterized filters allow you to implement customized filters for each subscriber session. In parameterized filters, you use variables to define a filter. When services are activated for a subscriber, actual values such as policing rates, destination addresses, or ports are substituted for the variables and are used to create filters.

Ascend-Data-Filters allow you to create dynamic filters based on values received from the RADIUS server in the Ascend-Data-Filter attribute (RADIUS attribute 242). The filter is configured on the RADIUS server and contains rules that specifically match conditions for traffic and define an action for the router to perform. When services are activated for a subscriber, a filter is created based on the values in the RADIUS attribute. You can also use Ascend-Data-Filters to create static filters by configuring the Ascend-Data-Filter attribute in a dynamic profile.

Fast update filters are similar to classic filters. However, fast update filters support subscriber-specific, rather than interface-specific, filter values. Fast update filters also allow individual filter terms to be incrementally added or removed from filters without requiring that the entire filter be recompiled for each modification. Fast update filters are essential for networking environments in which multiple subscribers share the same logical interface.

You configure firewall filters to determine whether to accept or reject traffic before it enters or exits an interface to which the firewall filter is applied. An *input* (or *ingress*) firewall filter is applied to packets that are entering a network. An *output* (or *egress*) firewall filter is applied to packets that are exiting a network. You can configure firewall filters to subject packets to filtering or class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority).

Classic Filters Overview

The dynamic firewall feature supports classic filters, which are static filters that are applied to an interface dynamically. They are compiled at commit time and then, when a service is activated, an interface-specific clone of the filter is created and attached to a logical interface. This dynamic application is performed by associating input or output filters with a dynamic profile.

Classic Filter Types

The following classic filter types are supported:

- Port (Layer 2) firewall filter – Port firewall filters apply to Layer 2 switch ports. You can apply port firewall filters only in the ingress direction on a physical port.
- VLAN firewall filter – VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, and leave a VLAN. You can apply VLAN firewall filters in both ingress and egress directions on a VLAN. VLAN firewall filters are applied to all packets that are forwarded to or forwarded from the VLAN.
- Router (Layer 3) firewall filter – You can apply a router firewall filter in both ingress and egress directions on Layer 3 (routed) interfaces.

Classic Filter Components

When creating a classic filter, you first define the family address type (inet or inet6) and then you define one or more terms that specify the filtering criteria and the action to take when a match occurs.

Each term, or rule, consists of the following components:

- Match conditions – Specifies values or fields that the packet must contain. You can define various match conditions, including:
- IP source address field
- IP destination address field
- Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field
- IP protocol field
- Internet Control Message Protocol (ICMP) packet type
- TCP flags
- Interfaces
- Actions – Specifies what to do when a match condition occurs. Possible actions are to accept or discard a packet. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.

Classic Filter Processing

The order of the terms within a classic filter is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. When a firewall filter contains multiple terms, the router takes a top-down approach and compares a packet against the first term in the firewall filter. If the packet matches the first term, the router executes the action defined by that term to either accept or reject the packet, and no other terms are evaluated. If the router does not find a match between the packet and first term, it then compares the packet to the next term in the firewall filter by using the same match process. If no match occurs between the packet and the second term, the router continues to compare the packet to each successive term defined in the firewall filter until a match is found. If a packet does not match any terms in a firewall filter, the default action is to discard the packet.

You can also specify a precedence (from 0 through 255) for input and output filters within a dynamic profile to force filter processing in a particular order. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. Filters with lower precedence values are applied to interfaces before filters with higher precedence values. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in random order.

NOTE Dynamic filters do not process outbound packets that are sourced from the routing engine. To filter outbound packets that are sourced from the routing engine, you can create static outbound filters for each interface.

Guidelines for Creating and Applying Classic Filters for Subscriber Interfaces

Dynamic configuration of firewall filters is supported. However, you can also continue to create static firewall filters for interfaces as you do normally, and then dynamically apply those filters to statically created interfaces using dynamic profiles. You can also use dynamic profiles to attach input and output filters through RADIUS.

When creating and applying filters, keep the following in mind:

- Dynamic application of only input and output filters is supported.
- The filters must be interface-specific.
- You can create family-specific inet and inet6 filters.
- You can create interface-specific filters at the unit level that apply to any family type (inet or inet6) configured on the interface.
- You can add or remove both IPv4 and IPv6 filters with the same service activation or deactivation.
- You can remove one filter type without impacting the other type of filter. For example, you can remove IPv6 filters and leave the current IPv4 filters active.
- You can chain up to five input filters and four output filters together.

- If you do not configure and apply a filter, the interface uses the default group filter configuration.
- You cannot modify or delete a firewall filter while subscribers on the same logical interface are bound.

Parameterized Filters Overview

Parameterized filters allow you to implement customized filters for each subscriber session. In parameterized filters, you use variables called unique identifiers (UIDs) to define your filter. When services are activated for a subscriber, actual values are substituted for the variables and are used to create filters.

Parameterized filters are configured under a dynamic profile. You can configure a general baseline filter under a dynamic profile and then provide specific variables of that filter when a dynamic session is activated. These variables can include policing rates, destination addresses, ports, and other items.

To provide better scaling, the system analyzes a dynamic profile, and then determines whether the set of variables for one session is the same as for a previous session. If a matching filter already exists, the session creates an interface-specific filter copy of that filter template. If the filter does not already exist, the session reads the configuration and compiles a new filter. This filter is installed as a template with an interface-specific filter copy for the current session pointing to it.

Ascend-Data-Filter Policies for Subscriber Management Overview

Subscriber management enables you to use Ascend-Data-Filters to create policies for subscriber traffic. An Ascend-Data-Filter is a binary value that is configured on the RADIUS server. The filter contains rules that specify match conditions for traffic and an action for the router to perform (such as accept or discard the traffic). The match conditions might include the source and destination IP address or port, the protocol, the filter direction, the traffic class, and policer information.

Subscriber management uses a dynamic profile to obtain the Ascend-Data-Filter attribute (RADIUS attribute 242) from the RADIUS server and apply the policy to a subscriber session. Dynamic profiles support Ascend-Data-Filters for inet and inet6 family types, and both families can be present in a dynamic profile. You include Junos OS predefined variables in the dynamic profiles – $junos-adf-rule-v4 for family inet and $junos-adf-rule-v6 for inet6. The Ascend-Data-Filter attribute can include rules for both address families. The predefined variables map the Ascend-Data-Filter rules for the respective family to the Junos OS firewall filter process. A firewall filter is created and attached to the subscriber’s logical interface.

You can also configure a static Ascend-Data-Filter by manually entering the required binary data as a hexadecimal string in a dynamic profile. A statically configured Ascend-Data-Filter in a dynamic profile takes precedence over an Ascend-Data-Filter attribute that is received from RADIUS. The static method is time-consuming to configure; it is typically used only for testing purposes.

The Ascend-Data-Filter attribute is supported in RADIUS Access-Accept and Change of Authorization (CoA) messages.

CoA updates existing filters based on the Ascend-Data-Filter Type field, as shown in the following list:

- If the Type field is 1, IPv4 rules are updated and IPv6 rules are unchanged. The opposite is true if the Type field is 3.
- If both Type 1 and 3 are specified, then all rules are updated.
- If the CoA has no Ascend-Data-Filter rules, then the existing rules are unchanged.

Filter Naming Conventions

Each Ascend-Data-Filter has a unique name, which is assigned by the dynamic firewall process, dfwd. The assigned names are displayed in the results of the `show subscriber extensive` and `show firewall` commands. Ascend-Data-Filters use the following naming convention:

```
__junos_adf_session#-interfacename-family-direction
```

For example:

```
__junos_adf_33847-ge/1/0/4.53-init-in
```

Each Ascend-Data-Filter rule maps to a single term, and the term names are simply t0, t1, ..., t*n*. If you configure the `counter` option, the router adds a count action to each term that is created. The counter names are a combination of the term names with -cnt appended. For example t0-cnt and t1-cnt.

Use of Multiple Sessions with Ascend-Data-Filters on an Interface

An interface can have multiple subscriber sessions, each session using its own Ascend-Data-Filter rules. When an Ascend-Data-Filter is applied to a subscriber session, the rules are created independently of any other filters and are added to the interface filter list. The Ascend-Data-Filter rules for the other sessions on the same interface are also added to the filter list. All packets that are processed for the interface must go through all filters, and the filters are applied according to the precedence you set.

Because the filter list can be a combination of several rules, you must consider how the multiple filters coexist. You must ensure that the filters are designed and applied correctly in order to provide the desired filtering and resulting action. For example, a session might have a filter that accepts traffic from Subscriber-A and discards all other traffic. However, a second session on the same interface might have a filter that accepts traffic from Subscriber-B only and discards other traffic. When the two filters are combined in the filter list, traffic from Subscriber-B is discarded by the first filter, and traffic from Subscriber-A is discarded by the second filter. As a result, no traffic is accepted on the interface because the two filters essentially cancel out each other and discard all traffic.

Optional ADF Filter Requirement for Some Subscribers

When you include either of the predefined variables – $junos-adf-rule-v4 or $junos-adf-rule-v6 – in the dynamic profile, by default the RADIUS reply message must include the Ascend-Data-Filter attribute (RADIUS attribute 242) for each subscriber. If the attribute is not included, the router reports an error.

A service provider might apply the same dynamic profile to a mixed pool of subscribers, such that the attribute is included by RADIUS for some of the subscribers and is not included for others. By default, the router returns an error for each of the subscribers without the attribute, consuming system resources. You can configure the dynamic profile to accommodate such a mixture of subscribers by making the attribute requirement optional. To do so, and to suppress attribute error reporting, specify the not-mandatory option with the adf statement at the [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter] hierarchy level. With this configuration, the Ascend-Data-filter is simply not created when the Ascend-Data-Filter attribute is not present.

Fast Update Filters Overview

Fast update filters provide more efficient filter processing over classic static filters when dynamic services are implemented for multiple subscribers that share the same logical interface.

Fast update filters support subscriber-specific filter values, as opposed to classic filters, which are interface-specific. Fast update filters allow individual filter terms, or rules, to be added or removed from filters without requiring the router to recompile the filter after each modification – terms are added and removed when subscriber services are added and removed.

Using the fast update filters feature involves three distinct operations:

- Creating the filter – You define fast update filters under the [edit dynamic-profiles *profile-name* firewall family *family*] hierarchy. The dynamic-profiles stanza enables you to use dynamic variables to create subscriber-specific configurations for the filter's match terms.
- Associating the filter with a dynamic profile – You use the [edit dynamic-profiles *profile-name* interface *interface-name* unit *unit-number* family *family* hierarchy to associate the filter with a dynamic profile. This is the same procedure used for classic filters.
- Attaching the filter to an interface – When a subscriber logs in, the dynamic profile instantiates the subscriber session and applies the properties of the profile, including the fast update filter, to the session interface. This is the same procedure used for classic filters. Also, similar to classic filters, the name of fast update filters can be provided in a user's RADIUS file.

When a dynamic profile instantiates a subscriber session and applies a fast update filter, the router verifies that the filter is not already present on the session interface. If the filter is not present, the router adds the filter. If the filter is already present on the interface, the router simply adds any new terms that are not in the existing filter. This procedure is reversed when subscriber sessions are deleted. Any terms that were added by a session are then removed when the session is deleted. The filter is deleted when the last subscriber session is deleted.

NOTE You can optionally specify that a term can be added only once and cannot be modified.

Fast Update Filter Components

When creating a fast update filter, you define one or more terms that specify the filtering criteria and the action to take when a match occurs.

Each term consists of the following components:

- Match condition – Specifies values or fields that the packet must contain. You can match a maximum of five fields in a fast update filter. A match condition can contain a single value or range. This differs from classic filters, in which terms can have multiple values. However, you can use additional terms to specify multiple ranges. The order in which the terms appear in a fast update filter is not important, because the router examines the most specific term first. (Classic filters examine the terms in the order in which the terms are listed.)
- Action – Specifies what to do when a packet matches the match condition. If no action is specified for a term, the default action is to accept the packet.

Terms that are added to the filter during session instantiation must have a unique set of match conditions. Two terms overlap, or conflict, if a packet can match both sets of conditions – as a result, there are two different actions for the packet. You can ensure that terms are unique by using the $junos-subscriber-ip-address variable as the source-address (for an input filter) or destination-address (for an output filter) in the from statement. You must then supply the source-address or destination-address condition, as appropriate, as the first condition in the match-order statement.

Fast Update Filter Processing

You must use the match-order statement to explicitly specify the order in which the router examines filter match conditions. Also, the router examines only those conditions that you include in the match-order statement. When a fast update filter contains multiple terms, the router compares a packet against the terms starting with the most specific condition first. When the packet first matches a condition, the router performs the action defined in the term to either accept or reject the packet, and then no other terms are evaluated. If the router does not find a match between the packet and first term, it then compares the packet to the next term in the filter. The router continues to compare the packet to the next specified term until a match is found. If there is no match after all terms have been examined, the router silently drops the packet.

You can specify a precedence (from 0 through 255) for input and output filters within a dynamic profile to force filter processing in a particular order. Setting a lower precedence value for a filter gives it a higher precedence within the dynamic profile. Filters with lower precedence values are applied to interfaces before filters with higher precedence values. A precedence of zero (the default) gives the filter the highest precedence. If no precedence is specified, the filter receives a precedence of zero (highest precedence). Filters with matching precedence (zero or otherwise) are applied in random order.

Fast Update Filter Names

When a filter is attached to an interface, the router first searches for a classic filter with the specified name, and then uses the classic filter. If no classic filter exists with that name, the router then searches in the dynamic profile for a fast update filter with the specified name, and uses that filter.

If two different dynamic profiles include a fast update filter with the same name, the match-order specification of the two filters must be identical. If the two filters are activated on the same interface, the terms are added together.

The router includes the filter name in show firewall command results. The router also creates unique names for filter terms and counters for the `show firewall` command.

When a fast update filter is created by the activation of a dynamic profile, the router creates an interface-specific name for the filter. The name uses the following format, which is also used for classic filters:

<filter-name>-<interface-name>.<subunit>-<direction>

For example, an input filter named httpFilter on interface ge-1/0/0.5 is named as follows (in indicates an input filter and out indicates an output filter):

http-filter-ge-1/0/0.5–in

The router creates unique names for the filter terms and counters by appending the session ID to all term and counter names. Terms that use the only-at-create statement have a session-id of 0. Terms and counters use the following format:

<term-name>-<session-id>

<counter-name>-<session-id>

Guidelines for Creating and Applying Fast Update Filters

Fast update filters enable you to create subscriber-specific firewall filters and dynamically apply these filters to statically created interfaces using dynamic profiles. Individual terms can be added to, or removed from, a filter without requiring that the entire filter be recompiled.

When creating and applying fast update filters, keep the following in mind:

- Dynamic application of input and output filters is supported.
- You cannot use the same fast update filter as both an input and output filter in the same dynamic profile attached to an interface.
- Fast update filters must always include terms that permit DHCP traffic to pass.
- You can create family inet and inet6 filters.
- You can add or remove both IPv4 and IPv6 filters with the same service activation or deactivation.
- You can remove one filter type without impacting the other type of filter. For example, you can remove IPv6 filters and leave the current IPv4 filters active.
- The interface-specific statement is required for all fast update filters.
- The match-order statement is required – you must explicitly state the order of the match fields in a fast update filter.
- The match-order statement uses an implied wildcard for conditions that you specify in the statement. If you specify a condition that is not also configured in the from specification of a filter term, the router considers that a wildcard for that condition.

