Junos® OS

Broadband Subscriber Access Protocols Feature Guide

Release
14.1

Modified: 2015-09-03

Copyright © 2016, Juniper Networks, Inc.

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Junos® OS Broadband Subscriber Access Protocols Feature Guide
14.1
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

 Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

List of Figures

Copyright © 2016, Juniper Networks, Inc.

List of Tables

 Copyright © 2016, Juniper Networks, Inc.

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

For the features described in this document, the following platforms are supported:

- MX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

Copyright © 2016, Juniper Networks, Inc.

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xxi defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xxi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** *community-ids* **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		static {
		route default {
		nexthop *address*;
		retain;
		}
		}
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

Copyright © 2016, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/

Copyright © 2016, Juniper Networks, Inc.

- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 1

Broadband Subscriber Access Network Overview

Subscriber Access Network Overview

A subscriber access environment can include various components, including 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 1 on page 26 shows an example of a basic subscriber access network.

Copyright © 2016, Juniper Networks, Inc.

Figure 1: Subscriber Access Network Example

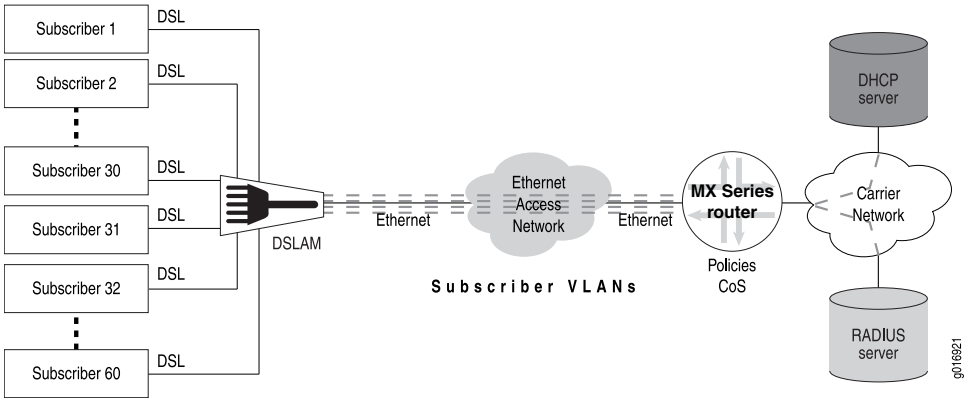

Related Documentation

- *Subscriber Management Overview*

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

Copyright © 2016, Juniper Networks, Inc.

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.

In choosing a MSAN type, refer to Figure 2 on page 27:

Figure 2: Choosing an MSAN Type

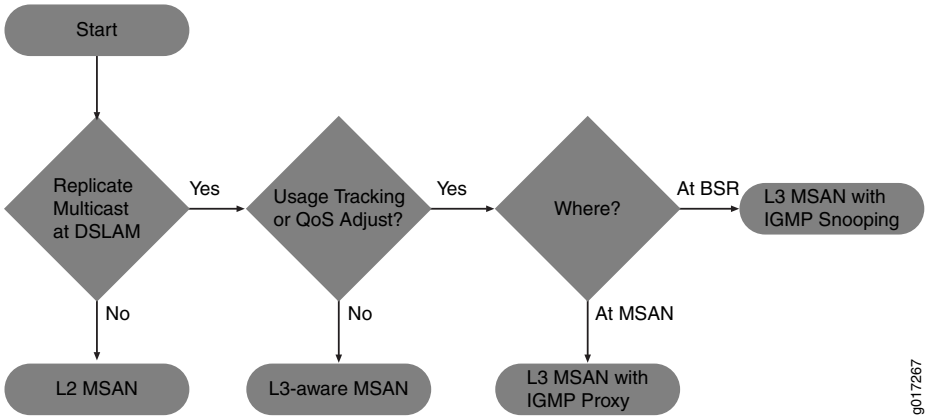

Related Documentation

- Ethernet MSAN Aggregation Options on page 27

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 3 on page 27 lists the possible MSAN aggregation methods and under what conditions they are used.

Table 3: Ethernet MSAN Aggregation Methods

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.

Copyright © 2016, Juniper Networks, Inc.

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.

Direct Connection

In the direct connection method, each MSAN has a point-to-point connection to the broadband services router. If an intermediate central office exists, traffic from multiple MSANs can be combined onto a single connection using wave-division multiplexing (WDM). You can also connect the MSAN to a video services router. However, this connection method requires that you use a Layer 3 MSAN that has the ability to determine which link to use when forwarding traffic.

When using the direct connection method, keep the following in mind:

- We recommend this approach when possible to simplify network management.
- Because multiple MSANs are used to connect to the services router, and Layer 3 MSANs generally require a higher equipment cost, this method is rarely used in a multiedge subscriber management model.
- Direct connection is typically used when most MSAN links are utilized less than 33 percent and there is little value in combining traffic from multiple MSANs.

Ethernet Aggregation Switch Connection

An Ethernet aggregation switch aggregates traffic from multiple downstream MSANs into a single connection to the services router (broadband services router or optional video services router).

When using the Ethernet aggregation switch connection method, keep the following in mind:

- Ethernet aggregation is typically used when most MSAN links are utilized over 33 percent or to aggregate traffic from lower speed MSANs (for example, 1 Gbps) to a higher speed connection to the services router (for example, 10 Gbps).
- You can use an MX Series router as an Ethernet aggregation switch. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Ring Aggregation Connection

In a ring topology, the remote MSAN that connects to subscribers is called the remote terminal (RT). This device can be located in the outside plant (OSP) or in a remote central office (CO). Traffic traverses the ring until it reaches the central office terminal (COT) at the head-end of the ring. The COT then connects directly to the services router (broadband services router or video services router).

Copyright © 2016, Juniper Networks, Inc.

NOTE: **The RT and COT must support the same ring resiliency protocol.**

You can use an MX Series router in an Ethernet ring aggregation topology. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Related Documentation

- Multiservice Access Node Overview on page 26

Broadband Access Service Delivery Options

Four primary delivery options exist today for delivering broadband network service. These options include the following:

- Digital Subscriber Line on page 29
- Active Ethernet on page 29
- Passive Optical Networking on page 30
- Hybrid Fiber Coaxial on page 30

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.

Copyright © 2016, Juniper Networks, Inc.

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.

 Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Broadband Delivery and FTTx on page 31

Broadband Delivery and FTTx

Many implementations use existing copper cabling to deliver signal to the premises, but fiber-optic cable connectivity is making its way closer to the subscriber. Most networks use a combination of both copper and fiber-optic cabling. The term *fiber to the x* (FTTx) describes how far into the network fiber-optic cabling runs before a switch to copper cabling takes place. Both PON and Active Ethernet can use fiber-optic portion of the network, while xDSL is typically used on the copper portion. This means that a single fiber-optic strand may support multiple copper-based subscribers.

Increasing the use of fiber in the network increases cost but it also increases network access speed to each subscriber.

The following terms are used to describe the termination point of fiber-optic cable in a network:

- Fiber to the Premises (FTTP), Fiber to the Home (FTTH), Fiber to the Business (FTTB)—Fiber extends all the way to the subscriber. PON is most common for residential access, although Active Ethernet can be efficiently used in dense areas such as apartment complexes. Active Ethernet is more common for delivering services to businesses.
- Fiber to the Curb (FTTC)—Fiber extends most of the way (typically, 500 feet/150 meters or less) to the subscriber. Existing copper is used for the remaining distance to the subscriber.
- Fiber to the Node/Neighborhood (FTTN)—Fiber extends to within a few thousand feet of the subscriber and converted to xDSL for the remaining distance to the subscriber.
- Fiber to the Exchange (FTTE)—A typical central office-based xDSL implementation in which fiber is used to deliver traffic to the central office and xDSL is used on the existing local loop.

Related Documentation

- Broadband Access Service Delivery Options on page 29

Copyright © 2016, Juniper Networks, Inc.

PART 1

Configuring the DHCP Access Network

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 2

Configuring Services for DHCP Subscribers

DHCP and Subscriber Management 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 and Subscriber Management Overview

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 and Subscriber Management Overview

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 and Subscriber Management 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. 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.

Related Documentation

- *Extended DHCP Local Server Overview*
- *Extended DHCP Relay Agent Overview*
- *DHCP Relay Proxy Overview*
- *Address-Assignment Pools Overview*

Subscriber Access Operation Flow Using DHCP Relay

The subscriber management feature requires that a subscriber (for example, a DHCP client) send a discover message to the router interface to initialize dynamic configuration of that interface.

Figure 3 on page 37 shows the flow of operations that occurs when the router is using DHCP relay to enable access for a subscriber.

 Copyright © 2016, Juniper Networks, Inc.

Figure 3: Subscriber Access Operation Flow

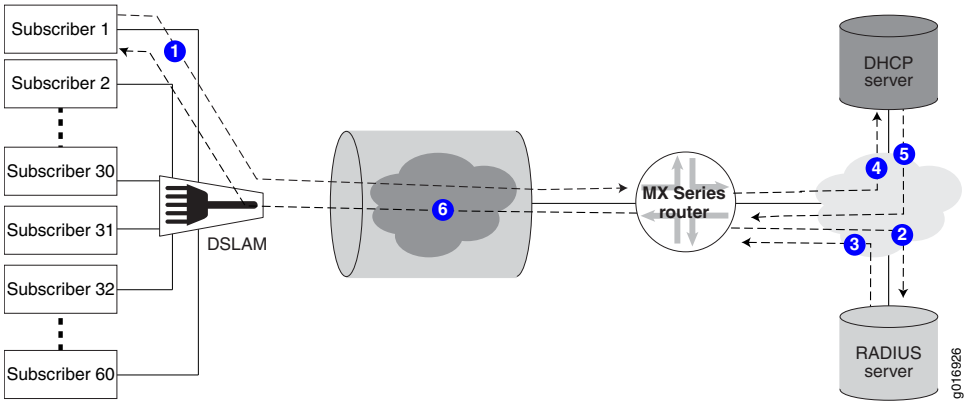

The following general sequence occurs during access configuration for a DHCP client:

1. The client issues a DHCP discover message.
2. The router issues an authorization request to the RADIUS server.
3. The RADIUS server issues an authorization response to the router.
4. The router passes the DHCP discover message through to the DHCP server.
5. The DHCP server issues an IP address for the client.
6. The router DHCP component sends an acknowledgement back to the client.

The subscriber now has access to the network and the authorized service.

Related Documentation

- *Subscriber Management Overview*
- *Configuring Subscriber Access*

Defining Various Levels of Services for DHCP Subscribers

This topic discusses how to create dynamic profiles to define various levels of service for DHCP clients.

Before you configure dynamic profiles for client services:

1. Create a basic dynamic profile.

   See *Configuring a Basic Dynamic Profile*.

2. Configure a dynamic profile that enables DHCP clients access to the network.

   See *Configuring Dynamic DHCP Client Access to a Multicast Network*

NOTE: **You can create a basic dynamic profile that contains both access configuration and some level of basic service.**

3. Ensure that the router is configured to enable communication between the client and the RADIUS server.

   See *Specifying the Authentication and Accounting Methods for Subscriber Access*.

4. Configure all RADIUS values that you want the profiles to use when validating DHCP clients.

   See *Configuring RADIUS Server Parameters for Subscriber Access*

To configure an initial client access dynamic profile:

1. Access the desired service profile.

   user@host# **set dynamic-profiles basic-service-profile**

2. (Optional) Define any IGMP protocols values as described for creating a basic access profile to combine a basic service with access in a profile.

   See *Configuring Dynamic DHCP Client Access to a Multicast Network*.

3. (Optional) Specify any filters for the interface.

   See *Dynamically Attaching Statically Created Filters for Any Interface Type*, *Dynamically Attaching Statically Created Filters for a Specific Interface Family Type*, or *Dynamically Attaching Filters Using RADIUS Variables*.

4. Define any CoS values for the service level you want this profile to configure on the interface.

Related Documentation

- *Configuring a Basic Dynamic Profile*
- *Dynamic Profiles Overview*

Example: Configuring a Tiered Service Profile for Subscriber Access

This example shows how to configure a tiered service profile for subscribers.

The profile contains three services:

- Gold—Subscribers that pay for this service are allocated 10M bandwidth for data, voice, and video services.
- Silver—Subscribers that pay for this service are allocated 5M bandwidth for data, voice, and video services.
- Bronze—Subscribers that pay for this service are allocated 1M bandwidth for the data service only.

Each subscriber is allocated a VLAN that is created statically. Subscribers log in using DHCP and authenticate using RADIUS. The subscribers can migrate from one service to another when they change subscriptions.

Copyright © 2016, Juniper Networks, Inc.