- A filter term can have only a single value or range; however, you can configure multiple terms to specify multiple ranges.
- You can match a maximum of five match conditions in a filter.

Unicast RPF in Dynamic Profiles for Subscriber Interfaces

Unicast reverse-path forwarding (RPF) provides a way to reduce the effect of denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on IPv4 and IPv6 interfaces. When you configure unicast RPF on an interface, it checks the packet source address. Packets that pass the check are forwarded. Packets that fail the check are dropped, or if a fail filter is configured, are passed to the filter for further evaluation.

Unicast RPF has two behavioral modes, strict and loose. When you configure unicast RPF in a dynamic profile, strict mode is the default. In strict mode, unicast RPF checks whether the source address of the incoming packet matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. In loose mode, unicast RPF checks only whether the source address has a match in the routing table. It does not check whether the interface expects to receive a packet from a specific source address.

For both modes, when an incoming packet fails the unicast RPF check, the packet is not accepted on the interface. Instead, unicast RPF counts the packet and sends it to an optional fail filter, if present. The fail filter determines what further action is taken on the packet. In the absence of a fail filter, the packet is silently discarded.

Firewall Filters and Enhanced Network Services Mode Overview

Under normal conditions, every firewall filter is generated in two different formats – compiled and term-based. The compiled format is used by the routing engine (RE) kernel, FPCs, and MS-DPs. The term-based format is used by MPCs. Compiled firewall filters are duplicated for each interface or logical interface to which they are applied. Term-based filters, instead of being duplicated, are referenced by each interface or logical interface.

When a combination of MPCs and any other cards populate a chassis, the creation of both firewall filter file formats is necessary. In most networks, the creation of both filter formats and any amount of duplication for compiled firewall filters has no effect on the router. However, in subscriber management networks that include thousands of statically configured subscriber interfaces, creating filters in multiple formats and duplicating those filters for each interface can utilize a large portion of router memory resources. You can use either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode to improve the scaling and performance specific to routing filters in a subscriber access network that uses statically configured subscriber interfaces.

In configurations where interfaces are created either statically or dynamically and firewall filters are applied dynamically, you must configure the chassis network services to run in enhanced mode. In configurations where interfaces are created statically and firewall filters are applied statically, you must configure chassis network services to run in enhanced mode and also configure each firewall filter for enhanced mode.

NOTE Do not use enhanced mode for firewall filters that are intended for control plane traffic. Control plane filtering is handled by the Routing Engine kernel, which cannot use the term-based format of the enhanced mode filters.

Table 7 shows the configuration options when determining enhanced network services mode usage.

Interface and Filter Configuration	Chassis Enhanced Mode Required	Firewall Filter Enhanced Mode Required
Dynamically-created interfaces and dynamically-applied filters	Yes	No
Statically-created interfaces and dynamically-applied filters	Yes	No
Statically-created interfaces and statically-applied filters	Yes	Yes

Table 7 Enhanced Network Services Mode and Firewall Filter Use Case Determination

To achieve significant resource savings for the router, combine chassis and filter enhanced mode configuration as follows:

- Install only MPCs in the chassis.

NOTE Configuring chassis network services to run one of the enhanced network services modes results in the router enabling only MPCs and MS-DPCs. Because MS-DPCs use compiled firewall filter format, a router chassis that is configured for one of the enhanced network services modes, configuring standard (non-enhanced) firewall filters for use with any MS-DPCs can decrease optimal resource efficiency.

- When configuring static interfaces on the router, configure chassis network services to run either Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode.
- When statically applying firewall filters to statically-created interfaces, configure any firewall filters for enhanced mode to limit the filter creation to only term-based format.

NOTE Any firewall filters that are not configured for enhanced mode are created in both compiled and term-based format, even if the chassis is running one of the enhanced network services modes. Only term-based (enhanced) firewall filters will be generated, regardless of the setting of the enhanced-mode statement at the [edit chassis network-services] hierarchy level, if any of the following are true:

- Flexible filter match conditions are configured at the [edit firewall family *family-name* filter *filter-name* term *term-name* from] or [edit firewall filter *filter-name* term *term-name* from] hierarchy levels.

- A tunnel header push or pop action, such as GRE encapsulate or de-encapsulate is configured at the [edit firewall family *family-name* filter *filter-name* term *term-name* then] hierarchy level.
- Payload-protocol match conditions are configured at the [edit firewall family *family-name* filter *filter-name* term *term-name* from] or [edit firewall filter *filter-name* term *term-name* from] hierarchy levels.
- An extension-header match is configured at the [edit firewall family *family-name* filter *filter-name* term *term-name* from] or [edit firewall filter *filter-name* term *term-name* from] hierarchy levels.
- A match condition is configured that only works with MPC cards, such as firewall bridge filters for IPv6 traffic.

WARNING Any firewall filter meeting the previous criteria will not be applied to the loopback, lo0, interface of DPC based FPCs. This means that term-based (enhanced) filters configured for use on the loopback interface of a DPC based FPC will not be applied. This will leave the Routing Engine unprotected by that filter.

Dynamic Service Sets Overview

A service set is a collection of services to be performed by an Adaptive Services (AS) or Multiservices PIC. You configure a service-set definition at the [edit services] hierarchy level. You can then apply the service set to one or more interfaces on the router. The service set can be applied either dynamically or statically.

To dynamically associate a service set to interfaces you include the service-set statement with the input or output statement at the [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* service] hierarchy level.

To statically associate a defined service set with an interface, you include the service-set statement with the input or output statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *family* service] hierarchy level.

Methods for Regulating Traffic by Applying Hierarchical Policers

You can deploy policers to enforce service level agreements limiting the input rate at the edge, and at the boundary between domains, to guarantee an equitable deployment of the service among the different domains. Policers determine whether each packet conforms (falls within the traffic contract), exceeds (using up the excess burst capacity), or violates (totally out of the traffic contract rate) the configured traffic policies, and then sets the prescribed action.

Hierarchical policers rate-limit premium traffic separately from the aggregate traffic on an interface as determined by different configured rates. You can 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 traffic or packets are classified for expedited forwarding (EF) or for a lower priority, such as non-expedited forwarding (non-EF).

Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine. You can apply a hierarchical policer for premium and aggregate (premium plus normal) traffic levels to a logical interface.

Hierarchical policing uses two token buckets, one for premium (EF) traffic and one for aggregate (non-EF) traffic, as shown in Figure 8.

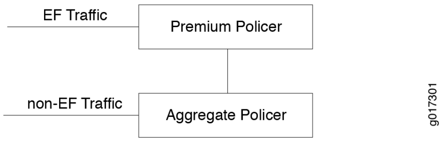

Figure 8 Hierarchical Policer

The class-of-service (CoS) configuration determines which traffic is EF and which is non-EF. Logically, hierarchical policing is achieved by chaining two policers.

Premium policer – You configure the premium policer with traffic limits for high-priority EF traffic only: a guaranteed bandwidth and a corresponding burst-size limit. EF traffic is categorized as nonconforming when its average arrival rate exceeds the guaranteed bandwidth and its average packet size exceeds the premium burst-size limit. For a premium policer, the only configurable action for nonconforming traffic is to discard the packets.

Aggregate policer – You configure the aggregate policer (also known as a logical interface policer) with an aggregate bandwidth (to accommodate both high-priority EF traffic up to the guaranteed bandwidth and normal-priority non-EF traffic) and a burst-size limit for non-EF traffic only. Non-EF traffic is categorized as nonconforming when its average arrival rate exceeds the amount of aggregate bandwidth not currently consumed by EF traffic and its average packet size exceeds the burst-size limit defined in the aggregate policer. For an aggregate policer, the configurable actions for nonconforming traffic are to discard the packets, assign a forwarding class, or assign a packet loss priority (PLP) level.

NOTE You must configure the bandwidth limit of the premium policer at or below the bandwidth limit of the aggregate policer. If the two bandwidth limits are equal, then only non-EF traffic passes through the interface unrestricted; no EF traffic arrives at the interface.

Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer. EF traffic is guaranteed the bandwidth specified as the premium bandwidth limit, while non-EF traffic is rate-limited to the amount of aggregate bandwidth not currently consumed by the EF traffic. Non-EF traffic is rate-limited to the entire aggregate bandwidth only while no EF traffic is present.

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. In Figure 8, the premium policer policies EF traffic and the aggregate policer polices non-EF traffic. In the sample configuration that follows, the hierarchical policer is configured with the following components:

Premium policer has a bandwidth limit set to 2 Mbps, burst-size limit set to 50 KB, and nonconforming action set to discard packets.

Aggregate policer has a bandwidth limit set to 10 Mbps, burst-size limit set to 100 KB, and nonconforming action set to mark high PLP.

```
[edit]
user@host# show dynamic-profiles firewall
hierarchical-policer policer-agg-prem {
    aggregate {
        if-exceeding {
            bandwidth-limit 10m;
            burst-size-limit 100k;
        }
        then {
            loss-priority high;
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit 2m;
            burst-size-limit 50k;
        }
        then {
            discard;
        }
    }
}
```

EF traffic is guaranteed a bandwidth of 2 Mbps. Bursts of EF traffic – EF traffic that arrives at the interface at rates above 2 Mbps – can also pass through the interface, provided that sufficient tokens are available in the 50 KB burst bucket. When no tokens are available, EF traffic is rate-limited using the discarded action associated with the premium policer.

Non-EF traffic is metered to a bandwidth limit that ranges between 8 Mbps and 10 Mbps, depending on the average arrival rate of the EF traffic. Bursts of non-EF traffic – non-EF traffic that arrives at the interface at rates above the current limit for non-EF traffic – also pass through the interface, provided that sufficient tokens are available in the 100 KB bandwidth bucket. Aggregate traffic in excess of the currently configured bandwidth or burst size are rate-limited using the action specified for the aggregate policer, which in this example is set to a high PLP.

The premium traffic is policed by both the premium policer and aggregate policer. Although the premium policer rate-limits the premium traffic, the aggregate policer decrements the credits but does not drop the packets. The aggregate policer rate-limits the non-premium traffic. Therefore, the premium traffic is assured to have the bandwidth configured for premium, and the non-premium traffic is policed to the remaining bandwidth.

Understanding Dynamic Multicast

Dynamic IGMP Configuration Overview

The Internet Group Management Protocol (IGMP) is a host to router signaling protocol for IPv4 used to support IP multicasting. This protocol manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

Subscriber access supports the configuration of IGMP within the dynamic profiles hierarchy. By specifying IGMP statements within a dynamic profile, you can dynamically apply IGMP configuration when a subscriber connects to an interface using a particular access technology (DHCP), enabling the subscriber to access a carrier (multicast) network.

Subscriber Management IGMP Model Overview

In an IPTV network, channel changes occur when a set-top box (STB) sends IGMP commands that inform an upstream device (for example, a multiservice access node [MSAN] or services router) whether to start or stop sending multicast groups to the subscriber. In addition, IGMP hosts periodically request notification from the STB about which channels (multicast groups) are being received.

You can implement IGMP in the subscriber management network in the following ways:

- Static IGMP – All multicast channels are sent to the MSAN. When the MSAN receives an IGMP request to start or stop sending a channel, it adds the subscriber to the multicast group and then discards the IGMP packet.
- IGMP Proxy – Only multicast channels currently being viewed are sent to the MSAN. If the MSAN receives a request to view a channel that is not currently being forwarded to the MSAN, it forwards the request upstream. However, the upstream device does not see all channel change requests from each subscriber, limiting bandwidth control options.
- IGMP Snooping – Only multicast channels currently being viewed are sent to the MSAN. The MSAN forwards all IGMP requests upstream, unaltered, even if it is already receiving the channel. The upstream device sees all channel change requests from each subscriber. Using IGMP snooping enables the broadband services router to determine the mix of services and the bandwidth requirements of each subscriber and adjust the bandwidth made available to each service.
- IGMP Passthrough – The MSAN transparently passes IGMP packets upstream to the broadband services router.

IGMP hosts (sources) also periodically verify that they are sending the correct traffic by requesting that each client send information about what multicast groups it wants to receive. The responses to this *IGMP query* can result in a substantial upstream traffic burst.

IGMPv2 is the minimum level required to support IPTV, and is the most widely deployed. Emerging standards specify IGMPv3.

Dynamic MLD Configuration Overview

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each router maintains a list of host multicast addresses that have listeners for each subnet, as well as a timer for each address. However, the router does not need to know the address of the listeners – just the address of the hosts. The router provides addresses to the multicast routing protocol it uses; this ensures that multicast packets are delivered to all subnets where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) protocol.

Subscriber access supports the configuration of MLD within the dynamic profiles hierarchy for dynamically created interfaces. By specifying MLD statements within a dynamic profile, you can dynamically apply MLD configuration when a subscriber connects to an interface using a particular access technology (DHCP), enabling the subscriber to access a carrier (multicast) network.

Understanding Subscriber Secure Policy

Subscriber secure policy enables you to mirror traffic on a per-subscriber basis. You can mirror the content of subscriber traffic as well as monitor events related to the subscriber session that is being mirrored.

Subscriber secure policy (SSP) mirroring can be based on information provided by either RADIUS or Dynamic Tasking Control Protocol (DTCP), and can mirror both IPv4 and IPv6 traffic. Configuration of subscriber secure policy mirroring is independent of the actual mirroring session – you can configure the mirroring parameters at any time. Also, you can use a single RADIUS or DTCP server to provision mirroring operations on multiple routers in a service provider's network. To provide security, the ability to configure, access, and view the subscriber secure policy components and configuration is restricted to authorized users.

After subscriber secure policy is triggered, both the subscriber incoming and outgoing traffic are mirrored. The original traffic is sent to its intended destination and the mirrored traffic is sent to a mediation device for analysis. The actual mirroring operation is transparent to subscribers whose traffic is being mirrored. A special UDP/IP header is prepended to each mirrored packet sent to the mediation device. The mediation device uses the header to differentiate multiple mirrored streams that arrive from different sources.

Support for Intercepting Both Layer 2 and Layer 3 Datagrams

When DTCP- or RADIUS-initiated SSP intercepts traffic on logical subscriber interfaces and VLAN subscriber interfaces, it sends both Layer 2 and Layer 3 datagrams to the mediation device. When you enable subscriber secure policy for these interfaces, traffic for all configured families (inet, inet6) including Layer 2 and Layer 3 control traffic is mirrored.

Traffic Filtering for DTCP-Initiated Subscriber Secure Policy Mirrored Traffic

You can filter mirrored traffic before it is sent to a mediation device. With this feature, service providers can reduce the volume of traffic sent to a mediation device. For some types of traffic, such as IPTV or video on demand, you do not need to mirror the entire content of the traffic because the content may already be known or controlled by the service provider.

Mirroring-Related Event Reporting

Subscriber secure policy also supports the use of SNMPv3 traps to report events related to the mirroring operation to an external device. Types of information sent in traps include identifying information for subscribers, such as username or IP address, and subscriber session events, such as login or logout events or mirroring session activation or deactivation. The traps map to messages defined in the *Lawfully Authorized Electronic Surveillance (LAES) for IP Network Access, American National Standard for Telecommunications*.

Support for L2TP Subscribers

Both DTCP-initiated and RADIUS-initiated SSP can be applied to Point-to-Point Protocol (PPP) subscribers whose traffic is tunneled with Layer 2 Tunneling Protocol (L2TP). DTCP SSP supports subscribers only at the L2TP network server (LNS), whereas RADIUS-initiated SSP supports subscribers at the L2TP access concentrator (LAC) or the LNS.