To configure a profile for a tiered service:

1. Configure the VLAN interfaces associated with each subscriber. Enable hierarchical scheduling for the interface.

```
interfaces {
  ge-2/0/0 {
    description subscribers;
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      vlan-tags outer 100 inner 100;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
    }
    unit 2 {
      family inet {
        vlan-tags outer 101 inner 101;
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
    }
    unit 3 {
      vlan-tags outer 102 inner 102;
      family inet {
        unnumbered-address lo0.0 preferred-source-address 100.0.0.1;
      }
    }
  }
}
```

2. Configure the static CoS parameters.

   In this example, each offering (video, voice, and data) is assigned a queue, and each service (Gold, Silver, and Bronze) is assigned a scheduler.

```
class-of-service {
  forwarding-classes {
    queue 0 data;
    queue 1 voice;
    queue 2 video;
  }
  scheduler-maps {
    bronze_service_smap {
      forwarding-class data scheduler data_sch;
    }
    silver_service_smap {
      forwarding-class data scheduler data_sch;
      forwarding-class voice scheduler silver_voice_sch;
      forwarding-class video scheduler silver_video_sch;
    }
    gold_service_smap {
      forwarding-class data scheduler data_sch;
      forwarding-class voice scheduler gold_voice_sch;
      forwarding-class video scheduler gold_video_sch;
    }
  }
```

```
    schedulers {
      data_sch {
        transmit-rate percent 20;
        buffer-size remainder;
        priority low;
      }
      silver_voice_sch {
        transmit-rate percent 30;
        buffer-size remainder;
        priority high;
      }
      silver_video_sch {
        transmit-rate percent 30;
        buffer-size remainder;
        priority medium;
      }
      gold_voice_sch {
        transmit-rate percent 40;
        buffer-size remainder;
        priority high;
      }
      gold_video_sch {
        transmit-rate percent 40;
        buffer-size remainder;
        priority medium;
      }
    }
  }
```

3. Configure the dynamic profile for the service.

   The scheduler maps configured for each service are referenced in the dynamic profile.

```
  dynamic-profiles {
    subscriber_profile {
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            family inet;
          }
        }
      }
      class-of-service {
        traffic-control-profiles {
          subscriber_tcp {
            scheduler-map $smap;
            shaping-rate $shaping-rate;
            guaranteed-rate $guaranteed-rate;
            delay-buffer-rate $delay-buffer-rate;
          }
        }
        interfaces {
          "$junos-interface-ifd-name" {
            unit "$junos-underlying-interface-unit" {
              output-traffic-control-profile subscriber_tcp;
            }
          }
```

 Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
```

4. Configure access for the subscribers.

   The DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You use DHCP relay to obtain configuration parameters, including an IP address, for subscribers. In this example, one DHCP server, address 100.20.42.1, can be used by subscribers.

   The DHCP relay configuration is attached to an active server group named service_provider_group.

   The subscribers are grouped together within the subscriber_group, and identifies characteristics such as authentication, username info, and the associated interfaces for the group members. In this example, it also identifies the active server group and the dynamic interface that is used by the subscribers in the group.

```
forwarding-options {
  dhcp-relay {
    server-group {
      service_provider_group {
        100.20.42.1;
      }
    }
    group subscriber_group {
      active-server-group service_provider_group;
      dynamic-profile subscriber_profile;
      interface ge-2/0/0.1;
      interface ge-2/0/0.2;
      interface ge-2/0/0.3;
    }
  }
}
```

Related Documentation

- For more information about configuring CoS for subscriber access, see *CoS for Subscriber Access Overview*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 3

Applying RADIUS Route Attributes to Subscribers or to Access Networks

Access and Access-Internal Routes for Subscriber Management

DHCP and PPP on the router use both access routes and access-internal routes to represent either the subscriber or the networks behind the attached router. An access route represents a network behind an attached router, and is set to a preference of 13. An access-internal route is a **/32** route that represents a directly attached subscriber, and is set to a preference of 12.

Access routes typically are used to apply the values of the RADIUS Framed-Route attribute [22] for IPv4 routes and the Framed-IPv6-Route attribute [99] for IPv6 routes. A framed route consists of a prefix that represents a public network behind the CPE, a next-hop gateway, and optional route attributes consisting of a combination of metric, preference, and tag. The only mandatory component of the framed route is the prefix. The next-hop gateway can be specified explicitly in the framed route. Alternatively, the absence of the gateway address implies address 0.0.0.0, which must resolve using the CPE's IP address. In either case, the convention is that the next-hop gateway is the CPE IP address.

You can configure a dynamic profile to use predefined variables to dynamically configure access routes using the values specified in the RADIUS attribute. To configure access routes include the **access** stanza at the **[edit dynamic-profiles *profile-name* routing-options]** hierarchy level. To configure access-internal routes, include the **access-internal** stanza at the same hierarchy level.

Consider the following rules for resolving the next-hop gateway to determine when each stanza is required:

- If the RADIUS framed route always specifies the next-hop gateway, only the **access** stanza is required in the dynamic profile. The **access-internal** stanza is not required.
- If the RADIUS framed route does not specify the next-hop gateway—as is more common—the variable representing the next-hop, $junos-framed-route-nexthop,

defaults to 0.0.0.0. This value implies that the CPE IP address is to be used. For this case, the **access-internal** stanza is required to resolve $junos-framed-route-nexthop to the CPE IP address (represented in the **access-internal** stanza by $junos-subscriber-ip-address) and the logical interface (represented as a qualified next-hop by $junos-interface-name).

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101
- *RADIUS IETF Attributes Supported by the AAA Service Framework*

Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management

You can dynamically configure access-internal routes. Configuring support for access-internal variables is optional, but it ensures that values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

DHCP subscriber interfaces require the qualified-next-hop to identify the interface and the MAC address.

To dynamically configure access-internal routes:

1. Specify that you want to configure the access-internal route.

   user@host# **edit dynamic-profiles** ***profile-name*** **routing-options**

2. Configure the IP address and the qualified next-hop address as variables.

   [edit dynamic-profiles *profile-name* routing-options]
   user@host# **edit** access-internal route **$junos-subscriber-ip-address** qualified-next-hop **$junos-interface-name**

   NOTE: the variable used for qualified-next-hop is $junos-interface-name.

Copyright © 2016, Juniper Networks, Inc.

3. Configure the MAC address for the qualified next-hop as a variable.

```
[edit dynamic-profiles profile-name routing-options access-internal route
  $junos-subscriber-ip-address qualified-next-hop $junos-underlying-interface]
user@host# set mac-address $junos-subscriber-mac-address
```

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 91

Configuring Dynamic Access Routes for Subscriber Management

You can dynamically configure access routes for DHCP and PPP subscribers based on the values specified in the following RADIUS attributes:

- For IPv4 access routes, use the variable, **$junos-framed-route-ip-address-prefix**. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, **$junos-framed-route-ipv6-address-prefix**. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].

To dynamically configure access routes:

1. Configure the route prefix for the access route as a variable.

   For IPv4:

   ```
   [edit dynamic-profiles profile-name routing-options]
   user@host# edit access route $junos-framed-route-ip-address-prefix
   ```

   For IPv6:

   ```
   [edit dynamic-profiles profile-name routing-options]
   user@host# edit access route $junos-framed-route-ipv6-address-prefix
   ```

2. Configure the next-hop address as a variable.

   For IPv4:

   ```
   [edit dynamic-profiles profile-name routing-options access route
     "$junos-framed-route-ip-address-prefix"]
   user@host# set next-hop $junos-framed-route-nexthop
   ```

   For IPv6:

   ```
   [edit dynamic-profiles profile-name routing-options access route
     "$junos-framed-route-ipv6-address-prefix"]
   user@host# set next-hop $junos-framed-route-ipv6-nexthop
   ```

3. Configure the metric as a variable.

   For IPv4:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set metric $junos-framed-route-cost
```

For IPv6:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set metric $junos-framed-route-ipv6-cost
```

4. Configure the preference as a variable (IPv4 only).

   For IPv4:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set preference $junos-framed-route-distance
```

   For IPv6:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set preference $junos-framed-route-ipv6-distance
```

5. Configure the tag as a variable (IPv4 only).

   IPv4:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set tag $junos-framed-route-tag
```

   IPv6:

```
[edit dynamic-profiles profile-name routing-options access route
  "$junos-framed-route-ip-address-prefix"]
user@host# set tag $junos-framed-route-ipv6-tag
```

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101
- Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 91
- *RADIUS IETF Attributes Supported by the AAA Service Framework*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 4

Suppressing DHCP Access, Access-Internal, and Destination Routes

Suppressing DHCP Access, Access-Internal, and Destination Routes

During the DHCP client binding operation, the DHCP process adds route information for the DHCP sessions by default. The DHCP process adds access-internal and destination routes for DHCPv4 sessions, and access-internal and access routes for DHCPv6 sessions. In some scenarios, you might want to override the default behavior and prevent DHCP from automatically installing the route information. For example, DHCP relay installs destination (host) routes by default—this action is required in certain configurations to enable address renewals from the DHCP server to work properly. However, the default installation of destination routes might cause a conflict when you configure DHCP relay with static subscriber interfaces. To avoid such configuration conflicts you can override the default behavior and prevent DHCP relay from installing the routes.

NOTE: **You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.**

You can configure both DHCP local server and DHCP relay agent to override the default route installation behavior, and you can specify the override for both DHCPv4 and DHCPv6 sessions. You can override the route installation globally or for named interface groups. For DHCPv4 you can override the installation of destination routes only or access-internal routes (the access-internal option prevents installation of both destination and access-internal routes). For DHCPv6 you can specify access routes, access-internal routes, or both.

Related Documentation

- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 48
- *Extended DHCP Local Server Overview*

- *DHCPv6 Local Server Overview*
- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*

Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default

You can configure both DHCP local server and DHCP relay agent to override the default installation of access, access-internal, and destination routes. For DHCPv4 you can override the installation of destination routes only or access-internal routes (the access-internal option prevents installation of both destination and access-internal routes). For DHCPv6 you can specify access routes, access-internal routes, or both. You can configure the override globally or for named interface groups.

NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.

NOTE: The **no-arp** statement is deprecated and the function is replaced by the **route-suppression** statement.

To configure route suppression and prevent DHCP from installing specific types of routes:

- For DHCP local server route suppression (for example, a global configuration):

  [edit system services dhcp-local-server]
  user@host# set route-suppression (DHCP Local Server and Relay Agent) access-internal

- For DHCP relay (for example, a group-specific configuration):

  [edit forwarding-options dhcp-relay group southeast]
  user@host# set route-suppression (DHCP Local Server and Relay Agent) destination

- For DHCPv6 local server (for example, a group-specific configuration):

  [edit system services dhcp-local-server group southern3]
  user@host# set dhcpv6 route-suppression (DHCP Local Server and Relay Agent) access access-internal

- For DHCPv6 relay (for example, a global configuration):

  [edit forwarding-options dhcp-relay]
  user@host# set dhcpv6 route-suppression (DHCP Local Server and Relay Agent) access

Related Documentation

- Suppressing DHCP Access, Access-Internal, and Destination Routes on page 47
- *Extended DHCP Local Server Overview*
- *DHCPv6 Local Server Overview*

Copyright © 2016, Juniper Networks, Inc.

- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 5

Providing Security in the DHCP Network

DHCP Snooping Support

DHCP snooping provides DHCP security on the router or switch by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. The router discards snooped packets by default. To enable normal processing of snooped packets, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

  In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.

Copyright © 2016, Juniper Networks, Inc.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

  In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

  Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

Related Documentation

- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 52
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 53
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 58
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 60
- Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent on page 62

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

You can configure how DHCP local server handles DHCP snooped packets. Depending on the configuration, DHCP local server either forwards or drops the snooped packets it receives.

Table 4 on page 53 indicates the action the router takes for DHCP local server snooped packets.

NOTE: Configured interfaces are those interfaces that have been configured with the group statement in the [edit system services dhcp-local-server] hierarchy. Non-configured interfaces are those that are in the logical system/routing instance but have not been configured by the group statement.

Copyright © 2016, Juniper Networks, Inc.

Table 4: Actions for DHCP Local Server Snooped Packets

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	forwarded	forwarded
configured-interfaces	forwarded	dropped
non-configured-interfaces	dropped	forwarded

To configure DHCP snooped packet forwarding for DHCP local server:

1. Specify that you want to configure DHCP local server.

   [edit]
   user@host# **edit system services dhcp-local-server**

2. Enable DHCP snooped packet forwarding for DHCP local server.

   [edit system services dhcp-local-server]
   user@host# **edit** forward-snooped-clients

3. Specify the interfaces that are supported for snooped packet forwarding.

   [edit system services dhcp-local-server forward-snooped-clients]
   user@host# **set (all-interfaces | configured-interfaces | non-configured-interfaces)**

For example, to configure DHCP local server to forward DHCP snooped packets on only configured interfaces:

```
[edit]
system {
  services {
    dhcp-local-server {
      forward-snooped-clients configured-interfaces;
    }
  }
}
```

Related Documentation

- DHCP Snooping Support on page 51

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes the first procedure, in which you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration.

The second procedure, which applies only to DHCPv4 relay agent, is described in "Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent" on page 58, and configures the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

You can enable or disable DHCP globally for DHCP relay, for a group of interfaces, or for a specific interface in a group.

By default, DHCP snooping is disabled for DHCP relay. To enable or disable DHCP snooping support globally:

1. Specify that you want to configure DHCP relay agent.
   - For DHCP relay agent:

     ```
     [edit]
     user@host# edit forwarding-options dhcp-relay
     ```

   - For DHCPv6 relay agent:

     ```
     [edit]
     user@host# edit forwarding-options dhcp-relay dhcpv6
     ```

2. Specify that you want to override the default configuration.
   - For DHCP relay agent:

     ```
     [edit forwarding-options dhcp-relay]
     user@host# edit overrides
     ```

   - For DHCPv6 relay agent:

     ```
     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# edit overrides
     ```

3. Enable or disable DHCP snooping support.
   - To enable DHCP snooping:
     - For DHCP relay agent:

       ```
       [edit forwarding-options dhcp-relay overrides]
       user@host# set allow-snooped-clients
       ```

     - For DHCPv6 relay agent:

       ```
       [edit forwarding-options dhcp-relay dhcpv6 overrides]
       user@host# set allow-snooped-clients
       ```

   - To disable DHCP snooping:
     - For DHCP relay agent:

       ```
       [edit forwarding-options dhcp-relay overrides]
       user@host# set no-allow-snooped-clients
       ```

     - For DHCPv6 relay agent:

       ```
       [edit forwarding-options dhcp-relay dhcpv6 overrides]
       user@host# set no-allow-snooped-clients
       ```


Copyright © 2016, Juniper Networks, Inc.