At the LAC, both subscriber ingress traffic (from the subscriber into the tunnel) and subscriber egress traffic (from the tunnel to the subscriber) are mirrored at the subscriber-facing ingress interface. The ingress traffic is mirrored after PPPoE de-encapsulation and before L2TP encapsulation. The egress traffic is mirrored after L2TP de-encapsulation. The mirrored packet includes the complete HDLC frame sent to the LNS rather than only the IP datagram.

At the LNS, both subscriber ingress traffic (from the LAC to the LNS) and subscriber egress traffic (from the LNS to the LAC) are mirrored at the inline services (si) interface corresponding to the subscriber. Ingress traffic is mirrored after de-encapsulation of L2TP, HDLC, and PPP headers. The egress traffic is mirrored before the IP datagram is encapsulated. The mirrored traffic contains only the IP datagram belonging to the subscriber.

There is no specific L2TP SSP configuration.

RADIUS-Initiated Subscriber Secure Policy Overview

RADIUS-initiated mirroring creates secure policies based on RADIUS VSAs and uses RADIUS attributes to identify the subscriber whose traffic is to be mirrored. Mirroring is initiated without regard to the subscriber location, router, interface, or type of traffic.

The mirroring operation can be initiated by RADIUS messages as follows:

- Subscriber login – Mirroring starts when the subscriber logs in and the router receives the trigger in a RADIUS Access-Accept message. Using triggers in RADIUS Access-Accept messages enables you to mirror per-subscriber traffic without regard to how often the subscriber logs in or out, or which router or interface the subscriber uses.
- In-session – Mirroring starts when the router receives the trigger in a RADIUS change of authorization request (CoA-Request) message. Using triggers in CoA-Request messages enables you to immediately mirror traffic of a subscriber who is already logged in.

DTCP-Initiated Subscriber Secure Policy Overview

Dynamic Tasking Control Protocol (DTCP)-initiated mirroring creates secure policies to mirror traffic for the subscriber based on DTCP messages. The attributes in a DTCP ADD message trigger the router to start mirroring traffic and specify the interface on which the mirroring takes place. The mirroring operations can be initiated by DTCP messages as follows:

- Subscriber login – Mirroring starts on the specified interface when the subscriber logs in. The DTCP ADD message must be sent to the router before the subscriber logs in.
- In-session – Mirroring starts for all subscribers that match the trigger supplied in the DTCP ADD message when the router receives a DTCP ADD message.

Subscriber Secure Policy Support for IPv4 Multicast Traffic

IP multicast traffic is used for applications such as audio or video streaming, IPTV, video conferencing, or online gaming. Multicast traffic is sent to multiple subscribers who have joined a multicast group.

Secure subscriber policy allows for the mirroring of IPv4 multicast traffic sent to a specific subscriber. If multiple subscribers whose traffic requires mirroring join the same multicast session, the subscriber secure policy feature mirrors each subscriber's traffic and forwards it separately to the mediation device with the proper prepended header.

Mirroring of multicast traffic is supported only for subscribers in the default logical system.

You can enable and disable the mirroring of multicast traffic on a per-chassis basis. You cannot enable or disable it on a per-subscriber basis.

Triggering the Mirroring of IPv4 Multicast Traffic

Multicast traffic being sent towards a subscriber does not contain much of the identifying information used to trigger mirroring of a subscriber's unicast traffic. For example, the multicast packet contains the multicast group address in the destination address of the packet instead of the subscriber's IP address. It also does not contain the user name or MAC address of the subscriber, and does not include information obtained by RADIUS or DHCP. Therefore, methods of identifying multicast traffic that is received by a subscriber are not the same as methods of identifying a subscriber's unicast traffic or multicast traffic that is sent by a subscriber.

To join a multicast group, a subscriber sends an IGMP join request, and it receives a reply. The reply contains the multicast groups to which the subscriber is registered. Triggering the mirroring of multicast traffic is based on the sending of the IGMP join request and the information in the IGMP reply. If the subscriber's unicast traffic is already being mirrored either through DTCP-initiated or RADIUS-initiated traffic mirroring, and the subscriber sends an IGMP join request, mirroring of multicast traffic sent to the subscriber is initiated. The traffic being mirrored is based on the groups contained in the IGMP reply.

Understanding Hierarchical Class of Service

Hierarchical class of service (HCoS) is the ability to apply traffic schedulers and shapers to a hierarchy of *scheduler nodes*. Each level of the scheduler hierarchy can be used to shape traffic based on different criteria such as application, user, VLAN, and physical port.

This allows you to support the requirements of different services, applications, and users on the same physical device and physical infrastructure.

HCoS is implemented primarily using traffic classifiers at the ingress and hierarchical schedulers and shapers at the egress.

A classifier is a filter that labels traffic at the device ingress based on configurable parameters such as application or destination. Traffic is classified into what is called a forwarding equivalence class (FEC). The FEC defines a class of traffic that receives common treatment.

Schedulers, and their associated shapers, are the function that controls the traffic bandwidth, jitter (delay variation), and packet loss priority at the egress of the device.

Hierarchical schedulers are used to apply multiple levels of scheduling and shaping with each level applied to different classifications such as forwarding equivalence class, VLAN, and physical interface (port) as shown in Figure 9.

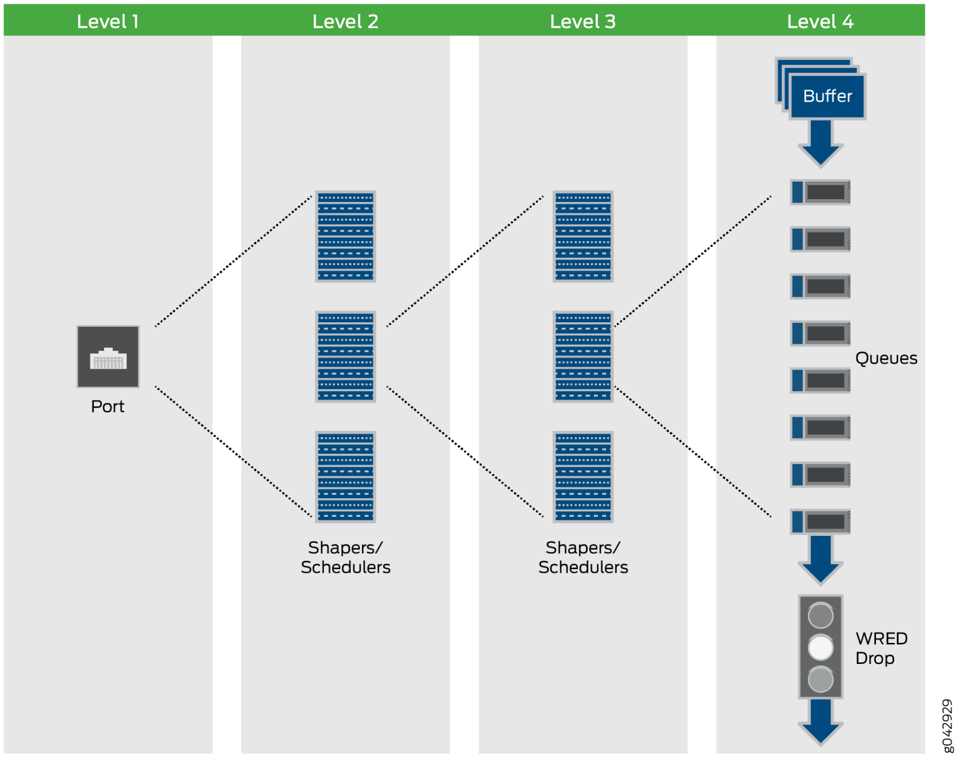

Figure 9 Hierarchical Scheduler(s)

NOTE Hierarchical class of service is also referred to as Hierarchical Quality of Service (HQoS) in other vendor's documentation.

A typical application of HCoS is to configure multiple levels of egress schedulers and shapers, at the subscriber edge, using dynamic profiles to provide traffic shaping and prioritization at the subscriber VLAN level and for multiple classes of traffic.

Dynamic profiles are a mechanism that allows you to dynamically apply schedulers and shapers to individual subscribers or groups of subscribers.

The Junos OS hierarchical schedulers support up to five levels of scheduler hierarchies on MX Series devices when using enhanced queuing Dense Port Concentrators (DPCs) or fine-grained queuing Modular Port Concentrators (MPCs), and Modular Interface Cards (MICs). It is important to know the capabilities of your hardware with respect to HCoS. The following are a few tips to help you:

- Only certain hardware supports the five-level scheduler hierarchy of HCoS.
- The number of queues and logical interfaces supported is dependent upon exactly what hardware you are using.
- The MX Series Packet Forwarding Engine handles guaranteed bandwidth and scheduler node weight differently than other Packet Forwarding Engines.
- The fine-grained queuing MPCs and MICs have a certain granularity with respect to the shaping and delay buffer values. The values used are not necessarily exactly the values configured.

MORE? To learn more about platform support for HCoS, use the Juniper Networks Feature Explorer (http://pathfinder.juniper.net/feature-explorer/). In the Feature Explorer, search on *hierarchical schedulers*.

In addition, it is important to note the following:

- HCoS is most frequently used to enforce service level agreements at the subscriber edged using dynamic traffic control profiles.
- Hierarchical schedulers can also be applied to Ethernet pseudowire interfaces, aggregated Ethernet interfaces, Layer 2 Tunnel Protocol (L2TP) network server (LNS) inline services, and GRE tunnels.
- Hierarchical ingress policing is a feature that is complimentary to and often used in conjunction with HCoS.
- There are other features in Junos OS that have similar sounding names.

NOTE Hierarchical class of service is also referred to as Hierarchical Quality of Service (HQoS) in other vendor's documentation.

Hierarchical Class of Service Network Scenarios

Hierarchical class of service (HCoS) can be used to provide granular control of traffic for a variety of different applications.

NOTE Hierarchical class of service is also referred to as Hierarchical Quality of Service (HQoS) in other vendor's documentation.

Hierarchical class of service is most frequently used in the following scenarios:

Services to Subscribers

Multiservice network operators face a challenge to provide different types of services on the same infrastructure to residential and business subscribers. The network operator needs to make sure each subscriber gets the network resources they paid for and each service gets the network resources it needs to operate properly.

If no CoS is applied, one service could consume most of the bandwidth of the transmission infrastructure and starve the other services.

Using hierarchical class of service, the network edge device can have up to five levels of scheduling and prioritization. So the traffic can be shaped and prioritized per customer and per service type. Controlling traffic in this way provides the ability to deliver the required service level for each subscriber for each service type.

By allowing network operators to consolidate different services and multiple customers on the same physical infrastructure, hierarchical class of service helps maximize the ability to offer revenue generating services while simultaneously minimizing capital cost.

Services to Businesses

Hierarchical class of service is a valuable tool for service providers that support business customers who are running applications with different prioritization and scheduling requirements over the same infrastructure. In this scenario hierarchical class of service allows lower priority traffic to fully utilize the available bandwidth on a port, while simultaneously ensuring low latency and guaranteed bandwidth to higher priority traffic on the same port.

This allows a provider to consolidate different services on the same physical device and physical infrastructure thus optimizing network resources while maintaining the required level of service.

All of this maximizes revenue and minimizes cost.

Wireless Backhaul

In a cellular network the operator might want to offer business services along with its cell tower traffic. One of the main challenges is to make sure that the time-sensitive cell traffic is not affected by the business services running on the same infrastructure. Each type of traffic has its own priority flows and bandwidth constraints. For example, wireless backhaul is very sensitive to fluctuations in the packet stream (Jitter) because it relies on synchronization.

In this scenario, hierarchical class of service allows each type of traffic to receive the required resources and quality of service while being delivered over the same infrastructure.

By consolidating different services on the same physical infrastructure, HCoS helps maximize revenue and minimize cost.

Understanding Hierarchical Scheduling

Hierarchical class of service (HCoS) is a set of capabilities that enable you to apply unique CoS treatment for network traffic based on criteria such as user, application, VLAN, and physical port.

This allows you to support the requirements of different services, applications, and users on the same physical device and physical infrastructure.

Scheduler Node-Level Designations in Hierarchical Scheduling

Scheduler hierarchies are composed of nodes and queues. Queues terminate the hierarchy. Nodes can be either root nodes, leaf nodes, or internal (non-leaf) nodes. Internal nodes are nodes that have other nodes as “children” in the hierarchy.

Scheduler hierarchies consist of levels, starting with Level 1 at the physical port. This article establishes a four-level scheduler hierarchy which, when fully configured, consists of the physical interface (Level 1), the interface set (Level 2), one or more logical interfaces (Level 3), and one or more queues (Level 4).

NOTE Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies. The concepts presented in this topic apply similarly to five scheduler hierarchy levels.

Table 8 describes the possible combinations of scheduler nodes and their corresponding node level designations for a hierarchical queuing MIC or MPC.

Scheduler Configuration for Hierarchical CoS	Hierarchical CoS Scheduler Nodes			
	Root Node	Internal (Non-Leaf) Nodes		Leaf Node
	Level 1	Level 2	Level 3	Level 4
One or more traffic control profiles configured on logical interfaces, but no interface-sets configured	Physical interface	–	One or more logical interfaces	One or more queues
Interface-sets (collections of logical interfaces) configured, but no traffic-control profiles configured on logical interfaces	Physical interface	–	Interface-set	One or more queues
Fully configured scheduler nodes	Physical interface	Interface-set	One or more logical interfaces	One or more queues

Table 8 Node Levels Designations in Hierarchical Scheduling

Table 8 illustrates how the configuration of an interface set or logical interface affects the terminology of hierarchical scheduler nodes. For example, suppose you configure an interface-set statement with logical interfaces (such as unit 0 and unit 2) and a queue. In this case, the interface-set is an internal node at Level 2 of the scheduler node hierarchy. However, if there are no traffic control profiles attached to logical interfaces, then the interface set is at Level 3 of the hierarchy.

Hierarchical Scheduling at Non-Leaf Nodes

Whereas standard CoS scheduling is based on the scheduling and queuing characteristics of a router's egress ports and their queues, hierarchical CoS scheduling is based on the scheduling and queuing characteristics that span a hierarchy of *scheduler nodes* over a port. The hierarchy begins at Level 1, a *root node* at the physical interface (port) level of the CLI hierarchy and terminates at Level 4, a *leaf node* at the queue level. Between the root and leaf nodes of any scheduler hierarchy are one or more *internal nodes*, which are non-root nodes that have other nodes as "children" in the hierarchy.

Whereas you configure standard CoS scheduling by applying a scheduler map to each egress port to specify a forwarding class and a queue priority level, you configure hierarchical CoS scheduling with additional parameters. To configure hierarchical CoS scheduling, you apply a scheduler map to the queue level (Level 4) of a scheduler hierarchy, and you can apply a different traffic control profile at each of the other levels. A traffic control profile specifies not only a scheduler map (forwarding class and queue priority level) but also optional shaping rate (PIR), guaranteed transmit rate (CIR), burst rate, delay buffer rate, and drop profile.

Hierarchical Schedulers and Traffic Control Profiles

When used, the interface set level of the hierarchy falls between the physical interface level (Level 1) and the logical interface (Level 3). Queues are always Level 4 of the hierarchy.

NOTE Beginning with Junos OS Release 16.1, certain MPCs on MX Series devices support up to five levels of scheduler hierarchies. The concepts presented in this topic apply similarly to five scheduler hierarchy levels.

Hierarchical schedulers add CoS parameters to the interface-set level of the configuration. They use traffic control profiles to set values for parameters such as shaping rate (the peak information rate [PIR]), guaranteed rate (the committed information rate [CIR] on these interfaces), scheduler maps (assigning queues and resources to traffic), and so on.

The following CoS configuration places the following parameters in traffic control profiles at various levels:

- Traffic control profile at the port level (tcp-port-level1):
- A shaping rate (PIR) of 100 Mbps
- A delay buffer rate of 100 Mbps
- Traffic control profile at the interface set level (tcp-interface-level2):
- A shaping rate (PIR) of 60 Mbps
- A guaranteed rate (CIR) of 40 Mbps

- Traffic control profile at the logical interface level (tcp-unit-level3):
- A shaping rate (PIR) of 50 Mbps
- A guaranteed rate (CIR) of 30 Mbps
- A scheduler map called smap1 to hold various queue properties (level 4)
- A delay buffer rate of 40 Mbps

In this case, the traffic control profiles look like this:

```
[edit class-of-service traffic-control-profiles]
tcp-port-level1 { # This is the physical port level
   shaping-rate 100m;
   delay-buffer-rate 100m;
}
tcp-interface-level2 { # This is the interface set level
   shaping-rate 60m;
   guaranteed-rate 40m;
}
tcp-unit-level3 { # This is the logical interface level
   shaping-rate 50m;
   guaranteed-rate 30m;
   scheduler-map smap1;
   delay-buffer-rate 40m;
}
```

Once configured, the traffic control profiles must be applied to the proper places in the CoS interfaces hierarchy.

```
[edit class-of-service interfaces]
interface-set level-2 {
   output-traffic-control-profile tcp-interface-level-2;
}
ge-0/1/0 {
   output-traffic-control-profile tcp-port-level-1;
   unit 0 {
      output-traffic-control-profile tcp-unit-level-3;
   }
}
```

In all cases, the properties for level 4 of the hierarchical schedulers are determined by the scheduler map.

Understanding Hierarchical Scheduling for MIC and MPC Interfaces

Scheduler Node Scaling for MIC and MPC Interfaces

In per-unit scheduling, the logical interfaces share a common level 2 node (one per port). In hierarchical-scheduling, each logical interface has its own level 2 node. Thus, scaling is limited by the number of level 2 nodes.

To better control system resources in hierarchical-scheduling mode, you can limit the number of scheduler node levels to two. In this case, all logical interfaces and interface sets with CoS scheduling policy share a single level 2 node. Consequently, the maximum number of logical interfaces with CoS scheduling policies is increased (the interface sets must be at level 3).

To configure scheduler node scaling, include the hierarchical-scheduler statement and set the `maximum-hierarchy-levels` option to 2 at the [edit interfaces xe-*fpc/pic/port*] hierarchy level.

```
[edit interfaces]
xe-2/0/0 {
   hierarchical-scheduler {
      maximum-hierarchy-levels 2;
   }
}
```

NOTE The maximum-hierarchy-levels option supports level 3 interface sets but not level 2 interface sets. If you configure level 2 interface sets with the maximum-hierarchy-levels option, you generate Packet Forwarding Engine errors.

Hierarchical Scheduling Priority Levels for MIC and MPC Interfaces

The queuing model used by MIC and MPC interfaces supports three priority levels for guaranteed scheduling priority and two lower priority levels for excess scheduling priority. You can configure a queue with one guaranteed priority and one excess priority. For example, you can configure a queue for guaranteed low (GL) as the guaranteed priority and configure excess high (EH) as the excess priority.

You can associate a guaranteed level with only one excess level. You can associate an excess level with any number of guaranteed priority levels, including none.

Interface nodes maintain their guaranteed priority level (for example, guaranteed high, GH) as long as they do not exceed their guaranteed bandwidth. If the queue bandwidth exceeds the guaranteed rate, then the priority drops to the excess priority (for example, excess high, EH). Because excess level priorities are lower than their guaranteed counterparts, the bandwidth guarantees for each of the other levels can be maintained.

Guaranteed Bandwidth and Weight of an Interface Node on MIC and MPC Interfaces

The queuing model used by MIC and MPC interfaces separates the concepts of *guaranteed bandwidth* and *weight* of an interface node, although the two terms are often used interchangeably. The guaranteed bandwidth for an interface node is the bandwidth the node can use, independent of what is happening at the other nodes of the scheduling hierarchy. The weight of an interface node, on the other hand, is a value that determines how *excess bandwidth* is used. The weight of a node comes into play when other nodes at the same hierarchical scheduling level use less than the sum of their guaranteed bandwidths.

For some application traffic types (such as constant bit rate voice, where there is little concern about excess bandwidth), the guaranteed bandwidth dominates the node. For other types of application traffic (such as bursty data, where a well-defined bandwidth is not always possible), the concept of weight dominates the node.

Hierarchical Scheduling for MIC and MPC Interfaces in Oversubscribed PIR Mode

In contrast to the Intelligent Queuing Enhanced (IQE) and Intelligent Queuing 2 Enhanced (IQ2E) PICs, the interfaces on MICs and MPCs set the guaranteed rate to zero in oversubscribed peak information rate (PIR) mode for the per-unit scheduler. Also, the configured rate is scaled down to fit the oversubscribed value. For example, if there are two logical interface units with a shaping rate of 1 Gbps each on a 1-Gbps port (which is, therefore, oversubscribed 2 to 1), then the guaranteed rate on each unit is scaled down to 500 Mbps (scaled down by 2).

With hierarchical schedulers in oversubscribed PIR mode, the guaranteed rate for every logical interface unit is set to zero. This means that the queue transmit rates are always oversubscribed.

Because in oversubscribed PIR mode the queue transmit rates are always oversubscribed, the following are true:

- If the queue transmit rate is set as a percentage, then the guaranteed rate of the queue is set to zero; but the excess rate (weight) of the queue is set correctly.
- If the queue transmit rate is set as an absolute value and if the queue has guaranteed high or medium priority, then traffic up to the queue's transmit rate is sent at that priority level. However, for guaranteed low traffic, that traffic is demoted to the excess low region. This means that best-effort traffic well within the queue's transmit rate gets a lower priority than out-of-profile excess high traffic. This differs from the IQE and IQ2E PICs.

Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview

On MX Series routers, you can apply hierarchical schedulers on aggregated ethernet bundles using interface sets. This feature enables you to configure a group of virtual LANs (VLANs) and control their bandwidth. This feature is supported at egress only.

You can configure interface sets for aggregated Ethernet (AE) interfaces created under static configurations. You can configure class-of-service parameters on AE interfaces, in either link-protect or non-link-protect mode. You can configure these parameters at the AE physical interface level. The CoS configuration is fully replicated for all AE member links in link-protect mode. You can control the way these parameters are applied to member links in non-link-protect mode by configuring the AE interface to operate in scaled mode or replicate mode.

The link membership list and scheduler mode of the interface set is inherited from the underlying aggregated Ethernet interface over which the interface set is configured. When an aggregated Ethernet interface operates in link protection mode, or if scheduler mode is configured to replicate member links, the scheduling parameters of the interface set are copied to each of the member links.

If the scheduler mode of the aggregated Ethernet interface is set to scale member links, the scheduling parameters are scaled based on the number of active member links (scaling factor is 1/A where A is the number of active links in the bundle) and applied to each of the AE interface member links.

To configure an interface set, include the interface-set statement at the [edit class-of-service interfaces] hierarchy level.

To apply scheduling and queuing parameters to the interface set, include the output-traffic-control-profile *profile-name* statement at the [edit class-of-service interfaces *interface-name* interface-set *interface-set-name*] hierarchy level.

To apply an output traffic scheduling and shaping profile for the remaining traffic to the logical interface or interface set, include the output-traffic-control-profile-remaining *profile-name* statement at the [edit class-of-service interfaces *interface-name*] hierarchy level or the [edit class-of-service interfaces *interface-name* interface-set *interface-set-name*] hierarchy level.

Hierarchical Class of Service for Subscriber Management Overview

The hierarchical class-of-service (HCoS) architecture as supported on fine-grained queuing MPCs is a powerful feature designed to provide a flexible and scalable CoS solution in broadband network gateway (BNG) subscriber access applications where triple-play or business class offerings are enabled through IP CoS.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters (which can include a delay-buffer bandwidth) and packet transmission scheduling parameters (which can include buffer management parameters) to an individual subscriber interface rather than to all interfaces configured on the port. HCoS enables you to dynamically modify queues when subscribers require services.

The logical interface set construct in a five-level scheduler hierarchy is the key feature that enables HCoS. The interface set feature allows you to group subscribers into aggregate classes with specific guaranteed and peak rates that map to service classes. Service classes ultimately map to how much you can charge for the differentiated service levels.

HCoS can be applied dynamically through the use of dynamic traffic profiles and RADIUS vendor-specific attributes (VSAs).

Dynamic traffic profiles are used to dynamically apply CoS to individual subscribers or groups of subscribers. This enables you, as a service provider, to deploy a BRAS solution without having to manually provision each customer. In a dynamic traffic profile, variables are used to represent the values for things like shaping rate and drop priority.

Dynamic traffic profiles are used in conjunction with dynamic profiles. Dynamic profiles allow you to dynamically provision IP service definitions by creating a template configuration and having the specific variable values assigned in real time when the subscriber authenticates to the network.

Understanding Hierarchical CoS for Subscriber Interfaces

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on a port. Hierarchical CoS enables you to dynamically modify queues when subscribers require services.

Hierarchical CoS is supported on MX Series routers with either Enhanced Queuing DPCs or queuing MPCs/MICs installed. Beginning with Junos OS Release 16.1, five levels of hierarchy are supported on MPC5E 3D Q line cards.

Interfaces support up to a five-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), an interface set or underlying interface (level 2), one or more underlying logical interfaces (level 3), one or more session or customer VLANs (level 4), and one or more queues (level 5). Although all CoS scheduling hierarchies are five-level, level 1 is always the physical interface and level 5 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure – for example, a logical interface or an interface set – and where those interfaces reside in the scheduling hierarchy – level 2, level 3, or level 4. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling*, *three-level hierarchical scheduling*, *four-level hierarchical scheduling* in this topic.

Two-Level Hierarchical Scheduling

Two-level hierarchical scheduling limits the number of hierarchical levels in the scheduling hierarchy to two as shown in Figure 10. In this configuration, interface sets are not configured and only the logical interfaces have traffic control profiles (TCPs). Configuring two levels of hierarchy on MPCs that support more levels preserves resources and allows the system to scale higher.

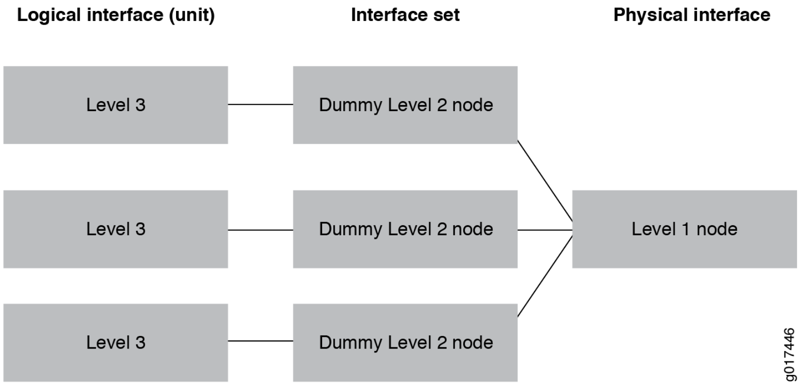

Figure 10 Two-Level Hierarchical Scheduling

In a two-level scheduling hierarchy, all logical interfaces and interface sets share a single node; no hierarchical relationship is formed.

You control two-level hierarchical scheduling by setting the maximum-hierarchy-levels option under the [edit interfaces *interface-name* hierarchical-scheduler] hierarchy to 2:

- If the maximum-hierarchy-levels option is not set, then interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile.
- If any member logical interface has a traffic-control profile, then the interface set is always a level 2 CoS scheduler node.
- If no member logical interface has a traffic-control profile, the interface set is always a level 3 CoS scheduler node.
- If the maximum-hierarchy-levels option is set, then the interface set can only be at level 3; it cannot be at level 2. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.

Table 9 summarizes the interface hierarchy and the CoS scheduler node levels for two-level hierarchical scheduling.

Level 1	Level 2	Level 3
Physical interface	Logical interface	One or more queues
Physical interface	Interface set	One or more queues
Physical interface	Logical interface	One or more queues

Table 9 Two-Level Hierarchical Scheduling – Interface Hierarchy Versus Scheduling Nodes

To configure two-level hierarchical scheduling, include the hierarchical-scheduler statement at the [edit interfaces *interface-name*] hierarchy level and set the maximum-hierarchy-levels option to 2.

```
[edit interfaces]
interface-name {
   hierarchical-scheduler {
      maximum-hierarchy-levels 2;
   }
}
```

Three-Level Hierarchical Scheduling

Three-level hierarchical scheduling is supported only on MX Series routers running MPC/MIC interfaces. Three-level hierarchical scheduling supports up to eight CoS queues. You can configure many different three-level scheduling hierarchies, depending on the location of the interface set or the use of underlying interfaces. In all variations, the physical interface is a level 1 CoS scheduler node and the queues reside at the highest level. Configuring three levels of hierarchy on MPCs that support more levels preserves resources and allows the system to scale higher.

NOTE Three-level hierarchical scheduling is supported only on subscriber interfaces and interface sets running over aggregated Ethernet interfaces on MPC/MIC interfaces in MX Series routers.

When you use three-level hierarchical scheduling, interface sets can reside at either level 3 or level 4. You can also configure an underlying logical interface at level 3 and a logical interface at level 4. Table 10 summarizes the most common cases of the interface hierarchy and the CoS scheduler node levels for three-level hierarchical scheduling.

Level 1	Level 2	Level 3	Level 4
Physical interface	Interface set	Logical interface	One or more queues
Physical interface	Logical interface	Interface set	One or more queues
Physical interface	Underlying logical interface	Logical interface	One or more queues

Table 10 Three-Level Hierarchical Scheduling – Interface Hierarchy Versus CoS Scheduling Node Levels

In three-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, and level 3 form a hierarchical relationship.

With a three-level hierarchical scheduling, logical interfaces can reside at level 2, or they can reside at level 3 if the logical interface at level 2 is an underlying logical interface. This is shown in Figure11.

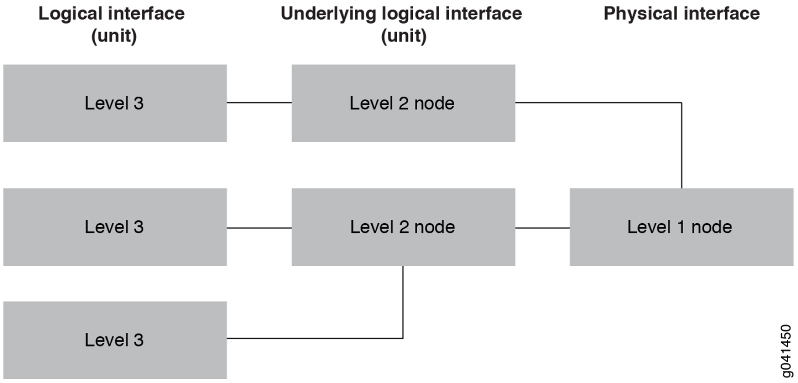

Figure 11 Three-Level Hierarchical Scheduling – Logical Interfaces at Level 3 with Underlying Logical Interfaces at Level 2

Another possible configuration for three-level hierarchical scheduling is shown in Figure 12. In this configuration, the logical interfaces are located at level 2 and the interface sets are located at level 3.