For example, to enable global DHCP snooping support :

```
forwarding-options {
  dhcp-relay {
    overrides {
      allow-snooped-clients;
    }
  }
}
```

To enable or disable DHCP snooping support for a group of interfaces:

1. Specify that you want to configure DHCP relay agent.
   - For DHCP relay agent:

     [edit]
     user@host# **edit forwarding-options** dhcp-relay

   - For DHCPv6 relay agent:

     [edit]
     user@host# **edit forwarding-options dhcp-relay dhcpv6**

2. Specify the named group.
   - For DHCP relay agent:

     [edit forwarding-options dhcp-relay]
     user@host# **edit group** *group-name*

   - For DHCPv6 relay agent:

     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# **edit group** *group-name*

3. Specify that you want to override the default configuration.
   - For DHCP relay agent:

     [edit forwarding-options dhcp-relay group *group-name*]
     user@host# **edit** overrides

   - For DHCPv6 relay agent:

     [edit forwarding-options dhcp-relay dhcpv6 group *group-name*]
     user@host# **edit** overrides

4. Enable or disable DHCP snooping support.
   - To enable DHCP snooping:
     - For DHCP relay agent:

       [edit forwarding-options dhcp-relay group *group-name* overrides]
       user@host# **set** allow-snooped-clients

     - For DHCPv6 relay agent:

       [edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides]
       user@host# **set** allow-snooped-clients

   - To disable DHCP snooping:

- For DHCP relay agent:

  [edit forwarding-options dhcp-relay group *group-name* overrides]
  user@host# **set** no-allow-snooped-clients

- For DHCPv6 relay agent:

  [edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides]
  user@host# **set** no-allow-snooped-clients

For example, to enable DHCP snooping support on all interfaces in group **boston**:

```
forwarding-options {
  dhcp-relay {
    group boston {
      overrides {
        allow-snooped-clients;
      }
    }
  }
}
```

To enable or disable DHCP snooping support on a specific interface:

1. Specify that you want to configure DHCP relay agent.

   - For DHCP relay agent:

     [edit]
     user@host# **edit forwarding-options** dhcp-relay

   - For DHCPv6 relay agent:

     [edit]
     user@host# **edit forwarding-options dhcp-relay dhcpv6**

2. Specify the named group containing the interface.

   - For DHCP relay agent:

     [edit forwarding-options dhcp-relay]
     user@host# **edit group** ***group-name***

   - For DHCPv6 relay agent:

     [edit forwarding-options dhcp-relay dhcpv6]
     user@host# **edit group** ***group-name***

3. Specify the interface for which you want to configure DHCP snooping.

   - For DHCP relay agent:

     [edit forwarding-options dhcp-relay group *group-name*]
     user@host# **edit interface** ***interface-name***

   - For DHCPv6 relay agent:

     [edit forwarding-options dhcp-relay dhcpv6 group *group-name*]
     user@host# **edit interface** ***interface-name***

4. Specify that you want to override the default configuration on the interface.

   - For DHCP relay agent:

Copyright © 2016, Juniper Networks, Inc.

[edit forwarding-options dhcp-relay group *group-name* interface *interface-name*]
user@host# **edit** overrides

- For DHCPv6 relay agent:

  [edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name*]
  user@host# **edit** overrides

5. Enable or disable DHCP snooping support.

   - To enable DHCP snooping:

     - For DHCP relay agent:

       [edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]
       user@host# **set** allow-snooped-clients

     - For DHCPv6 relay agent:

       [edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides]
       user@host# **set** allow-snooped-clients

   - To disable DHCP snooping:

     - For DHCP relay agent:

       [edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]
       user@host# **set** no-allow-snooped-clients

     - For DHCPv6 relay agent:

       [edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides]
       user@host# **set** no-allow-snooped-clients

For example, to disable DHCP snooping support on interface **ge-2/1/8.0** in group **boston**:

```
forwarding-options {
  dhcp-relay {
    group boston {
      interface ge-2/1/8.0 {
        overrides {
          no-allow-snooped-clients;
        }
      }
    }
  }
}
```

To enable DHCPv6 snooping support on interface **ge-3/2/1.1** in group **sunnyvale**:

```
forwarding-options {
  dhcp-relay {
    dhcpv6 {
      group sunnyvale {
        interface ge-3/2/1.1 {
          overrides {
```

```
                    allow-snooped-clients;
                }
            }
        }
    }
  }
}
```

Related Documentation

- DHCP Snooping Support on page 51
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 58
- Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 60
- *Overriding the Default DHCP Relay Configuration Settings*

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

You can configure how DHCP relay agent handles DHCP snooped packets. Depending on the configuration, DHCP relay agent either forwards or drops the snooped packets it receives.

DHCP relay uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes how you use the **forward-snooped-clients** statement to manage whether DHCP relay agent forwards or drops snooped packets, depending on the type of interface on which the packets are snooped. In the other part of the DHCP relay agent snooping configuration, which is described in "Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent" on page 53, you enable or disable the DHCP relay snooping feature.

Table 5 on page 59 shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the **allow-snooped-clients** statement. Table 6 on page 59 shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the **no-allow-snooped-clients** statement.

The router or switch also uses the configuration of the DHCP relay agent forwarding support to determine how to handle snooped BOOTREPLY packets. Table 7 on page 59 shows the action the router (or switch) takes for the snooped BOOTREPLY packets.

NOTE: Configured interfaces have been configured with the **group** statement in the **[edit forwarding-options dhcp-relay]** hierarchy. Non-configured interfaces are in the logical system/routing instance but have not been configured by the **group** statement.

Copyright © 2016, Juniper Networks, Inc.

Table 5: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	snooped packets result in subscriber (DHCP client) creation	dropped
all-interfaces	forwarded	forwarded
configured-interfaces	forwarded	dropped
non-configured-interfaces	snooped packets result in subscriber (DHCP client) creation	forwarded

Table 6: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	dropped	forwarded
configured-interfaces	dropped	dropped
non-configured-interfaces	dropped	forwarded

Table 7: Actions for Snooped BOOTREPLY Packets

forward-snooped-clients Configuration	Action
forward-snooped-clients not configured	snooped **BOOTREPLY** packets dropped if client is not found
forward-snooped-clients all configurations	snooped **BOOTREPLY** packets forwarded if client is not found

To configure DHCP snooped packet forwarding and BOOTREPLY snooped packet forwarding for DHCP relay agent:

1. Specify that you want to configure DHCP relay agent.

   [edit]
   user@host# **edit forwarding-options** dhcp-relay

2. Enable DHCP snooped packet forwarding.

   [edit forwarding-options dhcp-relay]
   user@host# **edit** forward-snooped-clients

3. Specify the interfaces that are supported for snooped packet forwarding.

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP relay agent to forward DHCP snooped packets on only configured interfaces:

```
[edit]
forwarding-options {
  dhcp-relay {
    forward-snooped-clients configured-interfaces;
  }
}
```

Related Documentation

- DHCP Snooping Support on page 51
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 53

Example: Configuring DHCP Snooping Support for DHCP Relay Agent

This example shows how to configure DHCP snooping support for DHCP relay agent.

Requirements

- Configure DHCP relay agent. See *Extended DHCP Relay Agent Overview*.

Overview

In this example, you configure DHCP snooping support for DHCP relay agent by completing the following operations:

- Override the default DHCP snooping configuration and enable DHCP snooping support for the interfaces in group **frankfurt**.
- Configure DHCP relay agent to forward snooped packets to only configured interfaces.

NOTE: **By default, DHCP snooping is disabled globally.**

Configuration

Step-by-Step Procedure

To configure DHCP relay support for DHCP snooping:

1. Specify that you want to configure DHCP relay agent.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```


Copyright © 2016, Juniper Networks, Inc.

2. Specify the named group of interfaces on which DHCP snooping is supported.

   [edit forwarding-options dhcp-relay]
   user@host# **edit group frankfurt**

3. Specify the interfaces that you want to include in the group. DHCP relay agent considers these as the configured interfaces when determining whether to forward or drop traffic.

   [edit forwarding-options dhcp-relay group frankfurt]
   user@host# **set interface fe-1/0/1.3 upto fe-1/0/1.9**

4. Specify that you want to override the default configuration for the group.

   [edit forwarding-options dhcp-relay group frankfurt]
   user@host# **edit** overrides

5. Enable DHCP snooping support for the group.

   [edit forwarding-options dhcp-relay group frankfurt overrides]
   user@host# **set** allow-snooped-clients

6. Return to the **[edit forwarding-options dhcp-relay]** hierarchy level to configure the forwarding action and specify that DHCP relay agent forward snooped packets on only configured interfaces:

   [edit forwarding-options dhcp-relay group frankfurt overrides]
   user@host# **up 2**

7. Enable DHCP snooped packet forwarding for DHCP relay agent.

   [edit forwarding-options dhcp-relay]
   user@host# **edit** forward-snooped-clients

8. Specify that snooped packets are forwarded on only configured interfaces (the interfaces in group **frankfurt**).

   [edit forwarding-options dhcp-relay forward-snooped-clients]
   user@host# **set configured-interfaces**

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group frankfurt.

```
[edit]
user@hostshow forwarding-options
dhcp-relay {
   forward-snooped-clients configured-interfaces;
   group frankfurt {
      overrides {
         allow-snooped-clients;
      }
      interface fe-1/0/1.3 {
         upto fe-1/0/1.9;
      }
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- DHCP Snooping Support on page 51
- *Configuring DHCP Snooping for DHCP Relay Agent*

Example: Enabling DHCP Snooping Support for DHCPv6 Relay Agent

Snooping support for DHCPv6 relay agent is disabled on the router by default. This example shows how to override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCPv6 snooping for a named group of interfaces and for a specific interface within a different named group.

NOTE: **You can also enable DHCPv6 snooping support globally by using the allow-snooped-clients statement at the [edit forwarding-options dhcp-relay dhcpv6 overrides] hierarchy level.**

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers

Before you begin:

- Configure DHCPv6 relay agent.

  See *DHCPv6 Relay Agent Overview*

- Configure named DHCPv6 relay agent interface groups to which you want to apply a common DHCP configuration.

  See *Grouping Interfaces with Common DHCP Configurations*

Overview

In this example, you override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCP snooping for both of the following:

- All of the interfaces in the group named **boston**
- Interface **ge-3/2/1.1** in the group named **sunnyvale**

Copyright © 2016, Juniper Networks, Inc.

Configuration

To override the default DHCPv6 relay agent snooping configuration to explicitly enable DHCPv6 snooping for a named group of interfaces and for a specific interface within a named group, perform these tasks:

- Enabling DHCPv6 Snooping Support for a Named Group of Interfaces on page 63
- Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group on page 64

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set forwarding-options dhcp-relay dhcpv6 group boston overrides allow-snooped-clients
set forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1 overrides
  allow-snooped-clients
```

Enabling DHCPv6 Snooping Support for a Named Group of Interfaces

Step-by-Step Procedure

To enable DHCPv6 snooping support for a named group of interfaces:

1. Specify that you want to configure DHCPv6 relay agent.

   ```
   [edit]
   user@host# edit forwarding-options dhcp-relay dhcpv6
   ```

2. Specify the named group of interfaces for which you want to enable DHCPv6 snooping.

   ```
   [edit forwarding-options dhcp-relay dhcpv6]
   user@host# edit group boston
   ```

3. Specify that you want to override the default DHCPv6 configuration for the interfaces in that group.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group boston]
   user@host# edit overrides
   ```

4. Enable DHCPv6 snooping support for all interfaces in group **boston**.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group boston overrides]
   user@host# set allow-snooped-clients
   ```

Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options dhcp-relay]
user@host# show
dhcpv6 {
  group boston {
    overrides {
```

```
            allow-snooped-clients;
        }
    }
}
```

If you are done configuring the router, enter **commit** from configuration mode.

Enabling DHCPv6 Snooping Support for a Specific Interface in a Named Group

Step-by-Step Procedure

To enable DHCPv6 snooping support for a specific interface within a named group of interfaces:

1. Return to the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level to specify that you want to configure DHCPv6 relay agent.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group boston overrides]
   user@host# up 2
   ```

2. Specify the named group containing the interface.

   ```
   [edit forwarding-options dhcp-relay dhcpv6]
   user@host# edit group sunnyvale
   ```

3. Specify the interface in group **sunnyvale** for which you want to enable DHCPv6 snooping.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group sunnyvale]
   user@host# edit interface ge-3/2/1.1
   ```

4. Specify that you want to override the default DHCPv6 configuration for interface **ge-3/2/1.1** in group **sunnyvale**.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1]
   user@host# edit overrides
   ```

5. Enable DHCPv6 snooping support for interface **ge-3/2/1.1** in group **sunnyvale**.

   ```
   [edit forwarding-options dhcp-relay dhcpv6 group sunnyvale interface ge-3/2/1.1
     overrides]
   user@host# set allow-snooped-clients
   ```

Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options dhcp-relay]
user@host# show
dhcpv6 {
  group boston {
    overrides {
      allow-snooped-clients;
    }
  }
  group sunnyvale {
    interface ge-3/2/1.1 {
      overrides {
        allow-snooped-clients;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

If you are done configuring the router, enter **commit** from configuration mode.