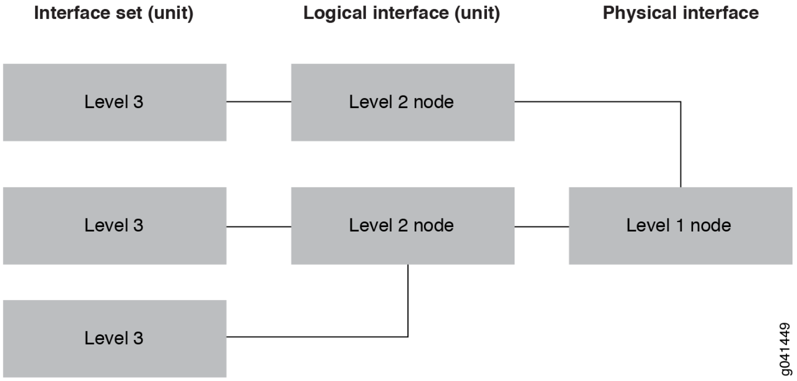

Figure 12 Three-Level Hierarchical Scheduling – Logical Interfaces at Level 2 with Interface Sets at Level 3

To configure three-level hierarchical scheduling, include the implicit-hierarchy option at the [edit interfaces *interface-name* hierarchical-scheduler] hierarchy level and optionally set the maximum-hierarchy-levels option to 3. (The default value for maximum-hierarchy-levels is 3.)

```
[edit interfaces]
```

```
interface-name {
   hierarchical-scheduler {
   implicit-hierarchy;
      maximum-hierarchy-levels 3;
   }
}
```

Interface Hierarchy Versus CoS Hierarchy

An interface hierarchy and a CoS scheduling hierarchy are distinctly different. Interface hierarchy refers to the relationship between the various interfaces – for example, the relationship between logical interfaces and an interface set, the relationship between a logical interface and an underlying logical interface, or the relationship between the physical interface and the logical interface. CoS scheduling hierarchy refers to the hierarchical relationship between the CoS scheduler nodes. In two-level hierarchical scheduling, no hierarchy is formed between the CoS scheduler nodes – the logical interface and interface set share a single level 2 scheduler node. However, when you use the implicit-hierarchy option for three-level hierarchical scheduling, the CoS scheduler nodes form a scheduling hierarchy.

Figure 13 and Figure 14 provide two scenarios for this discussion. Figure 13 shows an interface hierarchy where a Gigabit Ethernet interface (ge-1/0/0) is the physical interface. Two logical interfaces (ge-1/0/0.100 and ge-1/0/0.101) are configured on the physical interface:

- Logical interface ge-1/0/0.100 is a member of a PPPoE interface set and a Demux interface set.
- Logical interface ge-1/0/0.101 is a member of a demux interface set.

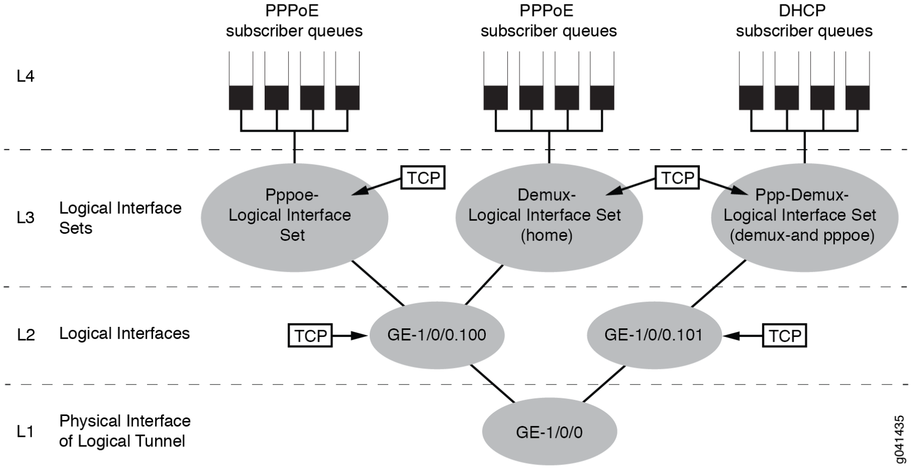

Figure 13 Logical Interfaces at Level 2 and Interface Sets at Level 3

Each interface set has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (interface sets) form a scheduling hierarchy.

To configure this scenario, you must include the implicit-hierarchy option under the hierarchical-scheduler statement on physical interface ge-1/0/0 and configure and apply traffic-control profiles on each interface set and underlying logical interface.

Figure 14 shows an interface hierarchy where Gigabit Ethernet interface ge-1/0/0 is the physical interface. Three logical interfaces are configured:

- Two logical interfaces (Pp0.100 and Demux0.100) reside on the underlying logical interface ge-1/0/0.100.
- A third logical interface (Pp0.101) resides on the underlying logical interface ge-1/0/0.101.

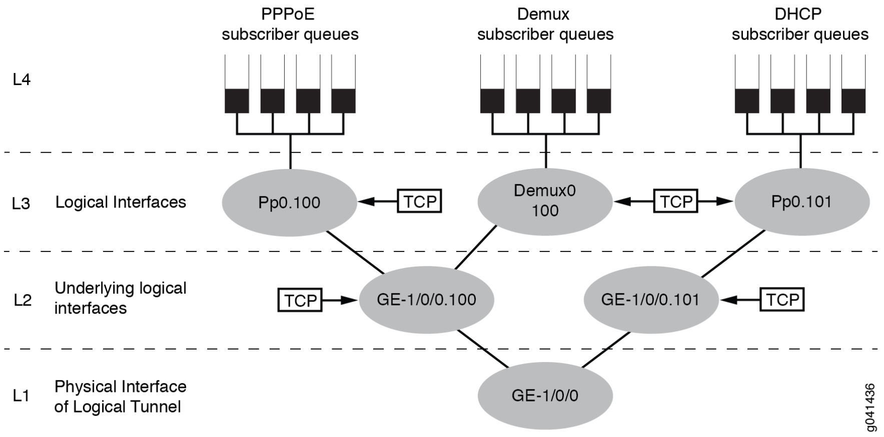

Figure 14 Logical Interfaces at Level 3 and Underlying Logical Interfaces at Level 2

Each logical interface has a dedicated queue. The CoS scheduler nodes at level 1 (physical interface), level 2 (underlying logical interfaces), and level 3 (logical interfaces) form a scheduling hierarchy.

To configure this scenario, you must include the implicit-hierarchy option under the hierarchical-scheduler statement on physical interface GE-1/0/0 and configure and apply traffic-control profiles on each logical interface and underlying logical interface.

You can configure many different three-level scheduling hierarchies; Figure 13 and Figure 14 present just two possible scenarios. Table 10 summarizes the possible interface locations and CoS scheduler nodes.

Four-Level Hierarchical Scheduling

Beginning with Junos OS Release 16.1, four-level hierarchical scheduling is supported on MX Series routers running MPC5E 3D Q interfaces. Four-level hierarchical scheduling supports up to eight class of service queues. In four-level scheduling hierarchies, the physical interface is a level 1 CoS scheduler node and the queues reside at level 5.

NOTE Four-level hierarchical scheduling is not supported agent circuit identifier (ACI) or aggregated Ethernet (AE) interfaces.

When you use four-level hierarchical scheduling, interface sets reside at either level 2 and logical interfaces reside at levels 3 and 4. Table 11 summarizes the most common case of the interface hierarchy and the CoS scheduler node levels for four-level hierarchical scheduling.

Level 1	Level 2	Level 3	Level 4	Level 5
Physical interface	Interface set	Customer VLAN (C-VLAN)	Session Logical Interface (ppp or dhcp)	One or more queues

Table 11 Four-Level Hierarchical Scheduling – Interface Hierarchy Versus CoS Scheduling Node Levels

In four-level hierarchical scheduling, the CoS scheduler nodes at level 1, level 2, level 3, and level 4 form a hierarchical relationship.

To configure four-level hierarchical scheduling, include the implicit-hierarchy option at the [edit interfaces *interface-name* hierarchical-scheduler] hierarchy level and set the maximum-hierarchy-levels option to 4.

```
[edit interfaces]
interface-name {
   hierarchical-scheduler{
      implicit-hierarchy;
      maximum-hierarchy-levels 4;
   }
}
```

Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview

Junos OS supports two aspects of CoS for MPLS pseudowire subscriber interfaces. You can apply CoS rewrite rules and behavior aggregate (BA) classifiers to MPLS pseudowire subscriber interfaces. In addition, CoS performs egress hierarchical shaping towards the subscriber on MPLS pseudowire subscriber interfaces.

Hierarchical CoS enables you to apply traffic scheduling and queuing parameters and packet transmission scheduling parameters to an individual subscriber interface rather than to all interfaces configured on the port. Hierarchical CoS is supported on MX Series routers with either EQ DPCs or MPC/MICs installed.

On Juniper Networks MX Series routers, MPC/MIC and EQ DPC interfaces support a four-level CoS scheduling hierarchy that, when fully configured, consists of the physical interface (level 1), the interface set or the underlying interface (level 2), one or more logical interfaces (level 3), and one or more queues (level 4). Although all CoS scheduling hierarchies are four-level, level 1 is always the physical interface and level 4 is always the queue. Hierarchical scheduling configurations consist of the type of interfaces you configure; for example, a logical interface or an interface set and where those interfaces reside in the scheduling hierarchy, either level 2 or level 3. Because many hierarchical scheduling configurations are possible, we use the terms *two-level hierarchical scheduling* and *three-level hierarchical scheduling* in this discussion.

CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces

CoS supports two-level and three-level hierarchies for MPLS pseudowire subscriber interfaces.

To configure two-level scheduling, include the maximum-hierarchy-levels 2 option under the [edit interfaces *interface-name* hierarchical-scheduler] statement on the physical interface of the logical tunnel anchor point.

To configure three-level hierarchical scheduling, include the implicit-hierarchy option under the [edit interfaces *interface-name* hierarchical-scheduler] statement on the physical interface of the logical tunnel anchor point. Use the following guidelines for configuring the implicit-hierarchy option:

- If an output traffic-control profile is configured on the pseudowire transport interface and on a pseudowire service interface, the two interfaces form a scheduling hierarchy. The pseudowire transport interface resides in a level 2 scheduler node and the pseudowire service interface resides in a level 3 scheduler node.
- If an output traffic-control profile is configured on the pseudowire services interface but not on a pseudowire transport interface, the pseudowire services interface resides in a level 3 scheduler node.
- If an output traffic-control profile is only configured on the pseudowire transport interface and not on the pseudowire services interface, the pseudowire transport interface resides in a level 3 scheduler node and all pseudowire traffic uses this node.
- If the implicit-hierarchy option is not set on the logical tunnel anchor point, logical interfaces behave normally with the hierarchical-scheduler mode configured with or without the hierarchical-scheduler maximum-hierarchy-levels option under the [edit interfaces *interface-name* hierarchical-scheduler] statement. In this case, when you apply a traffic-control profile to the pseudowire and service logical interfaces, they both reside in level 3 scheduler nodes and do not form a scheduling hierarchy, which might not be the desirable behavior. In business edge, where only the pseudowire logical interfaces need to be shaped, applying the traffic-control profile at just the transport logical interface may be sufficient.

When configuring the logical tunnel physical interface for the maximum hierarchy level, all pseudowire logical interfaces operating on the physical interface use the same hierarchy model. If you want to mix two-level and three-level scheduling hierarchies, you can group the pseudowires together by hierarchy levels and share the same logical tunnel anchor point or you can use three-level scheduling for all pseudowires over the anchor point.

To specify rewrite rules and classifiers on pseudowire interfaces, reference the pseudowire device under the [edit class-of-service interfaces] hierarchy level and specify the rewrite rules and classifiers for the pseudowire interfaces.

To control all pseudowire traffic using the same logical tunnel interface, apply CoS policies at the physical interface for the anchor logical tunnel.

CoS for L2TP LAC Subscriber Interfaces Overview

You can apply CoS to the Layer 2 Tunnel Protocol (L2TP) access concentrator (LAC) component.

In Layer 2 Tunnel Protocol (L2TP) configurations, IP and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the L2TP network server (LNS). You can manage the IP header by configuring classifiers and rewrite-rules that transfer the ToS (Type of Service) value or the 802.1p value from the *inner* IP header to the *outer* IP header of the L2TP packet.

Figure 15 shows the classifier `and rewrite rules that` you can configure from the LAC to the LNS, and from the LNS to the LAC.

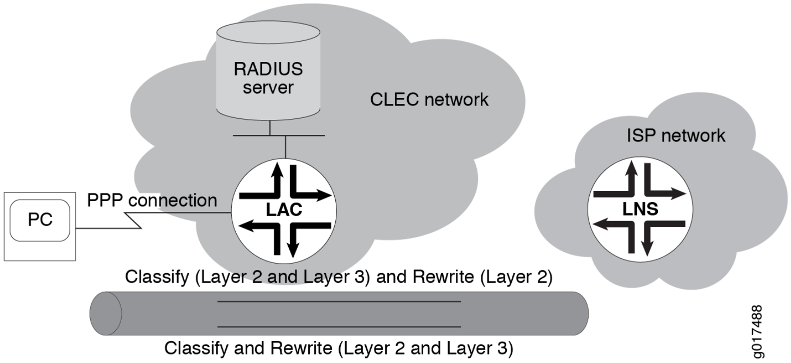
RADIUS server
CLEC network
ISP network
PPP connection
PC
LAC
LNS
Classify (Layer 2 and Layer 3) and Rewrite (Layer 2)
Classify and Rewrite (Layer 2 and Layer 3)
g017488

Figure 15 CoS Configuration for L2TP LAC Topology

Traffic from LAC to LNS

To set the ToS value or the 802.1p value on the inner IP header, you can configure both fixed and behavior aggregate (BA) classifiers for subscribers at Layer 2 or Layer 3 of the network.

Table 12 lists the configuration options for applying classifiers to a subscriber interface on an ingress LAC tunnel.

Classifier	Subscriber Interface
Fixed	Either of the following: PPP interface Underlying VLAN interface

Layer 2	Either of the following: PPP interface Underlying VLAN interface
Layer 3	Family of PPP interfaces

Table 12 Ingress LAC Tunnel Classifier Options

You cannot configure a Layer 2 and fixed classifier together.

The behavior of the Layer 2 and Layer 3 classifiers depends on the configuration. For example, a Layer 3 classifier for a family of PPP interfaces overrides a Layer 2 classifier configured at the PPP interface, except for the unknown packets and control packets.

If you do not configure a classifier for Layer 2, the system applies the default Layer 3 classifier so that tunneled and terminated subscribers have the same behavior. To prevent unknown packets and control packets from being discarded, the system assigns them to the best-effort forwarding class.

For egress tunnels, you configure rewrite rules at the PPP interface to set the ToS or 802.1p value of the outer IP header. Rewrite rules are applied accordingly to the forwarding class, packet loss priority (PLP), and code point.

LAC Tunnels: Traffic from LNS to LAC

On a LAC, mapping the inner IP header to the outer IP header of the L2TP packet depends on the classifier and rewrite-rule configurations. For example, Table 13 lists the values for the classifier and rewrite rules for a VLAN interface. For assured forwarding, the inner 802.1p value (ob001) is classified with the assured-forwarding class and low loss priority at the ingress interface. Based on the assured-forwarding class and low loss priority in the rewrite rule, the ToS value in the outer IP header is set to ob001.