Verification

To verify the DHCPv6 configuration in a multi-relay topology, perform this task:

- Verifying the Address Bindings for DHCPv6 Relay Agent Clients on page 65

Verifying the Address Bindings for DHCPv6 Relay Agent Clients

Purpose Verify the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Action Display detailed information about address bindings for DHCPv6 relay agent clients.

```
user@host > show dhcpv6 relay binding detail

Session Id:  13
     Client IPv6 Prefix:                    3000:0:0:8001::5/128
     Client DUID:                           LL0x1-00:00:65:03:01:02
     State:                                 BOUND(DHCPV6_RELAY_STATE_BOUND)
     Lease Expires:                         2011-11-21 06:14:50 PST
     Lease Expires in:                      293 seconds
     Lease Start:                           2011-11-21 06:09:50 PST
     Incoming Client Interface:             ge-3/2/1.1
     Server Address:                        unknown
     Next Hop Server Facing Relay:          4000::2
     Server Interface:                      none
     Client Id Length:                      10
     Client Id:                             /0x00030001/0x00006503/0x0102
```

Meaning The **Server Address** field in the **show dhcpv6 relay binding detail** command output typically displays the IP address of the DHCPv6 server. In this example, the value **unknown** in the **Server Address** field indicates that this is a multi-relay topology in which the DHCPv6 relay agent is not directly adjacent to the DHCPv6 server, and does not detect the IP address of the server.

In that case, the output instead includes the **Next Hop Server Facing Relay** field, which displays the next-hop address in the direction of the DHCPv6 server.

Related Documentation

- *DHCPv6 Relay Agent Overview*
- DHCP Snooping Support on page 51
- *Grouping Interfaces with Common DHCP Configurations*
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 53

Preventing DHCP Spoofing

A problem that sometimes occurs with DHCP is *DHCP spoofing*. In DHCP spoofing, an untrusted client floods a network with DHCP messages. Often these attacks utilize source IP address spoofing to conceal the true source of the attack.

DHCP snooping helps prevent DHCP spoofing by copying DHCP messages to the control plane and using the information in the packets to create anti-spoofing filters. The anti-spoofing filters bind a client's MAC address to its DHCP-assigned IP address and use this information to filter spoofed DHCP messages. In a typical topology, a carrier edge router (in this function also referred to as the broadband network gateway [BNG]) connects the DHCP server and the MX Series router (or broadband services aggregator [BSA]) performing the snooping. The MX Series router connects to the client and the BNG.

To configure DHCP snooping, you include the appropriate interfaces within a DHCP group. You can configure DHCP snooping for VPLS environments and bridge domains.

- In a VPLS environment, DHCP requests are forwarded over pseudowires. You configure DHCP snooping over VPLS at the **[edit routing-instances routing-instance-name]** hierarchy level.
- In bridge domains, DHCP snooping works on a per learning bridge basis. Each learning domain must have an upstream interface configured. This interface acts as the flood port for DHCP requests coming from the client side. DHCP requests are forwarded across learning domains in a bridge domain. You configure DHCP snooping on bridge domains at the **[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]** hierarchy level.

To configure DHCP relay to prevent DHCP spoofing:

1. Access the appropriate hierarchy for either a VPLS or bridge domain configuration.

   user@host# edit routing-instances blue

2. Specify that you want to configure DHCP relay.

   [edit routing-instances blue]
   user@host# edit forwarding-options dhcp-relay

3. Create the group and assign a name.

   [edit routing-instances blue forwarding-options dhcp-relay]
   user@host# edit group svl-10

4. Specify the names of one or more interfaces. DHCP will trust only the MAC addresses learned on the specified interfaces.

   [edit routing-instances blue forwarding-options dhcp-relay group svl-10]
   user@host# set interface fe-1/0/1.1
   user@host# set interface fe-1/0/1.2

Copyright © 2016, Juniper Networks, Inc.

NOTE: You can explicitly enable and disable interface support for DHCP snooped clients. See "Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent" on page 53.

Related Documentation

- *Extended DHCP Relay Agent Overview*
- For examples of DHCP snooping, see the *JUNOS MX Series Ethernet Services Routers Solutions Guide*.
- Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 53

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 6

Distinguishing Between Duplicate DHCP IPv4 Subscribers on the Same Subnet

DHCP Duplicate Client In Subnet Overview

In some network environments, client IDs and hardware addresses (MAC addresses) might not be unique, resulting in duplicate clients. A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same hardware address as an existing DHCP client—the existing client and the new client cannot exist simultaneously, unless you have configured the optional duplicate client support.

By default, DHCP local server and DHCP relay agent use the subnet information to differentiate between duplicate clients. However, in some cases, this level of differentiation is not adequate. For example, when multiple subinterfaces share the same underlying loopback interface with the same preferred source address, the interfaces appear to be on the same subnet.

You can enable support for duplicate clients in a subnet by configuring DHCP to use additional information to uniquely identify clients—the additional information is either the client incoming interface or the option 82 information in the DHCP packets. Using the option 82 information provides the following important benefits:

- You can configure DHCP relay to preserve and use the remotely created option 82.
- DHCP local server can support an environment in which an aggregation device is present between the client and the DHCP server.

When configured to support duplicate clients in the subnet, DHCP uses the following information to distinguish between the duplicate clients:

- The subnet on which the client resides

Copyright © 2016, Juniper Networks, Inc.

- The client ID or hardware address
- The duplicate clients option you configure—either the client incoming interface or the option 82 information in the client's incoming DHCP packets

Related Documentation

- Guidelines for Configuring Support for DHCP Duplicate Clients on page 70
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Their Incoming Interfaces on page 72
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Option 82 Information on page 70

Guidelines for Configuring Support for DHCP Duplicate Clients

When configuring DHCP duplicate client support, consider the following guidelines:

- The DHCP duplicate client support feature is used for DHCPv4 clients only.
- If you want to preserve the remotely-created option 82 information, use the **option 82** option with the **duplicate-clients-in-subnet** statement to distinguish between duplicate clients. If there is no remotely created option 82 in the incoming DHCP packets, the router locally creates the option 82 information.
- If you want to use the locally-created option-82, use the **incoming-interface** option with the **duplicate-clients-in-subnet** statement to distinguish between duplicate clients.
- DHCP relay agent and DHCP local server in the same routing instance must have the same the **duplicate-clients-in-subnet** configuration.
- For the Layer 3 wholesale model:
  - The wholesaler and retailer logical system/routing instances must have the same **duplicate-clients-in-subnet** statement configuration.
  - For DHCP relay, the wholesaler and the retailer routing contexts must both have the **relay-option-82** statement configured with the Agent Circuit ID suboption (suboption 1) in option 82.

Related Documentation

- DHCP Duplicate Client In Subnet Overview on page 69
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Their Incoming Interfaces on page 72
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Option 82 Information on page 70

Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Option 82 Information

Duplicate clients occur when two clients in a subnet have the same hardware address or the same client ID.

Copyright © 2016, Juniper Networks, Inc.

The following two procedures describe how to configure the router to use the option 82 information in the incoming packets to differentiate between duplicate clients. The first procedure describes the configuration for DHCP relay agent The second procedure is for DHCP local server.

To configure the DHCP relay agent to differentiate between duplicate clients based on option 82 information:

1. Specify that you want to configure DHCP relay agent.

   [edit forwarding-options]
   user@host# **edit** dhcp-relay

2. Configure DHCP relay to insert option 82 information if there is no remotely created option 82. Use the default setting, which inserts the interface ID rather than the optional interface description.

   [edit forwarding-options dhcp-relay]
   user@host# **set** relay-option-82 **circuit-id**

3. Configure the router to always accept DHCP client packets that contain option 82 information.

   [edit forwarding-options dhcp-relay]
   user@host# **set** overrides **trust-option-82**

   NOTE: The **trust-option-82** statement must always be enabled so the router can process incoming DHCP client packets that contain option 82 information when the packets have a gateway IP address (giaddr) of 0 (zero).

4. Configure DHCP relay to use the remotely created option 82 information to distinguish between duplicate clients. If there is no remotely created option 82 in the traffic, the router locally creates the option 82 information.

   [edit forwarding-options dhcp-relay]
   user@host# **set** duplicate-clients-in-subnet **option-82**

   NOTE: Make sure that the always-write-option-82 statement is *not* enabled, as the statement will overwrite the remotely created option 82.

To configure the DHCP local server to differentiate between duplicate clients based on the option 82 information:

1. Specify that you want to configure DHCP local server.

   [edit system services]
   user@host# **edit dhcp-local-server**

2. Configure the duplicate client support with the **option-82** option.

   [edit system services dhcp-local-server]
   user@host# **set** duplicate-clients-in-subnet option-82

Related Documentation

- DHCP Duplicate Client In Subnet Overview on page 69
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 70
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Their Incoming Interfaces on page 72

Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Their Incoming Interfaces

Duplicate clients occur when two clients in a subnet have the same hardware address or the same client ID.

The following two procedures describe how to configure the router to use the clients' incoming interface to differentiate between duplicate clients. The first procedure describes the configuration for DHCP relay agent: the second procedure is for DHCP local server.

To configure the DHCP relay agent to differentiate between duplicate clients based on the client incoming interface:

1. Specify that you want to configure DHCP relay agent.

   ```
   [edit forwarding-options]
   user@host# edit dhcp-relay
   ```

2. Configure the duplicate client support with the **incoming-interface** option.

   ```
   [edit forwarding-options dhcp-relay]
   user@host# set duplicate-clients-in-subnet incoming-interface
   ```

3. Configure DHCP relay to insert option 82 information if the information is not specified remotely. Use the default setting, which inserts the interface ID rather than the optional interface description.

   ```
   [edit forwarding-options dhcp-relay]
   user@host# set relay-option-82 circuit-id
   ```

4. Configure the router to overwrite any remotely supplied option 82 information in incoming packets.

   ```
   [edit forwarding-options dhcp-relay]
   user@host# set overrides always-write-option-82
   ```

5. Configure the router to always accept DHCP client packets that contain option 82 information.

   ```
   [edit forwarding-options dhcp-relay]
   user@host# set overrides trust-option-82
   ```

   NOTE: The *trust-option-82* statement must always be enabled so the router can process incoming DHCP client packets that contain option 82 information when the packets have a gateway IP address (giaddr) of 0 (zero).

Copyright © 2016, Juniper Networks, Inc.

To configure the DHCP local server to differentiate between duplicate clients based on the client incoming interface:

1. Specify that you want to configure DHCP local server.

   ```
   [edit system services]
   user@host# edit dhcp-local-server
   ```

2. Configure the duplicate client support with the **incoming-interface** option.

   ```
   [edit system services dhcp-local-server]
   user@host# set duplicate-clients-in-subnet incoming-interface
   ```

Related Documentation

- DHCP Duplicate Client In Subnet Overview on page 69
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 70
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Option 82 Information on page 70

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 7

Configuring High Availability in the DHCP Access Network

DHCP Liveness Detection Overview

Unlike PPP, DHCP does not define a native keepalive mechanism as part of either the DHCPv4 or DHCPv6 protocols. Without a keepalive mechanism, DHCP local server, DHCP relay, or DHCP relay proxy is unable to quickly detect if it has lost connectivity with a subscriber or a DHCP client; and it must rely on standard DHCP subscriber session or DHCP client session termination messages.

DHCP clients often do not send DHCP release messages prior to exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often considered insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection protocol, if a given subscriber (or client) fails to respond to a

Copyright © 2016, Juniper Networks, Inc.

configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

DHCP liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

Using DHCP liveness detection, IP sessions are acted upon as soon as liveness detection checks fail. This faster response time serves to:

- Provide more accurate time-based accounting of subscriber (or DHCP client) sessions.
- Better preserve router (switch) resources.
- Help to reduce the window of vulnerability to some security attacks.

Examples of liveness detection protocols include Bidirectional Forwarding Detection (BFD) for both DHCPv4 and DHCPv6 subscribers, IPv4 Address Resolution Protocol (ARP) for DHCPv4 subscribers, and IPv6 Neighbor Unreachability Detection for DHCPv6 subscribers.

NOTE: **Only BFD for DHCPv4 and DHCPv6 liveness detection is supported.**

When configuring BFD liveness detection, keep the following in mind:

- You can configure DHCPv4 and DHCPv6 liveness detection either globally or per DHCPv4 or DHCPv6 group.
- DHCPv4 or DHCPv6 subscriber access clients that do not support BFD are not affected by the liveness detection configuration. These clients can continue to access the network (once validated) even if BFD liveness detection is enabled on the router (or switch).
- When configured, DHCPv4 or DHCPv6 initiates liveness detection checks for relevant clients (that is, clients that support BFD) when those clients enter a bound state.
- After protocol-specific messages are initiated for a BFD client, they are periodically sent to the subscriber (or client) IP address of the client and responses to those liveness detection requests are expected within a configured amount of time.
- If liveness detection responses are not received from clients that support BFD within the configured amount of time for a configured number of consecutive attempts, the liveness detection check is deemed to have failed and a configured failure action is implemented.

Related Documentation

- Configuring Detection of DHCP Local Server Client Connectivity on page 81
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 77

Copyright © 2016, Juniper Networks, Inc.

Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity

Liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.
   - For DHCP global configuration:

     [edit forwarding-options dhcp-relay]
     user@host# **edit** liveness-detection

   - For DHCP group configuration:

     [edit forwarding-options dhcp-relay group *group-name*]
     user@host# **edit** liveness-detection

   NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the liveness-detection statement, and any subsequent configuration statements, at the **[edit forwarding-options dhcp-relay dhcpv6]** or **[edit forwarding-options dhcp-relay dhcpv6 group *group-name*** hierarchy level.

2. (Optional) Specify that you want to use DHCP relay proxy mode.

   [edit forwarding-options dhcp-relay group *group-name*]
   user@host# **set** overrides proxy-mode

3. Specify that you want to configure the liveness detection method.
   - For DHCP global configuration:

     [edit forwarding-options dhcp-relay liveness-detection]
     user@host# **edit** method

   - For DHCP group configuration:

     [edit forwarding-options dhcp-relay group *group-name* liveness-detection]
     user@host# **edit** method

4. Specify the liveness detection method that you want DHCP to use.

   NOTE: The only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

   - For DHCP global configuration:

     [edit forwarding-options dhcp-relay liveness-detection method]

```
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method]
user@host# edit bfd
```

5. Configure the liveness detection method as desired.

   See "Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients" on page 78 for an example of how to globally configure DHCP relay liveness detection.