Inner .1p Value	Forwarding Class	Loss Priority	Code Point	Outer ToS Value
ob000	best-effort	Low	000	ob000
ob001	assured-forwarding	Low	001	ob001
ob101	expedited-forwarding	Low	101	ob101
ob111	network-control	Low	11	ob111

Table 13 Sample Result for the Classifier and Rewrite Rules for a VLAN Interface

CoS for L2TP LNS Inline Services Overview

You can apply hierarchical scheduling and per-session shaping to Layer 2 Tunnel Protocol (L2TP) network server (LNS) inline services using a static or dynamic CoS configuration.

This feature is supported on MIC and MPC interfaces on MX240, MX480, and MX960 routers.

Guidelines for Applying CoS to the LNS

In L2TP configurations, IP, UDP, and L2TP headers are added to packets arriving at a PPP subscriber interface on the L2TP access concentrator (LAC) before being tunneled to the LNS.

When a service interface is configured for an L2TP LNS session, it has an *inner* IP header and an outer IP header. You can configure CoS for an LNS session that corresponds to the inner IP header only. The *outer* IP header is used for L2TP tunnel processing only.

However, we recommend that you configure classifiers and rewrite-rules to transfer the ToS (type of service) value from the inner IP header to the outer IP header of the L2TP packet.

Figure 16 shows the classifier and rewrite rules that you can configure on an LNS inline service.

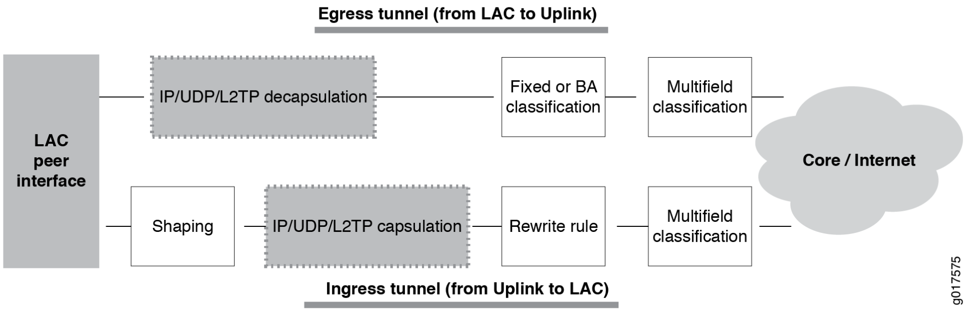
Egress tunnel (from LAC to Uplink)
LAC peer interface
IP/UDP/L2TP decapsulation
Fixed or BA classification
Multifield classification
Core / Internet
Shaping
IP/UDP/L2TP capsulation
Rewrite rule
Multifield classification
Ingress tunnel (from Uplink to LAC)
g017575

Figure 16 Processing of CoS Parameters in an L2TP LNS Inline Service

Hardware Requirements for Inline Services on the LNS

Hierarchical scheduling for L2TP LNS inline services is supported on MIC and MPC interfaces only. The services that you can configure depend on the hardware combination. Table 14 lists the supported inline services and peer interfaces for each MIC and MPC combination.

MPC Module	Inline Service Support–With Per-Session Shaping	Inline Service Support–Without Per-Session Shaping
MPC2E-3D-NG	No	Yes

MPC2E-3D-NG-Q MX80	Yes	Yes
MPC-3D-16XGE-SFPP	No	No

Table 14 Hardware Requirements for L2TP LNS Inline Services

Hierarchical CoS Shaping-Rate Adjustments Overview

This overview describes how MX Series 3D Universal Edge Routers installed in a subscriber access network can adjust hierarchical class-of-service (CoS) parameters to prevent bandwidth contention at subscriber interfaces.

Hierarchical CoS is supported only for subscriber interfaces on Enhanced Queueing (EQ) DPCs or MPC interfaces operating in hierarchical scheduler mode.

The characteristics of voice, data, and video applications vary widely in their requirements for traffic throughput, bandwidth management, delay and jitter tolerance, and buffer depth. To prevent bandwidth contention at subscriber interfaces, you can configure applications such as ANCP and Multicast to perform real-time adjustments to the shaping rate configured for subscriber interfaces for residential gateways. Enabling shaping-rate adjustments on the router can prevent bandwidth contention at the interface from causing degradation of the subscriber's voice, data, or video services.

Types of Shaping-Rate Adjustments

The ANCP application supports *absolute* adjustments to a specific shaping-rate value. You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

The OIF mapping and reverse OIF mapping multicast applications support *delta* adjustments that increase or decrease the current shaping rate by a certain value. The system adjusts traffic destined to the subscriber using reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

Levels of Shaping-Rate Adjustments

Both absolute and delta adjustments are made to a subscriber's aggregate shaping rate on a level 3 scheduler node.

Adjustments that occur on the scheduler node can also impact the shaping rates for all queues. This adjustment can be undesirable for service providers who want to provide a premium level of service on specific queues.

For delta-based adjustments by multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

Figure 17 shows a sample multicast network with shaping rates adjusted at the scheduler node level. The shaping rate is reduced by 4 Mbps (from 41 Mbps to 37 Mbps) at the scheduler node for subscriber interface 1, which reduces the rates of both the best effort and video on demand (VoD) service queues.

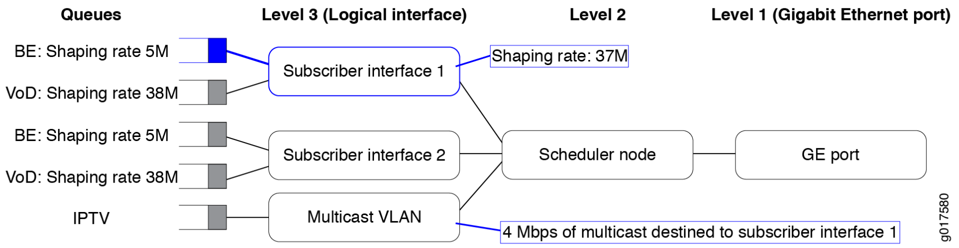

Figure 17 Scheduler Node and Queues with Adjusted Shaping Rates

Figure 18 shows the same network with queue-based adjustments enabled for the best-effort queue on subscriber 1. The shaping rate of the best-effort queue is reduced by 4 Mbps (from 5 Mbps to 1 Mbps). The VoD service queue is not affected.

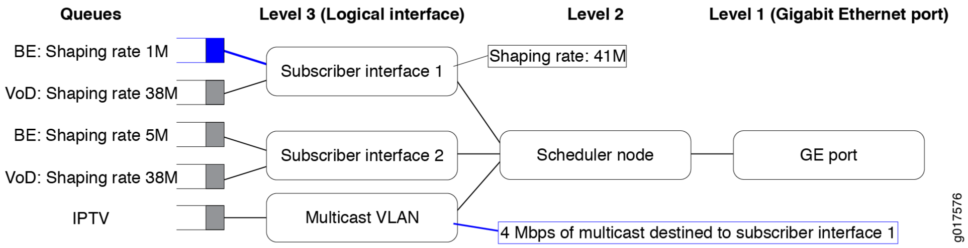

Figure 18 Queue with Adjusted Shaping Rate

Shaping Rate Adjustments for Subscriber Local Loops Overview

This overview describes how an MX Series 3D Universal Edge Router installed as an edge router can adjust hierarchical CoS policy for subscriber interfaces for subscriber local loops. You can configure the router to throttle the traffic sent to subscriber local loops so that the traffic does not exceed the current data transmission rate of those lines. This feature ensures that changes to subscriber local loop speeds do not cause bandwidth contention at the subscriber's residential gateway.

In a typical subscriber access network, traffic destined to a subscriber is delivered from the access network, through an edge router, to a DSLAM. The DSLAM multiplexes subscriber traffic through a DSL, also known as a *local loop*, to the subscriber's residential gateway. When line noise or cross talk in a subcarrier causes the error rate on a DSL to exceed a certain threshold, the DSLAM can adapt itself by lowering the data transmission rate to that carrier device. A lower data transmission rate is less susceptible to induced errors.

You can configure an MX Series router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. Whenever a DSLAM resynchronizes a subscriber local loop speed, the router adjusts the configured shaping rate for that line so that the aggregate egress traffic to those subscribers is shaped to the local loop speed before the traffic reaches the DSLAM. Unless the maximum amount of bandwidth allocated to the subscriber interface on the router is throttled to the local loop speed, bandwidth contention can occur at the subscriber's residential gateway, which can cause the DSLAM to drop packets. This type of shaping-rate adjustment requires the topology discovery and traffic-monitoring features of the Access Node Control Protocol (ANCP).

You can enable ANCP to communicate the subscriber local loop speed to CoS, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. The ANCP agent acquires unadjusted (net) subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS. You can also configure percentage and byte adjustments that the ANCP agent can make to the received net data rate for frame-mode DSL types before communicating the adjusted rate and overhead to CoS.

CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview

To control bandwidth at a household level in a subscriber access network, you can apply RADIUS dynamic class of service (CoS) traffic-shaping attributes to a dynamic interface set and its member subscriber sessions when the subscriber sessions are authenticated. (The dynamic interface set itself does not go through the authentication process.)

A *household* is represented by either a dynamic interface set or a dynamic agent-circuit-identifier (ACI) interface set from which the subscriber sessions originate. For this feature, dynamic interface sets and dynamic ACI interface sets are mapped to Level 2 of the Junos OS CoS scheduler hierarchy, which enables you to use CoS traffic-shaping to shape the bandwidth at the household (interface set) level.

The *subscriber sessions*, also referred to as *subscriber interfaces* or *client sessions*, can be dynamic VLAN, PPPoE, or IP demultiplexing (IP demux) subscriber interfaces. The subscriber interfaces are mapped to Level 3 of the Junos OS CoS scheduler hierarchy.

Supported Network Configurations

Applying RADIUS dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions is supported for the following network configurations:

- Dynamic IP demux subscriber interfaces (for DHCP subscribers) over either a dynamic interface set or a dynamic ACI interface set
- Dynamic PPPoE subscriber interfaces over either a dynamic interface set or a dynamic ACI interface set

Traffic-Control Profiles in Subscriber Interface Dynamic Profiles

To apply dynamic CoS traffic-shaping attributes to a dynamic interface set and its member subscriber sessions, you must define and attach the traffic-control profiles for *both* the dynamic interface set and the dynamic subscriber sessions within the dynamic profile for the subscriber interface.

At the [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles] hierarchy level in the dynamic profile, configure both of the following:

- Traffic-control profile for the dynamic VLAN, PPPoE, or IP demux subscriber interfaces
- Traffic-control profile for the dynamic interface set or dynamic ACI interface set to which the subscriber interfaces belong

RADIUS tag values for the Junos OS CoS traffic shaping predefined variables used in both traffic-control profiles must be in the 100s range.

At the [edit dynamic-profiles *profile-name* interfaces] hierarchy level in the dynamic profile, use the output-traffic-control-profile statement to apply the traffic-control profiles to the dynamic subscriber interface and the dynamic interface set or dynamic ACI interface set.

CoS Traffic Shaping Predefined Variables for Dynamic Interface Sets and Member Subscriber Sessions

The set of $junos-cos-*parameter* predefined dynamic variables has been duplicated and assigned a RADIUS tag value in the 100s range for use with this feature. The RADIUS tag value is the only difference between the existing CoS traffic-shaping predefined dynamic variables and the predefined dynamic variables that you must use with this feature.

Both RADIUS instances of the $junos-cos-*parameter* predefined dynamic variables are available, but you must use the dynamic variables with tag values in the 100s range to apply CoS traffic-shaping attributes to both the dynamic interface set and member subscriber sessions in a subscriber interface dynamic profile.

For example, the existing $junos-cos-shaping-rate predefined variable is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 2. To apply CoS traffic-shaping attributes to the dynamic interface set and its member subscriber sessions, you must instead use the $junos-cos-shaping-rate predefined variable that is assigned RADIUS vendor ID 4874, attribute number 108, and tag value 102.

NOTE Do not configure a combination of $junos-cos-*parameter* predefined dynamic variables with RADIUS tag values in the 100s range and $junos-cos-*parameter* predefined dynamic variables with tag values not in the 100s range in the same traffic-control profile. If you do so, the subscriber authentication process fails.

Understanding Broadband Subscriber Management Wholesale Networks

Layer 2 and Layer 3 Wholesale Overview

In general, wholesaling broadband services allows service providers to resell broadband services and allows other providers to deploy their own services over the incumbent network. There are different methods to partitioning an access network for resale. The two most common approaches are based on either Layer 2 or Layer 3 information. Wholesale access is the process by which the access network provider (the *wholesaler*) partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers (or *retailers*).

In a Layer 3 wholesale configuration, you partition the wholesaler access network at the network layer or the subscriber IP component by associating the IP component with a distinct Layer 3 domain. In a Layer 2 wholesale configuration, you partition the access network at the subscriber circuit or customer VLAN (C-VLAN) by backhauling the connection through the service provider backbone network to the subscribing retailer network where the access traffic can be managed at higher layers.

In a Junos OS Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) subscriber access configuration, wholesale partitioning is accomplished through the use of logical systems and routing instances within the router. Logical systems offer a stricter partitioning of routing resources than routing instances. The purpose behind the use of logical systems is to distinctly partition the physical router into separate administrative domains. This partitioning enables multiple providers to administer the router simultaneously, with each provider having access only to the portions of the configuration relevant to their logical system. Junos OS supports up to 15 named logical systems in addition to the default logical system (that is, inet.0). Unless otherwise specified in configuration, all interfaces belong to the default logical system.

NOTE This Junos OS release supports the use of only the default logical system. Partitioning currently occurs through the use of separate routing instances.

A logical system can have one or more routing instances. Typically used in Layer 3 VPN scenarios, a routing instance does not have the same level of administrative separation as a logical system because it does not offer administrative isolation. However, the routing instance defines a distinct routing table, set of routing policies, and set of interfaces.

Subscriber to Logical System and Routing Instance Relationship

As subscriber sessions are established, subscriber to logical system/routing instance memberships are established by the AAA framework configured for the default logical system. When configuring Layer 3 wholesaling, you typically configure global (wholesale) information within the default (master) logical system and default routing instance. Incoming subscribers must then be authenticated, but this authentication can be handled in one of two ways:

- Single (wholesaler only) authentication – Incoming subscribers are authenticated by the wholesaler RADIUS server. After authentication, the subscribers are assigned values specified by dynamic profiles (routing instances, interfaces, and any configuration values) specific to a particular retailer.
- Dual (wholesaler and retailer) authentication – Sometimes referred to as *double-dip authentication*. Incoming subscribers are initially authenticated by RADIUS using the wholesale configuration. Authenticated subscribers are then redirected to other routing instances associated with individual retailer network space. When you redirect subscribers, and those subscribers are to be authenticated by AAA servers owned by individual retailers, the subscribers must be authenticated again by the AAA servers before they are provided an address and any dynamic profile values are assigned. After reauthentication, however, the subscribers are managed normally using any values specific to the retailer routing instance to which they are assigned.

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview

You can use RADIUS to assign various values through the use of dynamic variables within dynamic profiles. However, the configuration of at least one of the two VSAs described in Table 15 is required for a wholesale network to function.

Attribute Number	Attribute Name	Description	Value
26-1	LSRI-Name	Client logical system/routing instance membership name. Allowed only from RADIUS server for "default" logical system/routing instance membership.	string: logical system:routing instance
26-25	Redirect-LSRI-Name	Client logical system/routing instance membership name indicating to which logical system/routing instance membership the request is redirected for user authentication.	string: logical system:routing instance

Table 15 Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution

Specifying the $junos-routing-instance dynamic variable in a dynamic profile triggers a RADIUS access-accept response of either the LSRI-Name VSA or the Redirect-LSRI-Name VSA. Returning an LSRI-Name attribute in the access-accept response provides the logical system and routing instance in which the logical interface is to be created and the router updates the session database with the specified routing instance value. Returning a Redirect-LSRI-Name attribute in the access-accept response results in the router immediately sending a second access-request message (sometimes referred to as a *double-dip*) to the RADIUS server specified by the logical system:routing instance attribute specified by the Redirect-LSRI-Name VSA.

NOTE Attributes returned as a result of a second access-request message to the logical system/routing instance membership specified by the Redirect-LSRI-Name VSA override any prior attributes returned by initial access-accept responses to the default logical system/routing instance membership.

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management DHCPv4 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv4 relay configuration. However, you can also implement DHCPv4 Relay Proxy or DHCPv4 Local Server configuration.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. Figure 19 illustrates a basic Layer 3 wholesale topology model from which you can expand.

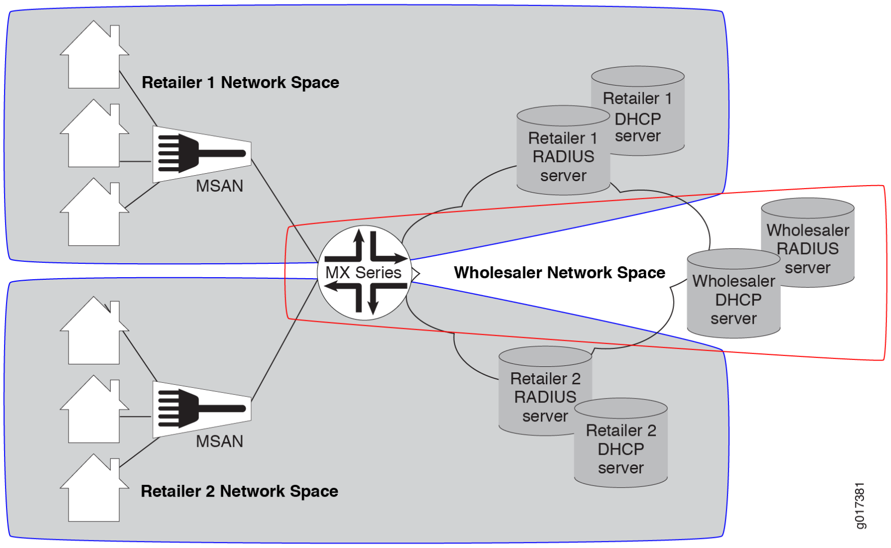

Figure 19 Basic Subscriber Management Layer 3 Wholesale Solution Topology

A DHCP Layer 3 wholesale network solution can use various combinations of the following configuration elements:

- Subscriber network VLAN configuration
- DHCPv4 configuration (DHCPv4 Relay, DHCPv4 Relay Proxy, or DHCPv4 Local Server)
- Addressing server or addressing server access configuration (if not using DHCPv4 Local Server)
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Dynamic profile configuration for retailer access (following subscriber redirection, if applicable)
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management PPPoE Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. Figure 20 illustrates a basic PPPoE Layer 3 wholesale topology model from which you can expand.

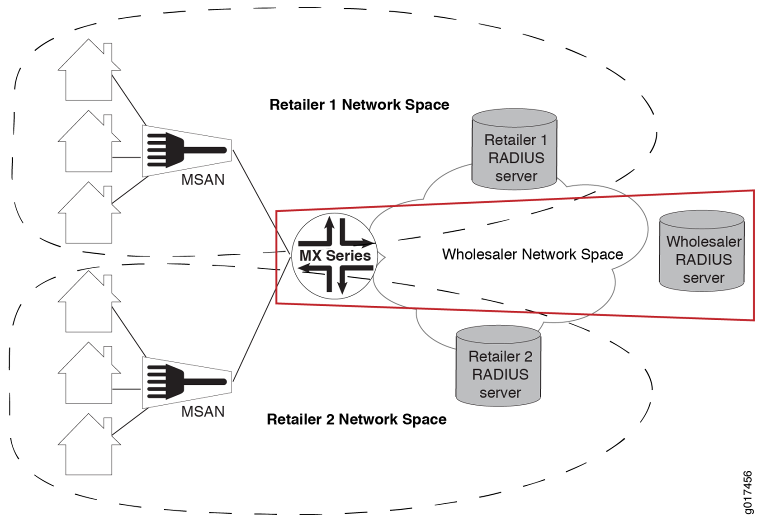

Figure 20 Basic Subscriber Management PPPoE Layer 3 Wholesale Solution Topology

When you are configuring a PPPoE Layer 3 wholesale network solution, the following configuration elements are required:

- Subscriber network VLAN configuration
- Addressing server or addressing server access configuration
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

This implementation of PPPoE Layer 3 wholesale supports the following:

- Dynamic PPPoE interface creation.
- Static VLAN use only.
- AAA server assignment of subscribers to different routing instances within the same (default) logical system only.

Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements

The network topology for the subscriber management Layer 2 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a Virtual Private LAN Service (VPLS) configuration.

Layer 2 wholesale networks are supported on MPC/MIC interfaces.

To explain the concept but limit complexity, this solution provides a configuration with one wholesaler and only two retailers. Figure 21 illustrates a basic Layer 2 wholesale topology model from which you can expand.

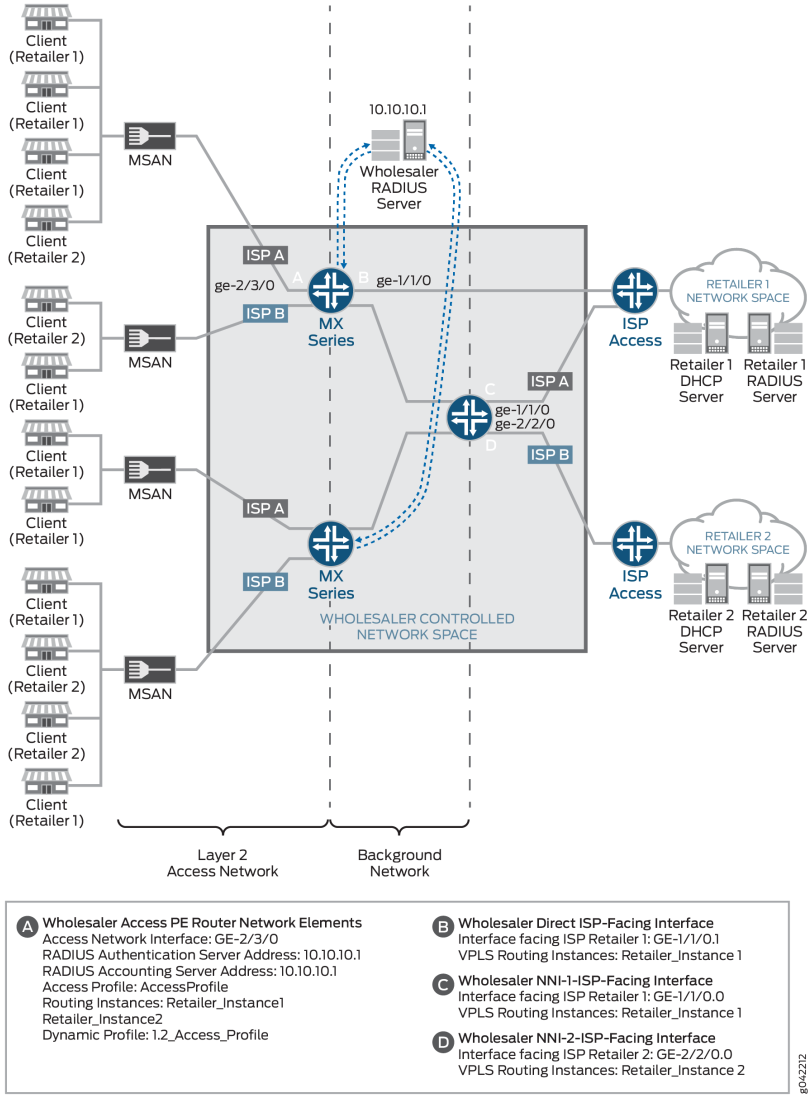

A **Wholesaler Access PE Router Network Elements**
Access Network Interface: GE-2/3/0
RADIUS Authentication Server Address: 10.10.10.1
RADIUS Accounting Server Address: 10.10.10.1
Access Profile: AccessProfile
Routing Instances: Retailer_Instance1
Retailer_Instance2
Dynamic Profile: 1.2_Access_Profile

B **Wholesaler Direct ISP-Facing Interface**
Interface facing ISP Retailer 1: GE-1/1/0.1
VPLS Routing Instances: Retailer_Instance 1

C **Wholesaler NNI-1-ISP-Facing Interface**
Interface facing ISP Retailer 1: GE-1/1/0.0
VPLS Routing Instances: Retailer_Instance 1

D **Wholesaler NNI-2-ISP-Facing Interface**
Interface facing ISP Retailer 2: GE-2/2/0.0
VPLS Routing Instances: Retailer_Instance 2

Figure 21 Basic Subscriber Management Layer 2 Wholesale Solution Topology

When you are configuring a Layer 2 wholesale network solution, the following configuration elements are required:

- Subscriber access dynamic VLAN configuration including dynamic profile configuration for retailer routing instances
- Routing instance configuration for individual retailers on provider edge (PE) routers and network-to-network interface (NNI) routers.
- VLAN interface configuration
- RADIUS server access configuration
- Core network configuration

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

The conventional mechanism for triggering autosensed dynamic VLANs relies on access line attributes provided by PPPoE or DHCP traffic in upstream control packets. Packets of a specified type are exceptioned and authorization depends on fields extracted from the packet as specified in a dynamic profile assigned to the autosensed VLAN range. However, for some wholesale networks, the traffic might not be PPPoE or DHCP. In this case, a different mechanism is required.

Figure 22 shows a sample topology with direct connections between the wholesaler's BNG and the NSP (network service provider) routers for the retailers. Each retailer's network resides in a dedicated routing instance. The wholesaler uses Layer 2 cross-connects to implement the retail networks with 1:1 autosensed, dynamic VLANs and VLAN tag swapping. Core-facing physical interfaces are dedicated to forwarding subscriber connections to the retailer's router. The traffic for an entire outer VLAN can be wholesaled this way. This direct-connect model supports any combination of wholesaler-owned and wholesaled connections for the entire access-facing VLAN range.

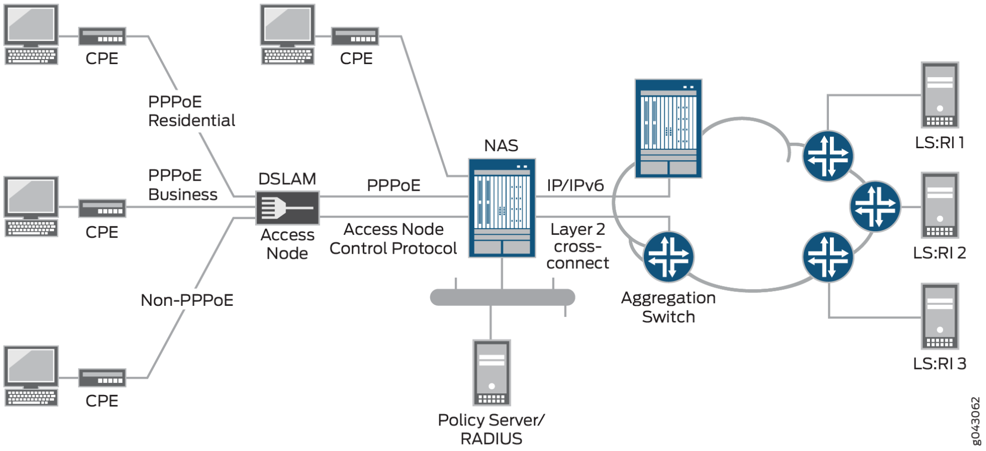

Figure 22 Sample Layer 2 Wholesale Access Topology

A wholesaler providing Layer 2 bitstream access to NSP partners might use this model. Bitstream access enables retailers to offer bidirectional transmission of broadband data and other high-speed services directly to customers across the wholesaler's network. In this topology, the PPPoE residential and subscriber customers are retained by the wholesaler (access provider). The non-PPPoE connections (here multiple connections and subscribers are represented by a single line) can be wholesaled to retail NSPs.

In this model, dynamic VLAN detection and creation for the wholesaled connections do not use in-band control packets. Instead, they rely on an out-of-band protocol, ANCP. ANCP Port Up messages both announce to the ANCP agent on the BNG that new access lines are operational and provide updates about previously announced lines. The messages include ANCP DSL attributes that correspond to Juniper Networks DSL VSAs and DSL Forum VSAs.

Starting in Junos OS Release 16.1, you can configure the ANCP agent to trigger the creation of an autosensed VLAN when the ANCP agent receives a Port Up message where the DSL-Line-State attribute has a value of Showtime. The Showtime state indicates that ports are configured, the subscriber is connected, and the DSL modem is online and ready to transfer data. The other possible values of the attribute, Idle and Silent, are ignored for this purpose and are used by the ANCP agent only to update the ANCP session database (SDB).

During VLAN authorization, RADIUS determines which traffic belongs to the access provider's own subscribers and which belongs to the wholesale customer (retail NSP) based on identification of the subscriber's access line by the agent remote identifier.

When the ANCP agent receives the Port Up message, the agent triggers the autoconfiguration daemon, autoconfd, to initiate the VLAN detection, authorization, and creation processes. Those processes require the following information:

- Three ANCP subscriber access loop attributes (TLVs) that identify the access line and are conveyed in the Port Up message:
  - Access-Loop-Circuit-ID – Access loop circuit identifier used by the ANCP agent to determine which logical interface or interface set corresponds to the subscriber; corresponds to the Juniper Networks Acc-Loop-Cir-ID VSA (26-110).
  - Access-Loop-Remote-ID – Unique identifier of the access line; corresponds to the Juniper Networks Acc-Loop-Remote-ID VSA (26-182).
  - Access-Aggregation-Circuit-ID-Binary – Identifier that represents the outer VLAN tag that the access node inserts on upstream traffic; corresponds to the Juniper Networks Acc-Aggr-Cir-Id-Bin VSA (26-111).
- The name of the physical interface facing the subscriber. This name derives from the local mapping of an ANCP neighbor to the corresponding subscriber-facing access port.
  - The Access-Aggregation-Circuit-ID-Binary attribute and the access-facing interface name together provide information equivalent to that used for conventional autosensed VLAN detection.
  - ANCP Port Down messages indicate that the subscriber access loop is not present or at least is no longer operational. This message triggers the automat-

ic destruction of the dynamic VLAN, regardless of the value of any other ANCP line attribute.

VLAN logical interfaces are created in the default routing-instance unless a nondefault routing instance is provided by local authorization (domain map) or external authorization (RADIUS). Multiple routing instances are required when both access-provider-owned and wholesaled connections are supported at the same time. One routing instance is required for the access provider's own subscribers. An additional routing instance is required for each retail NSP. Consequently, the routing-instance has to be specified when the VLAN is authorized. The RADIUS-based VLAN authorization process determines whether the subscriber access-loop identified by the attributes in the Port Up message is wholesaled to a partner NSP – and therefore maintained as a unique routing-instance – or managed as a subscriber owned by the access provider.