6. Configure the action the router takes when a liveness detection failure occurs.

   - For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit failure-action action
```

   - For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]
user@host# edit failure-action action
```

Related Documentation

- *Extended DHCP Relay Agent Overview*
- DHCP Liveness Detection Overview on page 75
- Configuring Detection of DHCP Local Server Client Connectivity on page 81
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

This example shows how to configure liveness detection for DHCP relay agent subscribers using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

Requirements

- Juniper Networks MX Series routers.
- Configure DHCP relay agent. See *Extended DHCP Relay Agent Overview*.

Copyright © 2016, Juniper Networks, Inc.

Overview

In this example, you configure liveness detection for DHCP relay agent subscribers by completing the following operations:

1. Enable liveness detection globally for DHCP relay subscribers.
2. Specify BFD as the liveness detection method for all dynamically created DHCP relay subscribers.
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router takes when a liveness detection failure occurs.

NOTE: This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the liveness-detection statement, and any subsequent configuration statements, at the **[edit forwarding-options dhcp-relay dhcpv6]** or **[edit forwarding-options dhcp-relay dhcpv6 group *group-name*]** hierarchy level.

Configuration

Step-by-Step Procedure

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.

   [edit forwarding-options dhcp-relay]
   user@host# edit liveness-detection

2. Specify that you want to configure the liveness detection method.

   [edit forwarding-options dhcp-relay liveness-detection]
   user@host# edit method

3. Specify BFD as the liveness detection method that you want DHCP to use.

   [edit forwarding-options dhcp-relay liveness-detection method]
   user@host# edit bfd

4. Configure the detection time threshold (in milliseconds) at which a trap is produced.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# set detection-time threshold **50000**

5. Configure the time (in milliseconds) for which BFD holds a session up notification.

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# set holddown-interval **50**

6. Configure the BFD minimum transmit and receive interval (in milliseconds).

   [edit forwarding-options dhcp-relay liveness-detection method bfd]
   user@host# set minimum-interval **45000**

7. Configure the minimum receive interval (in milliseconds).

   [edit forwarding-options dhcp-relay liveness-detection method bfd]

Copyright © 2016, Juniper Networks, Inc.

```
user@host# set minimum-receive-interval 60000
```

8. Configure a multiplier value for the detection time.

```
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set multiplier 100
```

9. Disable the ability for BFD interval timers to change or adapt to network situations.

```
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set no-adaptation
```

10. Configure the BFD session mode.

```
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set session-mode automatic
```

11. Configure the threshold and minimum interval for the BFD transmit interval.

```
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set transmit-interval threshold 60000 minimum-interval 45000
```

12. Configure the BFD protocol version you want to detect.

```
[edit forwarding-options dhcp-relay liveness-detection method bfd]
user@host# set version automatic
```

13. Configure the action the router takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.

```
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit failure-action action
```

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group frankfurt.

```
[edit]
user@host# show forwarding-options
dhcp-relay {
  liveness-detection {
    failure-action clear-binding-if-interface-up;
    method {
      bfd {
        version automatic;
        minimum-interval 45000;
        minimum-receive-interval 60000;
        multiplier 100;
        no-adaptation;
        transmit-interval {
          minimum-interval 45000;
          threshold 60000;
        }
        detection-time {
          threshold 50000;
        }
        session-mode automatic;
```

 Copyright © 2016, Juniper Networks, Inc.

```
            holddown-interval 50;
        }
    }
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- *Extended DHCP Relay Agent Overview*
- DHCP Liveness Detection Overview on page 75
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 77

Configuring Detection of DHCP Local Server Client Connectivity

Liveness detection for DHCP subscriber IP sessions or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

NOTE: You can also configure DHCP liveness detection for DHCP relay.

To configure liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.
   - For DHCP global configuration:

     [edit system services dhcp-local-server]
     user@host# **edit** liveness-detection

   - For DHCP group configuration:

     [edit system services dhcp-local-server group *group-name*]
     user@host# **edit** liveness-detection

     NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the liveness-detection statement, and any subsequent configuration statements, at the [edit system services dhcp-local-server dhcpv6] or [edit system services dhcp-local-server dhcpv6 group *group-name*] hierarchy level.

2. Specify that you want to configure the liveness detection method.
   - For DHCP global configuration:

     [edit system services dhcp-local-server liveness-detection]

```
user@host# edit method
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit method
```

3. Specify the liveness detection method that you want DHCP to use.

> NOTE: The only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection method]
user@host# edit bfd
```

4. Configure the liveness detection method as desired.

   See "Example: Configuring Group Liveness Detection for DHCP Local Server Clients" on page 83 for an example of how to configure DHCPv4 groups for DHCP local server liveness detection.

5. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit failure-action action
```

Related Documentation

- DHCP Liveness Detection Overview on page 75
- *Extended DHCP Local Server Overview*
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 77
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

Copyright © 2016, Juniper Networks, Inc.

Example: Configuring Group Liveness Detection for DHCP Local Server Clients

This example shows how to configure group liveness detection for DHCP local server subscribers or DHCP clients using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

Requirements

- Juniper Networks MX Series routers
- Juniper Networks EX Series switches
- Configure DHCP local server. See *Extended DHCP Local Server Overview*.

Overview

In this example, you configure group liveness detection for DHCP local server subscribers (clients) by completing the following operations:

1. Enable liveness detection for DHCP local server subscriber (or DHCP client) groups.
2. Specify BFD as the liveness detection method for all dynamically created DHCP local server subscribers (clients).
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router (switch) takes when a liveness detection failure occurs.

NOTE: This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the liveness-detection statement, and any subsequent configuration statements, at the **[edit system services dhcp-local-server dhcpv6]** or **[edit system services dhcp-local-server dhcpv6 group *group-name*]** hierarchy level.

Configuration

Step-by-Step Procedure

To configure group liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

```
[edit system services dhcp-local-server ]
user@host# edit liveness-detection
```

2. Specify that you want to configure liveness detection for a specific DHCP local server group.

```
[edit system services dhcp-local-server liveness-detection]
```

```
user@host# edit group local_group_1
```

3. Specify that you want to configure the liveness detection method.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit method
```

4. Specify BFD as the liveness detection method that you want DHCP to use.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method]
user@host# edit bfd
```

5. Configure the detection time threshold (in milliseconds) at which a trap is produced.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set detection-time threshold 30000
```

6. Configure the time (in milliseconds) for which BFD holds a session up notification.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set holddown-interval 50
```

7. Configure the BFD minimum transmit and receive interval (in milliseconds).

NOTE: You do not need to configure the BFD minimum transmit and receive interval if you configure the minimum-interval for the BFD transmit-interval statement and the minimum-receive-interval.

```
[edit system services dhcp-local-servergroup local_group_1 liveness-detection method
  bfd]
user@host# set minimum-interval 45000
```

8. Configure the minimum receive interval (in milliseconds).

NOTE: You do not need to configure the BFD minimum receive interval if you configure the BFD minimum transmit and receive interval.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set minimum-receive-interval 60000
```

9. Configure a multiplier value for the detection time.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set multiplier 100
```

10. Disable the ability for BFD interval timers to change or adapt to network situations.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set no-adaptation
```


Copyright © 2016, Juniper Networks, Inc.

11. Configure the BFD session mode.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set session-mode automatic
```

12. Configure the threshold and minimum interval for the BFD transmit interval.

NOTE: You do not need to configure the transmit interval values if you have already configured the minimum transmit and receive interval for BFD.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set transmit-interval threshold 60000 minimum-interval 45000
```

13. Configure the BFD protocol version you want to detect.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
  method bfd]
user@host# set version automatic
```

14. Configure the action the router (switch) takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit failure-action action
```

Results From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
user@host# show system
services {
  dhcp-local-server {
    group local_group_1 {
      liveness-detection {
        failure-action clear-binding-if-interface-up;
        method {
          bfd {
            version automatic;
            minimum-interval 45000;
            minimum-receive-interval 60000;
            multiplier 100;
            no-adaptation;
            transmit-interval {
              minimum-interval 45000;
              threshold 60000;
            }
            detection-time {
              threshold 30000;
            }
```

```
                    session-mode automatic;
                    holddown-interval 50;
                }
            }
        }
    }
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- *Extended DHCP Local Server Overview*
- DHCP Liveness Detection Overview on page 75
- Configuring Detection of DHCP Local Server Client Connectivity on page 81

High Availability Using Unified ISSU in the DHCP Access Network

The unified in-service software upgrade (unified ISSU) feature supports the DHCP access model used by subscriber management. This support ensures that the router preserves active DHCP subscriber sessions and session services after a unified ISSU has completed.

See *Unified ISSU Concepts* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

Unified ISSU supports the subscriber management DHCP access model, which includes DHCP local server, DHCPv6 local server, DHCP relay, and DHCP relay proxy.

Accounting, filter, and class-of-service (CoS) statistics for DHCP subscribers are preserved after a unified ISSU on MPC/MIC interfaces on MX Series routers.

Related Documentation

- Verifying and Monitoring Subscriber Management Unified ISSU State on page 91
- *Unified ISSU System Requirements*

Graceful Routing Engine Switchover for DHCP

For EX Series switches, only extended DHCP local server maintains the state of active DHCP client leases. The DHCP local 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 subscriber authentication. You can configure dynamic profile and authentication support on a global basis or for a specific group of interfaces. The extended DHCP local server also supports the use of Junos address-assignment pools or external authorities, such as RADIUS, to provide the client address and configuration information.

For MX Series routers, the extended DHCP local server and the DHCP relay agent applications both maintain the state of active DHCP client leases in the session database.

Copyright © 2016, Juniper Networks, Inc.

The extended DHCP application can recover this state if the DHCP process fails or is manually restarted, thus preventing the loss of active DHCP clients in either of these circumstances. However, the state of active DHCP client leases is lost if a power failure occurs or if the kernel stops operating (for example, when the router is reloaded) on a single Routing Engine.

You can enable graceful switchover support on both EX Series switches and MX Series routers. To enable graceful switchover support for the extended DHCP local server or extended DHCP relay agent on a switch, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To enable graceful Routing Engine switchover support on MX Series routers, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. You cannot disable graceful Routing Engine switchover support for the extended DHCP application when the router is configured to support graceful Routing Engine switchover.

For more information about using graceful Routing Engine switchover, see *Understanding Graceful Routing Engine Switchover*.

Related Documentation

- *Extended DHCP Local Server Overview*
- *Extended DHCP Relay Agent Overview*
- High Availability Using Unified ISSU in the PPP Access Network on page 115

Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to wait 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover (GRES) takes place.

- Benefits of Delaying Removal of Access Routes and Access-Internal Routes on page 87
- Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes on page 88
- Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes on page 88

Benefits of Delaying Removal of Access Routes and Access-Internal Routes

The 3-minute delay in removing access routes and access-internal routes after a graceful Routing Engine switchover provides sufficient time for the DHCP client process (jdhcpd), PPP client process (jpppd), or routing protocol process (rpd) to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

Copyright © 2016, Juniper Networks, Inc.

Configuring the router to delay removal of access routes and access-internal routes after a graceful Routing Engine switchover has the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine
- Prevents loss of subscriber traffic due to unavailable routes

Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale access routes and access-internal routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the jdhcpd or jpppd client process to reinstall all of the subscriber routes.

Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process (rpd) immediately purges the stale access routes and access-internal routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

Related Documentation

- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 88
- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- *Examples: Configuring Static Routes*

Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover

In subscriber networks configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to delay for 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover takes place.

To configure the router to delay removal (flushing) of access-routes and access-internal routes after a graceful Routing Engine switchover:

1. Specify that you want to configure subscriber management.

 Copyright © 2016, Juniper Networks, Inc.

```
[edit system services]
user@host# edit subscriber-management
```

2. Configure the router to wait 180 seconds before removing access-routes and access-internal routes after a graceful Routing Engine switchover.

```
[edit system services subscriber-management]
user@host# set gres-route-flush-delay
```

Related Documentation

- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 87
- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- *Examples: Configuring Static Routes*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 8

Monitoring and Managing DHCP for Subscriber Access

Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers

Purpose View configuration information for access routes and access-internal routes on DHCP subscribers.

Action

- To display extensive information about access routes and access-internal routes:

  user@host>**show route extensive**

- To display the configuration for access routes:

  user@host>**show route protocol access**

- To display the configuration for access-internal routes:

  user@host> **show route protocol access-internal**

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane queiscing has not started). The second example shows that unified ISSU

is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

Related Documentation

- High Availability Using Unified ISSU in the PPP Access Network on page 115
- High Availability Using Unified ISSU in the DHCP Access Network on page 86
- High Availability Using Unified ISSU in the L2TP Access Network on page 209
- *Unified ISSU Concepts*

Copyright © 2016, Juniper Networks, Inc.

PART 2

Configuring the PPP Access Network

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 9

Configuring PPP for Subscriber Access

Dynamic Profiles for PPP Subscriber Interfaces Overview

Subscriber management 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 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.

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 103

- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- Verifying and Managing PPP Configuration for Subscriber Management on page 119
- Example: Minimum PPPoE Dynamic Profile on page 99

Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests

On MX Series routers with Modular Port Concentrators/Modular Interface Cards (MPCs/MICs), the Packet Forwarding Engine on an MPC/MIC processes and responds to Link Control Protocol (LCP) Echo-Request packets that the PPP subscriber (client) initiates and sends to the router. LCP Echo-Request packets and LCP Echo-Reply packets are part of the PPP keepalive mechanism that helps determine whether a link is functioning properly.