Flat-File Accounting Overview

Accounting statistics can be collected from the Packet Forwarding Engine and reported in an XML flat file, which both contains and describes the data. Starting in Junos OS Release 16.1, you can use a flat-file profile that acts as a template to define attributes for accounting flat files.

Subscriber service accounting statistics are typically collected based on RADIUS Acct-Start and Acct-Stop messages that are sent to a RADIUS server individually or in bulk. Starting in Junos OS Release 17.1, you can alternatively configure service-filter-based accounting statistics to be recorded per subscriber in a local flat file that is not automatically forwarded to a RADIUS server. This configuration collects the running total service statistics per interface family. Service accounting is initiated when the service profile is attached to the interface, whether by a static configuration or a RADIUS Change of Authorization (CoA) message.

When the accounting file is created, a file header is also created if the file format is IP Detail Record (IPDR). The header is not created if the format is comma-separated variable (CSV). The file header includes the following information:

- XML namespace – Static link to the World Wide Web Consortium (W3C) organization's XML Schema Instance (XSI) definition.
- Schema version – Configurable name of the schema that defines the information conveyed in the accounting file. The schema version is associated with a specific XML format and output based on the flat-file profile configuration that is used for the business purpose. This structure enables the XML-formatted contents of the file to be correctly interpreted by the service provider's external file processor.
- NAS ID – Name of the BNG host (network access server) where the accounting statistics are collected.
- File creation timestamp – UTC time zone date and time when the accounting file was created.
- File ID – Number identifying the file. The ID is incremented when a new accounting file is created and can range from 1 through 2,147,483,647.

For example, consider the following sample header for an accounting file for Extensible Subscriber Services Manager (ESSM) business subscribers:

```
<BNGFile xmlns:xsi=http://"www.w3.org/2001/XMLSchema-instance"
 xsi:noNamespaceSchemaLocation="BNG_IPDR_20130423.xsd" NAS-ID="host-
mx480-x5"
 FileCreationTimeStamp="2015-10-09T08:25:50" FileID="29">
<IPDR>
......
......
</IPDR>
</BNGFile>
```

Table 16 lists the elements and their values in the sample header.

Description	Header Element	Value
XML namespace	xmlns	:xsi=http://"www.w3.org/2001/XMLSchema-instance"
schema version	xsi:noNamespaceSchemaLocation	BNG_IPDR_20130423.xsd
NAS ID	NAS-ID	host-mx480-x5
file creation timestamp	FileCreationTimeStamp	2015-10-09T08:25:50
File ID	FileID	29

Table 16 Value of Elements in Sample Accounting Flat File XML Header

You can configure the following options for flat-file accounting at the [edit accounting-options file *filename*] hierarchy level:

- Maximum size of the accounting file.
- Number of files that are saved before overwriting.
- One or more sites where the files are sent for archiving.
- Frequency at which the files are transferred to an archive site.
- Start time for file transfer.
- Compression for the transferred files.
- Local backup on the router for files when transfer fails.
- Whether accounting files are saved when a change in mastership occurs for both the new master Routing Engine and the new backup Routing Engine or for only the new master Routing Engine.
- How long files are kept before being deleted from the local backup directory.

You can also create one or more flat-file profiles at the [edit accounting-options flat-file-profile *profile-name*] hierarchy level that act as templates to specify the following attributes for new accounting files when they are created:

- Statistics fields that you want to collect, such as egress statistics or ingress statistics fields.
- Name and format of the accounting file.

- Frequency at which the Packet Forwarding Engine is polled for the statistics.
- Schema version.

Archive sites provide security and storage for the accounting files, which are transferred at regular intervals. When more than one archive site is configured, the router attempts to transfer the files to the first site on the list. If that fails, the router tries each of the other sites in turn until the transfer either succeeds for one site or fails for all sites. If you configure the last site in the list to be a local directory on the router rather than another remote site, then the files are backed up locally if all remote sites fail. The failed files are simply stored in the designated site. They are not automatically resubmitted to the archival sites. You must use an event script or some other means to have these files resubmitted. Any files remaining in the local directory are deleted when the cleanup-interval expires.

Alternatively, you can use the backup-on-failure statement at the [edit accounting-options file *filename*] hierarchy level to back up the files locally if all the remote attempts fail. If that occurs, the router compresses the accounting files and backs them up to the `/var/log/pfedBackup/` directory. Whenever any of the archive sites is reachable, the router attempts to transfer the data from `/var/log/pfedBackup/` to that site in compressed format. If the transfer of the backed-up files to the reachable site fails, the system tries to transfer the files to any other site that becomes reachable during the transfer interval. Any files that fail to transfer are compressed and kept in `/var/log/pfedBackup/` until an archival site is reachable and the files are successfully transferred. Any files that remain in that directory are deleted when the cleanup-interval expires.

BEST PRACTICE Use the backup-on-failure feature to reliably and automatically back up files and retransmit them to archives rather than relying on a local site listed as the last archive site.

If the backup Routing Engine does not have access to the archive site – for example, when the site is not connected by means of an out-of-band interface or when the path to the site is routed through a line card – you can ensure that the backup Routing Engine's accounting files are backed up by using the push-backup-to-master statement at the [edit accounting-options file *filename*] hierarchy level. When a change in mastership occurs, the new backup Routing Engine saves its files to the `/var/log/pfedBackup/` directory. The master Routing Engine subsequently includes these files when it sends its own accounting files to the archive site at every transfer interval.

To conserve resources during transfer of accounting files and at the archive site, use the compress statement at the [edit accounting-options file *filename*] hierarchy level to compress the files when they are transferred. This option is disabled by default.

A system logging message is generated when a transfer succeeds (transfer-file: Transferred *filename*) or fails (transfer-file failed to transfer). In the event of a failure, an error message is logged to indicate the nature of the failure.

Appendix: Links and URLs

Broadband Subscriber Mgmt Library: http://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/index.html

Subscriber Management and BNG Overview: http://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/subscriber-mgmt-getting-started.html

AAA Service Framework: http://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-aaa.html

http://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-access-activating-services.html

Class of Service: http://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-cos.html

DHCP-based Subscriber Access Configuration: https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/subscriber-management-access-services-overview.html

Components of a Dynamic Profile: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-access-activating-services.html

Understanding Subscriber Access Network: https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/subscriber-mgmt-access-network.html

Multiservice Access Node Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-msan.html

Ethernet MSAN Aggregation Options: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-msan-options.html

Broadband Access Service Delivery Options: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-delivery.html

DHCP Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-dhcp-access.html

Dynamic Profiles for PPP Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/ppp-subscriber-access-dynamic-profiles-attachment-overview.html

L2TP for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscribermanagement-l2tp-overview.html

PPP MPLS Pseudowire Subscriber Logical Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/pseudowire-subscriber-interfaces-overview.html

Understanding Subscriber Management VLAN Architecture: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-vlan-architecture.html

Dynamic 802.1Q VLAN Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/vlan-dynamic-subscriber-access.html

Understanding Subscriber Interfaces; DHCP Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-understanding-interfaces-profiles.html

Demultiplexing Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-ip-demux.html

PPPoE Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-pppoe.html

MLPPP Support for LNS and PPPoE Subscribers: https://www.juniper.net/documentation/en_US/junos/topics/concept/mlppp-lns-ppoe-subscriber-mx-series-support-overview.html

ATM for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-atm-overview.html

Understanding Subscriber Sessions; RADIUS Server Options for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/aaa-radius-server-options-overview.html

Global RADIUS Options for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/radius-options-overview.html

RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework: https://www.juniper.net/documentation/en_US/junos/topics/concept/aaa-service-framework-radius-vsa-overview.html

Specifying the Authentication and Accounting Methods for Subscriber Access https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/aaa-subscriber-access-authentication-accounting-methods.html

RADIUS Accounting Statistics for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/aaa-radius-accounting-overview.html

Understanding Session Options for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/access-profile-session-options-overview.html

Domain Mapping Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-domain-maps.html

Using RADIUS Dynamic Requests for Subscriber Access Management: https://www.juniper.net/documentation/en_US/junos/topics/concept/aaa-radius-dynamic-requests-overview.html

Extended DHCP Local Server Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/dhcp-extended-dhcp-local-server-overview.html

Extended DHCP Relay Agent Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/dhcp-extended-dhcp-relay-overview.html

DHCP Relay Proxy Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/dhcp-extended-dhcp-relay-proxy-overview.html

Default Subscriber Service Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-default-subscriber-service-overview.html

Junos OS Enhanced Subscriber Management Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-enhanced-overview.html

Address-Assignment Pools Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-address-assignment-pools-overview.html

DNS Name Server Address Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/dns-name-server-address-overview.html

CLI-Activated Subscriber Services: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-cli-activated-subscriber-service.html

Subscriber Services with Multiple Instances Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-multi-instance-services-oview.html

Diameter Base Protocol Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/diameter-overview.html

Understanding CoS for Subscriber Access: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access.html

CoS for Aggregated Ethernet Subscriber Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-ae.html

CoS for PPPoE Subscriber Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-pppoe.html

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-dedicated-queues-trio.html

Bandwidth Management for Downstream Traffic in Edge Networks Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-downstream.html

Changing CoS Services Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-dynamic-upgrade.html

CoS for Interface Sets of Subscribers Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-sets.html

Understanding Firewall Filters; Understanding Dynamic Firewall Filters:https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-dynamic-firewall-filter-overview.html

Classic Filters Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-dynamic-firewall-classic-filters-overview.html

Parameterized Filters Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-par-filt-overview.html

Ascend-Data-Filter Policies for Subscriber Management Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-ascend-data-filters.html

Fast Update Filters Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-dynamic-firewall-fast-update-filters-overview.html

Unicast RPF in Dynamic Profiles for Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/unicast-mx-series-dynamic-profiles.html

Firewall Filters and Enhanced Network Services Mode Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/chassis-services-enhanced-mode-overview.html

Dynamic Service Sets Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-dynamic-service-sets-overview.html

Methods for Regulating Traffic by Applying Hierarchical Policers: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-hierarchical-policer-overview.html

Understanding Dynamic Multicast; Dynamic IGMP Configuration Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-address-igmp-dynamic-config-overview.html

Subscriber Management IGMP Model Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-igmp.html

Dynamic MLD Configuration Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-mld-dynamic-config-overview.html

Understanding Subscriber Secure Policy: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-secure-policy-overview.html

RADIUS-Initiated Subscriber Secure Policy Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-secure-policy-overview-radius.html

DTCP-Initiated Subscriber Secure Policy Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-secure-policy-overview-dtcp.html

Subscriber Secure Policy Support for IPv4 Multicast Traffic: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-subscriber-secure-policy-multicast.html

Understanding Hierarchical Class of Service: https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/cos/config-guide-hierarchical-cos.html

Hierarchical Class of Service Network Scenarios: https://www.juniper.net/documentation/en_US/junos/topics/concept/hierarchical-cos-use-case.html

Understanding Hierarchical Scheduling: https://www.juniper.net/documentation/en_US/junos/topics/concept/hierarchical-scheduler-terms-cos-config-guide.html

Hierarchical Schedulers and Traffic Control Profiles: https://www.juniper.net/documentation/en_US/junos/topics/concept/hierarchical-scheduler-and-tcps-cos-config-guide.html

Understanding Hierarchical Scheduling for MIC and MPC Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/hw-cos-trio-scheduling-hierarchical.html

Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-hierarchical-scheduler-on-aggregated-ethernet-interfaces-overview.html

Hierarchical Class of Service for Subscriber Management Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/hierarchical-cos-subscriber-overview.html

Understanding Hierarchical CoS for Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/hierarchical-scheduling-mic-mpc-understanding.html

Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-mpls-access-pseudowire-overview.html

CoS Configuration Overview for MPLS Pseudowire Subscriber Interfaces: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-mpls-access-pseudowire-configuration-overview.html

CoS for L2TP LAC Subscriber Interfaces Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-l2tp.html

CoS for L2TP LNS Inline Services Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-l2tp-lns.html

Hierarchical CoS Shaping-Rate Adjustments Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-adjustments.html

Shaping Rate Adjustments for Subscriber Local Loops Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/cos-subscriber-access-adjustments-subscriber-loops-understanding.html

CoS Traffic Shaping Attributes for Dynamic Interface Sets and Member Subscriber Sessions Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-cos-traffic-shaping-interface-sets-oview.html

Understanding Broadband Subscriber Management Wholesale Networks
Layer 2 and Layer 3 Wholesale Overview: https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/subscriber-access/subscriber-mgmt-wholesale.html

Subscriber to Logical System and Routing Instance Relationship: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-dhcp-layer3-wholesale-ls-ri.html

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-wholesale-vsa-config-overview.html

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-wholesale-topology.html

Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-pppoe-wholesale-topology.html

Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements: https://www.juniper.net/documentation/en_US/junos/topics/concept/subscriber-management-solution-layer2-wholesale-topology.html

Layer 2 Wholesale with ANCP-Triggered VLANs Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/ancp-layer2-bitstream-access-overview.html

Flat-File Accounting Overview: https://www.juniper.net/documentation/en_US/junos/topics/concept/flat-file-accounting-overview.html

The Network Design and Architecture Center

The Network Design and Architecture Center provides all the resources you need to design and deploy your network, all in one place.

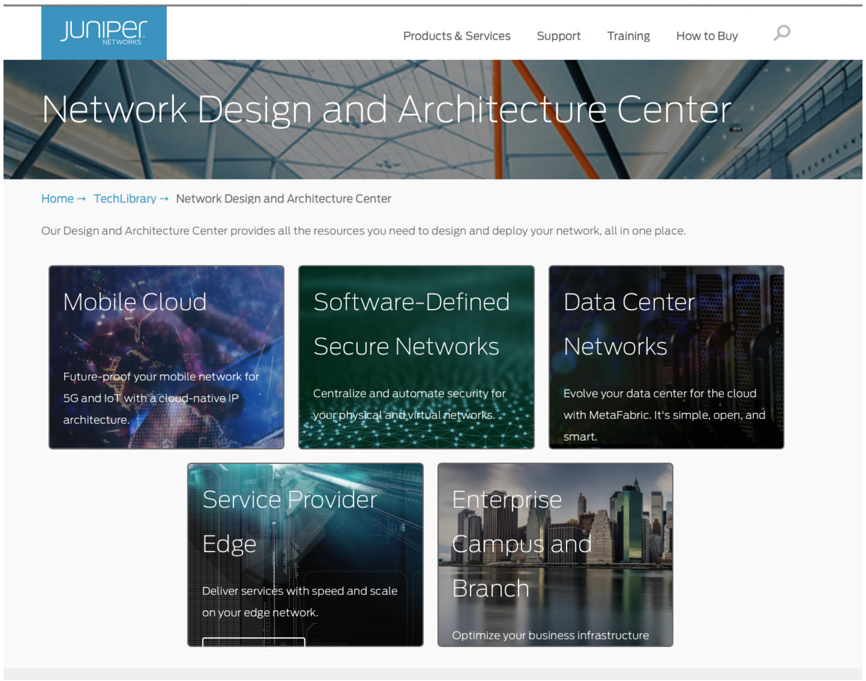

Network Design and Architecture Center

Home → TechLibrary → Network Design and Architecture Center

Our Design and Architecture Center provides all the resources you need to design and deploy your network, all in one place.

Mobile Cloud

Future-proof your mobile network for 5G and IoT with a cloud-native IP architecture.

Software-Defined Secure Networks

Centralize and automate security for your physical and virtual networks.

Data Center Networks

Evolve your data center for the cloud with MetaFabric. It's simple, open, and smart.

Service Provider Edge

Deliver services with speed and scale on your edge network.

Enterprise Campus and Branch

Optimize your business infrastructure

http://www.juniper.net/documentation/en_US/design-and-architecture/index.html