Previously, LCP Echo-Request packets and LCP Echo-Reply packets were handled on an MX Series router by the Routing Engine. Support for the PPP fast keepalive feature enables the Packet Forwarding Engine on the MPC/MIC to receive LCP Echo-Request packets from the PPP subscriber and transmit LCP Echo-Reply packets in response, without having to send the LCP packets to the Routing Engine for processing. The mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine is referred to as *PPP fast keepalive*.

Relieving the Routing Engine of having to process LCP Echo-Request packets provides increased bandwidth on the router to support a larger number of subscribers with improved performance.

- How PPP Fast Keepalive Processing Works on page 96
- Statistics Display for PPP Fast Keepalive on page 97
- Effect of Changing the Forwarding Class Configuration on page 97

How PPP Fast Keepalive Processing Works

You do not need any special configuration on an MX Series router with MPCs/MICs to enable processing of PPP fast keepalive requests on the Packet Forwarding Engine. The feature is enabled by default, and cannot be disabled.

The following sequence describes how an MX Series router processes LCP Echo-Request packets and LCP Echo-Reply packets on the Packet Forwarding Engine on the MPC/MIC:

1. The Routing Engine notifies the Packet Forwarding Engine when transmission of keepalive requests is enabled on a PPP logical interface. The notification includes the magic numbers of both the server and the remote client.
2. The Packet Forwarding Engine receives the LCP Echo-Request packet initiated by the PPP subscriber (client).

 Copyright © 2016, Juniper Networks, Inc.

3. The Packet Forwarding Engine validates the peer magic number in the LCP Echo-Request packet, and transmits the corresponding LCP Echo-Reply packet containing the magic number negotiated by the router.
4. If the Packet Forwarding Engine detects a loop condition in the link, it sends the LCP Echo-Request packet to the Routing Engine for further processing.

   The Routing Engine continues to process LCP Echo-Request packets until the loop condition is cleared.

Transmission of keepalive requests from the Packet Forwarding Engine on the router is not currently enabled.

Statistics Display for PPP Fast Keepalive

When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the **Keepalive statistics** field in the output of the **show interfaces pp0.*logical* statistics** operational command does not include statistics for the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.

Effect of Changing the Forwarding Class Configuration

To change the default queue assignment (forwarding class) for outbound traffic generated by the Routing Engine, you can include the **forwarding-class *class-name*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level.

For PPP fast (inline) keepalive LCP Echo-Request and LCP Echo-Reply packets transmitted between an MX Series router with MPCs/MICs and a PPP client, changing the forwarding class configuration takes effect immediately for both new PPP-over-Ethernet (PPPoE), PPP-over-ATM (PPPoA), and L2TP network server (LNS) subscriber sessions created after the configuration change, and for existing PPPoE, PPPoA, and LNS subscriber sessions established before the configuration change.

Related Documentation

- *Configuring Keepalives*
- *Disabling the Sending of PPPoE Keepalive Messages*
- *Changing the Default Queuing and Marking of Host Outbound Traffic*

Configuring Dynamic Profiles for PPP

A dynamic profile acts as a template that enables you to create, update, or remove a configuration that includes attributes for client access (for example, interface or protocol) or service (for example, IGMP). Using these profiles you can consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes simultaneously.

After they are created, the profiles reside in a profile library on the router. You can then use the **dynamic-profile** statement to attach profiles to interfaces. To assign a dynamic profile to a PPP interface, you can include the **dynamic-profile** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number ppp-options]
dynamic-profile profile-name;
```

To monitor the configuration, issue the **show interfaces *interface-name*** command.

For information about dynamic profiles, see *Dynamic Profiles Overview* in the *Junos Subscriber Access Configuration Guide*.

For information about creating dynamic profiles, see *Configuring a Basic Dynamic Profile* in the *Junos Subscriber Access Configuration Guide*.

For information about assigning a dynamic profile to a PPP interface, see "Attaching Dynamic Profiles to Static PPP Subscriber Interfaces" on page 98 in the *Junos Subscriber Access Configuration Guide*.

NOTE: **Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces for this release.**

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 103

Attaching Dynamic Profiles to Static PPP Subscriber Interfaces

You can attach a dynamic profile to a static PPP subscriber interface. When a PPP subscriber logs in, the specified dynamic profile is instantiated and the services defined in the profile are applied to the interface.

To attach a dynamic profile to a static PPP subscriber interface:

1. Specify that you want to configure PPP options.

   ```
   [edit interfaces pp0 unit 0]
   user@host# edit ppp-options
   ```

2. Specify the dynamic profile you want to associate with the interface.

   ```
   [edit interfaces pp0 unit 0 ppp-options]
   user@host# set dynamic-profile vod-profile-50
   ```


Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- Example: Minimum PPPoE Dynamic Profile on page 99
- Verifying and Managing PPP Configuration for Subscriber Management on page 119

Example: Minimum PPPoE Dynamic Profile

This example shows the minimum configuration for a dynamic profile that is used for static PPPoE interfaces. The configuration must include the **interfaces pp0** stanza.

```
dynamic-profiles {
  ppp-profile-1 {
    interfaces {
      pp0 {
        unit "$junos-interface-unit";
      }
    }
  }
}
```

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 10

Applying RADIUS Route Attributes to Subscribers or Access Networks

Configuring Dynamic Access-Internal Routes for PPP Subscriber Management

You can dynamically configure access-internal routes for PPP subscribers. Configuring support for access-internal variables is optional, but it ensures that values from the access-internal variables are used if the next-hop value is missing in the relevant RADIUS attribute—Framed-Route [22] for IPv4 and Framed-IPv6-Route [99] for IPv6.

BEST PRACTICE: **We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed route support.**

For PPP subscriber interfaces, you do not need to specify the MAC address for access-internal routes.

To dynamically configure access-internal routes for PPP:

1. Specify that you want to configure the access-internal route.

   user@host# **edit dynamic-profiles** *profile-name* **routing-options**

2. Specify the IP address as a variable.

   [edit dynamic-profiles *profile-name* routing-options]
   user@host# **edit access-internal route $junos-subscriber-ip-address**

3. Specify the qualified-next-hop as a variable.

   [edit dynamic-profiles *profile-name* routing-options access-internal route $junos-subscriber-ip-address]
   user@host# **set qualified-next-hop $junos-interface-name**

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Verifying the Configuration of Access and Access-Internal Routes for DHCP Subscribers on page 91

Verifying the Configuration of Access and Access-Internal Routes for PPP Subscribers

Purpose View configuration information for access routes and access-internal routes on PPP subscribers.

Action

- To display extensive information about access routes and access-internal routes:

  user@host>**show route extensive**

- To display the configuration for access routes:

  user@host>**show route protocol access**

- To display the configuration for access-internal routes:

  user@host> **show route protocol access-internal**

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 11

Configuring Authentication for PPP

Configuring Dynamic Authentication for PPP Subscribers

You can configure a dynamic profile that includes PPP authentication that enables PPP clients to dynamically access the network. You can specify either CHAP or PAP authentication. Optionally, you can also control the order in which the router negotiates the CHAP and PAP protocols.

For dynamic interfaces, the router supports unidirectional authentication only—the router always functions as the authenticator. When you configure PPP authentication in a dynamic profile, CHAP authentication supports the **challenge-length** option, which enables you to configure the minimum length and maximum length of the CHAP challenge message. Neither CHAP authentication nor PAP authentication supports any other configuration options, including the **passive** statement.

NOTE: **Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces.**

To configure authentication in a dynamic profile for PPP subscriber interfaces:

1. Name the dynamic profile.

   [edit]
   user@host# **edit dynamic-profiles vod-profile-25**

2. Configure the interfaces and unit for the dynamic profile. Use **pp0** for the interface type and the Junos predefined variable for the unit.

   [edit dynamic-profiles vod-profile-25]
   user@host# **edit interfaces pp0 unit $junos-interface-unit**

3. Configure PPP options.

   [edit dynamic-profiles vod-profile-25 interfaces pp0 unit “$junos-interface-unit”]
   user@host# **edit** ppp-options

4. Specify the authentication protocol used in the dynamic profile. You can configure either CHAP or PAP. There are no additional options for either authentication protocol.

   ```
   [edit dynamic-profiles vod-profile-25 interfaces pp0 unit "$junos-interface-unit"
     ppp-options]
   user@host# set chap
   ```

5. (Optional) Configure the minimum length and maximum length of the CHAP challenge message.

   See "Modifying the CHAP Challenge Length" on page 104.

6. (Optional) Configure the order in which the router negotiates the CHAP and PAP authentication protocols.

   See "Controlling the Negotiation Order of PPP Authentication Protocols" on page 110.

Related Documentation

- Modifying the CHAP Challenge Length on page 104
- Controlling the Negotiation Order of PPP Authentication Protocols on page 110
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- Example: Minimum PPPoE Dynamic Profile on page 99
- Verifying and Managing PPP Configuration for Subscriber Management on page 119

Modifying the CHAP Challenge Length

You can modify the default minimum length and maximum length of the Challenge Handshake Authentication Protocol (CHAP) challenge message that the router sends to a PPP client. The CHAP challenge message, which contains information that is unique to a particular PPP subscriber session, is used as part of the authentication mechanism between the router and the client to verify the identity of the client for access to the router.

By default, the minimum length of the CHAP challenge is 16 bytes, and the maximum length is 32 bytes. You can override this default to configure the CHAP challenge minimum length and maximum length in the range 8 bytes through 63 bytes.

BEST PRACTICE: **We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.**

Before you begin:

- Configure the CHAP protocol on the interface.

Copyright © 2016, Juniper Networks, Inc.

- For dynamic PPP subscriber interfaces, see "Configuring Dynamic Authentication for PPP Subscribers" on page 103.
- For static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

To configure the minimum and maximum length of the CHAP challenge message:

1. Specify that you want to configure PPP options.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]
     user@host# edit ppp-options

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number*]
     user@host# edit ppp-options

2. Specify that you want to configure CHAP options.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options]
     user@host# edit chap

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number* ppp-options]
     user@host# edit chap

3. Specify the minimum length and maximum length of the CHAP challenge.
   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options chap]
     user@host# set challenge-length minimum *minimum-length* maximum *maximum-length*

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number* ppp-options chap]
     user@host# set challenge-length minimum *minimum-length* maximum *maximum-length*

   For example, the following **challenge-length** statement in a dynamic profile named pppoe-client-profile sets the minimum length of the CHAP challenge to 20 bytes, and the maximum length to 40 bytes.

   [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit" ppp-options chap]
   user@host# set challenge-length minimum 20 maximum 40

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- *Configuring the PPP Challenge Handshake Authentication Protocol*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 12

Configuring PPP Network Control Protocol Negotiation

PPP Network Control Protocol Negotiation Mode Overview

The *Network Control Protocol* (NCP) is a mechanism used to establish and configure different Network Layer protocols for Point-to-Point Protocol (PPP) connections. On MX Series routers with Modular Port Concentrators (MPCs), you can configure *PPP NCP negotiation* to actively or passively control subscriber connections initiated by the router functioning as a PPP server.

PPP NCP Negotiation Modes

PPP NCP negotiation operates in either of the following modes:

- *Active PPP NCP negotiation mode*—The router sends an NCP Configuration Request message without waiting for the PPP client to do so.
- *Passive PPP NCP negotiation mode*—The router waits for the PPP client to send an NCP Configuration Request message before sending its own Configuration Request message. Dynamic subscriber interface connections and static subscriber interface connections use passive PPP NCP negotiation by default.

Copyright © 2016, Juniper Networks, Inc.

Router behavior for active mode and passive mode PPP NCP negotiation differs for dynamic PPP subscribers and static PPP subscribers, as summarized in Table 8 on page 108.

Table 8: PPP NCP Negotiation Mode Behavior for Dynamic and Static Subscribers

PPP Subscribers	PPP NCP Negotiation Mode	Router Behavior
Dynamic	Active	The router establishes the local network address and uses it to send the NCP Configuration Request message without waiting for the PPP client to send a Configuration Request.
Dynamic	Passive	The router establishes the local network address after it receives the NCP Configuration Request message from the PPP client.
Static	Active	The router sends the authentication acknowledgement to the PPP client, and then sends the NCP Configuration Request message without waiting for the PPP client to send its own Configuration Request.
Static	Passive	The router sends the authentication acknowledgement to the PPP client, and then waits for an NCP Configuration Request message from the client before sending a Configuration Request.

PPP NCP Negotiation Mode Supported Configurations

You can configure PPP Network Control Protocol (NCP) negotiation for the following single-stack and dual-stack subscriber configurations on MX Series routers with MPCs:

- Dynamic PPP subscriber connections terminated at the router
- Static PPP subscriber connections terminated at the router
- Dynamic tunneled PPP subscribers at the L2TP network server (LNS)
- Static tunneled PPP subscribers at the L2TP network server (LNS) on an inline service (**si**) interface

PPP NCP Active Negotiation Requirements for IPv4 Dynamic and Static PPP Subscribers

To configure active PPP IPv4 Network Control Protocol (IPNCP) negotiation for dynamic and static PPP subscribers in a single-stack or dual-stack configuration, make sure you meet the following requirements:

- Configure the IPv4 (**inet**) protocol family in a dynamic profile (for dynamic subscribers) or at the interface level (for static subscribers).
- Assign any of the following IPv4 address attributes for the subscriber during the authentication process:

 Copyright © 2016, Juniper Networks, Inc.

- Framed-IP-Address (RADIUS Attribute 8)—RADIUS explicit IPv4 address
- Framed-Pool (RADIUS Attribute 88)—RADIUS IPv4 adress pool name
- IPv4 attributes allocated from a locally configured address pool

When you have met these requirements, use the **initiate-ncp ip** statement to enable active IPNCP negotiation for dynamic and static subscribers in a single-stack or dual-stack configuration.

PPP NCP Active Negotiation Requirements for IPv6 Dynamic and Static PPP Subscribers

To configure active PPP IPv6 Network Control Protocol (IPv6NCP) negotiation for dynamic and static PPP subscribers in a single-stack or dual-stack configuration, make sure you meet the following requirements:

- Configure the IPv6 (**inet6**) protocol family in a dynamic profile (for dynamic subscribers) or at the interface level (for static subscriber).
- Assign any of the following IPv6 address attributes for the subscriber during the authentication process:
  - Delegated-IPv6-Prefix (RADIUS Attribute 123)—RADIUS explicit IPv6 address
  - Framed-IPv6-Prefix (RADIUS Attribute 97)—RADIUS explicit IPv6 prefix
  - Framed-IPv6-Pool (RADIUS Attribute 100)—RADIUS explicit IPv6 adress or prefix pool name
  - IPv6 attributes allocated from a locally configured Neighbor Discovery Router Advertisement (NDRA) pool

When you have met these requirements, use the **initiate-ncp ipv6** statement to enable active IPv6NCP negotiation for dynamic and static subscribers in a single-stack or dual-stack configuration.

PPP NCP Negotiation Requirements for IPv4 and IPv6 Dual-Stack Configurations

You can configure either active or passive PPP NCP negotiation for the IPv4 and IPv6 subscriber interfaces in a dual-stack configuration.

To configure active negotiation in a dual-stack configuration, do all of the following:

- Make sure you meet the IPv4 and IPv6 protocol and address family requirements.
- Use the **initiate-ncp ip** statement to enable active negotiation for the IPv4 subscriber interface.
- Use the **initiate-ncp ipv6** statement to enable active negotiation for the IPv6 subscriber interface.

To configure passive negotiation in a dual-stack configuration, do both of the following:

- Make sure you meet the IPv4 and IPv6 protocol and address family requirements.

- Use the **initiate-ncp dual-stack-passive** statement to enable passive negotiation for the dual-stack configuration. The **initiate-ncp dual-stack-passive** statement overrides the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured.

The following additional guidelines apply when you configure PPP NCP negotiation for dual-stack subscribers:

- Dual-stack subscribers configured for either active mode or passive mode PPP NCP negotiation continue to use the same negotiation mode when the NCP mechanism is renegotiated.
- Using the **on-demand-ip-address** statement to save IPv4 addresses for dual-stack PPP subscribers when you are not using the IPv4 service has no effect on configuration of the PPP NCP negotiation mode in a dual-stack configuration.

Related Documentation

- Configuring the PPP Network Control Protocol Negotiation Mode on page 112

Controlling the Negotiation Order of PPP Authentication Protocols

You can control the order in which the router tries to negotiate PPP authentication protocols when it verifies that a PPP client can access the network. By default, the router first tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication. If the the attempt to negotiate CHAP authentication is unsuccessful, the router then tries to negotiate Password Authentication Protocol (PAP) authentication.

You can modify this default negotiation order in any of the following ways:

- Specify that the router negotiate PAP authentication first, followed by CHAP authentication if PAP negotiation is unsuccessful.

  When you specify both authentication protocols in either order, you must enclose the set of protocol names in square brackets ([ ]).
- Specify that the router negotiate only CHAP authentication.
- Specify that the router negotiate only PAP authentication.

Before you begin:

- Configure the CHAP or PAP protocol on the interface.
  - For dynamic PPP subscriber interfaces, see "Configuring Dynamic Authentication for PPP Subscribers" on page 103.
  - For CHAP on static interfaces with PPP encapsulation, see *Configuring the PPP Challenge Handshake Authentication Protocol*.
  - For PAP on static interfaces with PPP encapsulation, see *Configuring the PPP Password Authentication Protocol*.

To control the order in which the router negotiates PPP authentication protocols:

1. Specify that you want to configure PPP options.

 Copyright © 2016, Juniper Networks, Inc.

- For dynamic PPP subscriber interfaces:

  [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit"]
  user@host# **edit** ppp-options

- For static interfaces with PPP encapsulation:

  [edit interfaces pp0 unit *logical-unit-number*]
  user@host# **edit** ppp-options

2. Specify the negotiation order for PPP authentication protocols on the router.

   - For dynamic PPP subscriber interfaces:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options]
     user@host# **set** authentication **[*authentication-protocols*]**

   - For static interfaces with PPP encapsulation:

     [edit interfaces pp0 unit *logical-unit-number* ppp-options]
     user@host# **set** authentication **[*authentication-protocols*]**

The following sample **authentication** statements in a dynamic profile named pppoe-client-profile show the different ways you can configure the negotiation order for PPP authentication protocols. (The **authentication** statements for configuring static interfaces are identical.)

- To specify that the router negotiate PAP authentication first, followed by CHAP authentication:

  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit" ppp-options]
  user@host# **set authentication [pap chap]**

- To specify that the router negotiate only CHAP authentication:

  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit" ppp-options]
  user@host# **set authentication chap**

- To specify that the router negotiate only PAP authentication:

  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit" ppp-options]
  user@host# **set authentication pap**

- To restore the default negotiation order for PPP authentication protocols after you have modified it:

  [edit dynamic-profiles pppoe-client-profile interfaces pp0 unit "$junos-interface-unit" ppp-options]
  user@host# **set authentication [chap pap]**

Related Documentation

- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring the PPP Password Authentication Protocol*

Configuring the PPP Network Control Protocol Negotiation Mode

Configuring PPP Network Control Protocol (NCP) negotiation enables you to actively or passively control subscriber connections initiated by the router functioning as a PPP server. Both dynamic and static subscriber interface connections use passive PPP NCP negotiation by default.

You can configure the PPP NCP negotiation mode (active or passsive) for the following subscriber configurations on MX Series routers with MPCs:

- Dynamic PPP subscriber connections terminated at the router, using a dynamic profile
- Static PPP subscriber connections terminated at the router, using a per-interface configuration
- Dynamic tunneled PPP subscribers at the L2TP network server (LNS), using a dynamic profile
- Static tunneled PPP subscribers at the LNS, using a per-inline service (**si**) interface configuration
- Dynamic and static tunneled PPP subscribers at the LNS, using a user-group profile

To configure PPP NCP negotiation mode:

1. Specify that you want to configure PPP-specific properties for the subscriber.
   - For dynamic PPP subscriber connections terminated at the router:

     [edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]
     user@host# **edit ppp-options**

   - For static PPP subscriber connections terminated at the router:

     [edit interfaces pp0 unit *logical-unit-number*]
     user@host# **edit ppp-options**

   - For dynamic tunneled PPP subscribers at the LNS:

     [edit dynamic-profiles *profile-name* interfaces “$junos-interface-ifd-name” unit “$junos-interface-unit”]
     user@host# **edit ppp-options**

   - For static tunneled PPP subscribers at the LNS:

     [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
     user@host# **edit ppp-options**

   - In a group profile for dynamic and static tunneled PPP subscribers at the LNS:

     [edit access group-profile *profile-name* ppp]
     user@host# **edit ppp-options**

2. Configure PPP NCP negotiation mode in any of the following ways:
   - To configure active PPP NCP negotiation for IPv4 subscribers in a single-stack or dual-stack configuration, use the **initiate-ncp ip** statement.

Copyright © 2016, Juniper Networks, Inc.

For example, to configure active negotiation for static IPv4 connections terminated at the router:

```
[edit interfaces pp0 unit logical-unit-number ppp-options]
user@host# initiate-ncp ip
```

- To configure active PPP NCP negotiation for IPv6 subscribers in a single-stack or dual-stack configuration, use the **initiate-ncp ipv6** statement.

  For example, to configure active negotiation for dynamic IPv6 connections terminated at the router:

  ```
  [edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit"
    ppp-options]
  user@host# initiate-ncp ipv6
  ```

- To configure passive PPP NCP negotiation for dynamic or static subscribers in an IPv4 and IPv6 dual-stack configuration, use the **initiate-ncp dual-stack-passive** statement, which overrides both the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured.

  For example, to configure passive negotiation for dynamic tunneled PPP subscribers at the LNS in an IPv4 and IPv6 dual-stack configuration:

  ```
  [edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit
    "$junos-interface-unit"]
  user@host# initiate-ncp dual-stack-passive
  ```

Related Documentation

- PPP Network Control Protocol Negotiation Mode Overview on page 107

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 13

Configuring High Availability in the PPP Access Network

High Availability Using Unified ISSU in the PPP Access Network

The unified in-service software upgrade (unified ISSU) feature supports the PPPoE access model used by subscriber management. This support ensures that the router preserves active PPPoE subscriber sessions and session services after a unified ISSU has completed.

See *Unified ISSU Concepts* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

Unified ISSU supports the subscriber management PPPoE access model for static and dynamic PPPoE access, and includes the following features:

- Terminated, non-tunneled PPPoE connections configured with static or dynamic PPP logical interfaces and static or dynamic underlying interfaces
- Subscriber services on single-link PPP interfaces
- Preservation of statistics for accounting, filter, and CoS on MPC/MIC interfaces

  Accounting statistics for PPPoE subscribers are *not* preserved after a unified ISSU on Enhanced Intelligent Queuing 2 (IQ2E) PICs on M120 and M320 routers.

Unified ISSU for the subscriber management PPPoE access model *does not support* Multilink Point-to-Point Protocol (MLPPP) bundle interfaces. MLPPP bundle interfaces require the use of an Adaptive Services PIC or Multiservices PIC to provide PPP subscriber services. These PICs do not support unified ISSU.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Verifying and Monitoring Subscriber Management Unified ISSU State on page 91
- *Unified ISSU System Requirements*

Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover

For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to wait 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover (GRES) takes place.

- Benefits of Delaying Removal of Access Routes and Access-Internal Routes on page 116
- Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes on page 116
- Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes on page 116

Benefits of Delaying Removal of Access Routes and Access-Internal Routes

The 3-minute delay in removing access routes and access-internal routes after a graceful Routing Engine switchover provides sufficient time for the DHCP client process (jdhcpd), PPP client process (jpppd), or routing protocol process (rpd) to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes for DHCP subscribers and PPP subscribers.

Configuring the router to delay removal of access routes and access-internal routes after a graceful Routing Engine switchover has the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine
- Prevents loss of subscriber traffic due to unavailable routes

Graceful Restart and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale access routes and access-internal routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the jdhcpd or jpppd client process to reinstall all of the subscriber routes.

Nonstop Active Routing and Delayed Removal of Access Routes and Access-Internal Routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process (rpd) immediately purges the stale

 Copyright © 2016, Juniper Networks, Inc.

access routes and access-internal routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic.

Configuring the delay in removing access and access-internal routes after a graceful Routing Engine switchover causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

Related Documentation

- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 88
- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- *Examples: Configuring Static Routes*

Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover

In subscriber networks configured with either nonstop active routing (NSR) or graceful restart, you can configure the router to delay for 180 seconds (3 minutes) before removing access routes and access-internal routes for DHCP and PPP subscriber management after a graceful Routing Engine switchover takes place.

To configure the router to delay removal (flushing) of access-routes and access-internal routes after a graceful Routing Engine switchover:

1. Specify that you want to configure subscriber management.

   ```
   [edit system services]
   user@host# edit subscriber-management
   ```

2. Configure the router to wait 180 seconds before removing access-routes and access-internal routes after a graceful Routing Engine switchover.

   ```
   [edit system services subscriber-management]
   user@host# set gres-route-flush-delay
   ```

Related Documentation

- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 87
- Access and Access-Internal Routes for Subscriber Management on page 43
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- *Examples: Configuring Static Routes*

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 14

Monitoring and Managing PPP for Subscriber Access

Verifying and Managing PPP Configuration for Subscriber Management

Purpose View or clear information about PPP configuration for subscriber management.

Action

- To display information about PPP interfaces:

  user@host> show ppp interface

- To display PPP statistics information:

  user@host> show ppp statistics

- To display PPP session summary information:

  user@host> show ppp summary

- To display PPP address-pool information:

  user@host>**show ppp address-pool**

Related Documentation

- Dynamic Profiles for PPP Subscriber Interfaces Overview on page 95
- CLI Explorer

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane queiscing has not started). The second example shows that unified ISSU is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
```

```
    Mastership           Master
    Database             Available
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0

user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

Related Documentation

- High Availability Using Unified ISSU in the PPP Access Network on page 115
- High Availability Using Unified ISSU in the DHCP Access Network on page 86
- High Availability Using Unified ISSU in the L2TP Access Network on page 209
- *Unified ISSU Concepts*

Copyright © 2016, Juniper Networks, Inc.

PART 3

Configuring the L2TP Access Network

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 15

L2TP and Subscriber Access Overview

L2TP for Subscriber Access Overview

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 4 on page 123 shows a simple L2TP topology.

Figure 4: Typical L2TP Topology

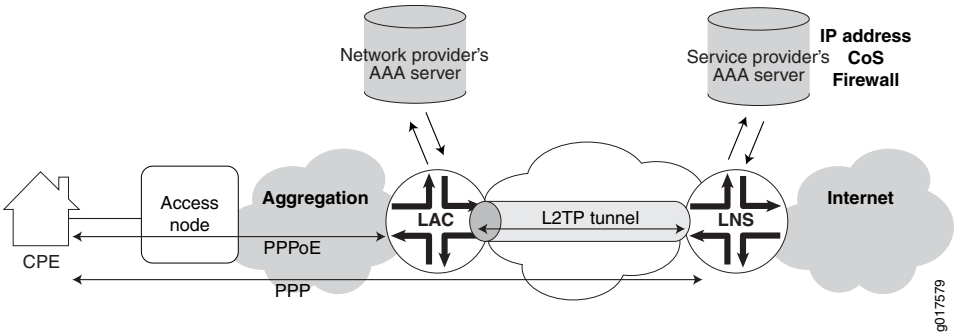

L2TP separates the termination of access technologies, such as cable or xDSL, from the termination of PPP and subsequent access to a network. This separation enables public ISPs to outsource their access technologies to competitive local exchange carriers (CLECs). L2TP provides ISPs the capability to 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

Copyright © 2016, Juniper Networks, Inc.

logical interfaces. Figure 5 on page 124 shows the protocol layer stacking for an L2TP pass-through connection.

Figure 5: Protocol Stacking for L2TP Subscribers in Pass-Through Mode

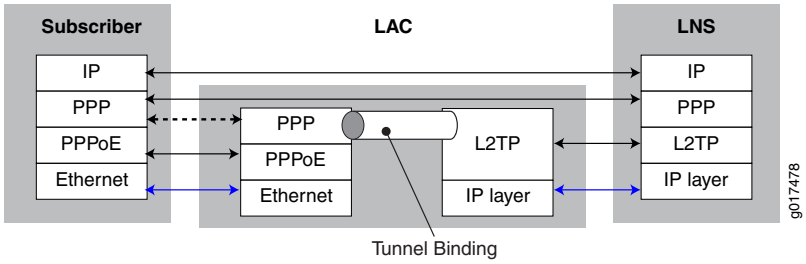

NOTE: On MX Series routers, the LAC and LNS functions are supported only on MPCs; they are not supported on any services PIC or MS-DPC. For details about MPC support for L2TP, see the *MX Series Interface Module Reference*

Certain M Series routers support LNS functions on services PICs. For more information about the L2TP implementation on M Series routers, see *L2TP Services Configuration Overview*.

The LAC dynamically creates tunnels based on AAA authentication parameters and transmits L2TP packets to the LNS by means of the IP/User Datagram Protocol (UDP). Traffic travels in an L2TP *session*; a tunnel is an aggregation of one or more sessions. You can also provision a domain map that is used by AAA to determine whether to tunnel or terminate the PPPoE subscriber on the LAC. A one-to-one mapping exists between each PPP subscriber tunneled to the LNS and an L2TP session.

When the LNS is an MX Series router, a LAC-facing peer interface on an MPC provides an IP address for the exchange of IP packets between the tunnel endpoints; the Routing Engine maintains the L2TP tunnels. The Packet Forwarding Engine hosts one or more inline services (**si**) interfaces. These interfaces function like a virtual physical interface and *anchor* the L2TP sessions on the LNS. The **si** interface enables L2TP services without requiring a special services PIC. Finally, another interface is used to transmit the subscriber data to and from the Internet.

The characteristics of the tunnel can originate either from a tunnel profile that you configure or from RADIUS tunnel attributes and vendor-specific attributes (VSAs) from the AAA server accessible at the LAC. You can include a tunnel profile in a domain map, which applies the tunnel profile before RADIUS authentication takes place. You can use RADIUS standard attributes and VSAs to override any or all characteristics configured by the tunnel profile in a domain map. Alternatively, RADIUS can itself apply a tunnel profile when the RADIUS Tunnel-Group VSA [26-64] is specified in the RADIUS login.

The Virtual-Router VSA [26-1] in the subscriber profile on the service provider AAA server (accessible from the LNS) determines the routing instance in which the L2TP session is brought up on the LNS. When this VSA is not present, the subscriber session comes up in the same routing instance as the tunnel, because the AAA server can be accessed only from the routing instance in which the tunnel terminates on the LNS.

Copyright © 2016, Juniper Networks, Inc.

This behavior is different than for DHCP and non-tunneled PPPoE subscribers, which come up in the default routing instance in the absence of the Virtual-Router VSA. For L2TP subscribers, you must include this VSA in the subscriber profile when you want the subscriber session to come up in a different routing instance than the tunnel routing instance.

The LAC supports RADIUS-initiated mirroring, which creates secure policies based on certain RADIUS VSAs, and uses RADIUS attributes to identify a subscriber whose traffic is to be mirrored. (This feature is not supported for an LNS configured on an MX Series router.)

The LAC and the LNS support unified ISSU. When an upgrade is initiated, the LAC completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade. Subscriber logouts are recorded during the upgrade and are completed after the upgrade has completed.

Related Documentation

- *RADIUS IETF Attributes Supported by the AAA Service Framework*
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- Configuring a Tunnel Profile for Subscriber Access on page 162
- *Domain Mapping Overview*
- *Unified ISSU Concepts*

L2TP Terminology

Table 9 on page 125 describes the basic terms for L2TP.

Table 9: L2TP Terms

Term	Description
AVP	Attribute value pair (AVP)—Combination of a unique attribute—represented by an integer—and a value containing the actual value identified by the attribute.
Call	A connection (or attempted connection) between a remote system and the LAC.
LAC	L2TP access concentrator (LAC)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LNS. The LAC sits between an LNS and a remote system and forwards packets to and from each.
LNS	L2TP network server (LNS)—A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. The LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC.
Peer	In the L2TP context, either the LAC or LNS. The LAC's peer is an LNS, and vice versa.

Copyright © 2016, Juniper Networks, Inc.

Table 9: L2TP Terms *(continued)*

Term	Description
Proxy authentication	PPP pre-authentication performed by the LAC on behalf of the LNS. The proxy data is sent by the LAC to the LNS containing attributes such as authentication type, authentication name, and authentication challenge. The LNS responds with the authentication results.
Proxy LCP	Link Control Protocol (LCP) negotiation that is performed by the LAC on behalf of the LNS. The proxy is sent by the LAC to the LNS containing attributes such as the last configuration attributes sent and received from the client.
Remote system	An end system or router attached to a remote access network, which is either the initiator or recipient of a call.
Session	A logical connection created between the LAC and the LNS when an end-to-end PPP connection is established between a remote system and the LNS. NOTE: There is a one-to-one relationship between established L2TP sessions and their associated PPP connections.
Tunnel	A connection between the LAC-LNS pair consisting of a control connection and 0 or more L2TP sessions.

Related Documentation

- L2TP for Subscriber Access Overview on page 123

L2TP Implementation

L2TP is implemented on four levels:

- Source—The local router acting as the LAC.
- Destination—The remote router acting as the LNS.
- Tunnel—A direct path between the LAC and the LNS.
- Session—A PPP connection in a tunnel.

When the router has established destinations, tunnels, and sessions, you can control the L2TP traffic. Making a change to a destination affects all tunnels and sessions to that destination; making a change to a tunnel affects all sessions in that tunnel. For example, closing a destination closes all tunnels and sessions to that destination.

Sequence of Events on the LAC

The router acting as the LAC creates destinations, tunnels, and sessions dynamically, as follows:

1. The client initiates a PPP connection with the router.
2. The router and the client exchange Link Control Protocol (LCP) packets. The LAC negotiates on behalf of the LNS; this is known as *proxy LCP*.

Copyright © 2016, Juniper Networks, Inc.

3. The LAC authenticates the client on behalf of the LNS; this is known as *proxy authentication*. By using either a local database related to the domain name or RADIUS authentication, the router determines either to terminate or to tunnel the PPP connection.
4. If the router discovers that it should tunnel the session, it does the following:
   a. Sets up a new destination or selects an existing destination.
   b. Sets up a new tunnel or selects an existing tunnel.

      When a shared secret is configured in either the tunnel profile or the RADIUS attribute Tunnel-Password [69]—depending on which method is used to configure the tunnel—the secret is used to authenticate the tunnel during the establishment phase. The LAC includes the Challenge AVP in the SCCRQ message sent to the LNS. The LNS returns the Challenge Response AVP in the SCCRP message. If the response from the LNS does not match the value expected by the LAC, then tunnel authentication fails and the tunnel is not established.
   c. Opens a new session.
5. The router forwards the results of the LCP negotiations and authentication to the LNS.

A PPP connection now exists between the client and the LNS.

NOTE: **The router discards received packets if the size of the variable-length, optional offset pad field in the L2TP header is too large. The router always supports packets that have an offset pad field of up to 16 bytes, and may support larger offset pad fields, depending on other information in the header. This restriction is a possible, although unlikely, cause of excessive discarding of L2TP packets.**

Sequence of Events on the LNS

A router acting as an LNS might be set up as follows:

1. The LAC initiates a tunnel with the router acting as the LNS.
2. The LNS verifies that a tunnel with this LAC is valid: the destination is configured, the hostname and the tunnel password are correct.
3. The LNS completes the tunnel setup with the LAC.
4. The LAC sets up a session and initiates a session request to the LNS.
5. The LNS uses a static interface or creates a dynamic interface to anchor the PPP session.
6. If they are enabled and present, the LNS accepts the proxy LCP and the proxy authentication data and passes them to PPP.
7. PPP processes the proxy LCP, if it is present, and, if the proxy LCP is acceptable, places LCP on the LNS in opened state without renegotiation of LCP.

Copyright © 2016, Juniper Networks, Inc.

8. PPP processes the proxy authentication data, if it is present, and passes the data to AAA for verification. (If the data is not present, PPP requests the data from the peer.)

   NOTE: When the proxy LCP is not present or not acceptable, the LNS negotiates LCP with the peer. When LCP renegotiation is enabled on the LNS, the LNS ignores any pre-negotiated LCP parameters and renegotiates both the LCP parameters and PPP authentication with the PPP client.

9. The LNS passes the authentication results to the peer.

Related Documentation

- L2TP for Subscriber Access Overview on page 123

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 16

Configuring L2TP Tunneling and Switching for Subscribers

L2TP Tunnel Switching Overview

L2TP tunnel switching, also known as L2TP multihop, simplifies the deployment of an L2TP network across multiple domains. A router that lies between a LAC and an LNS is configured as an *L2TP tunnel switch* (LTS)—sometimes referred to simply as a *tunnel switch* or a *tunnel switching aggregator* (TSA)—as shown in Figure 6 on page 130. The LTS is configured as both an LNS and a LAC. When a remote LAC sends encapsulated PPP packets to the LNS configured on the LTS, the LTS can forward or redirect the packets through a different tunnel to a different LNS beyond the LTS. The logical termination point of the original L2TP session is switched to a different endpoint.

For example, in the network shown in Figure 6 on page 130, packets from the subscriber provisioned by service provider A are initially targeted at the LNS configured on the LTS. The LTS might redirect those packets to LNS1.

Copyright © 2016, Juniper Networks, Inc.

Figure 6: L2TP Tunnel Switching Network Topology

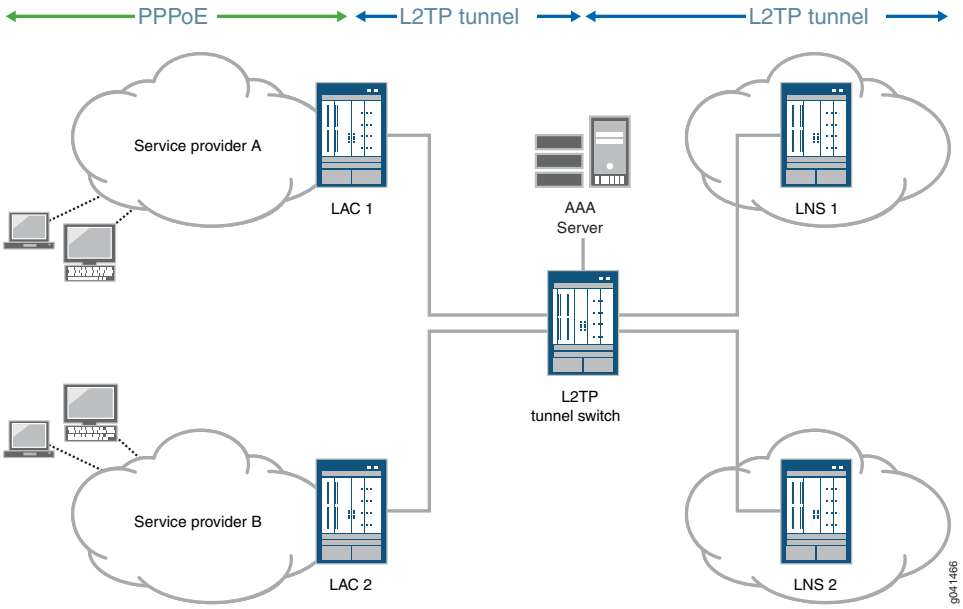

L2TP tunnel switching simplifies network configuration when the administrative domain of a LAC is different from that of the desired LNS. For example:

- The LTS acts as the LNS for multiple LACs. The individual LACs do not have to have the administrative control or capability required to identify the most appropriate LNS on which to terminate their sessions. The LTS performs that function is centralized in the LTS.
- The LTS acts as the LAC for multiple LNSs. When a new remote LAC is added to an ISP's network, the ISP does not have to reconfigure its LNS routers to accommodate the new LAC, because they connect to the LAC on the LTS.

In a Layer 2 wholesale network, the wholesaler can use L2TP tunnel switching to create a flatter network configuration that is easier to manage. The wholesaler bundles Layer 2 sessions from a LAC that are destined for different ISPs—and therefore different LNSs—onto a single L2TP tunnel. This configuration enables a common L2TP control connection to be used for the LAC.

Figure 7 on page 131 shows an example of L2TP tunnel switching for incoming calls with the following sequence of events:

1. The subscriber opens a PPP session to the LAC.
2. The LAC creates the first L2TP tunnel to the LNS configured on the LTS and the first L2TP session to carry the encapsulated PPP packets.
3. During authentication of this first session, the LTS determines whether to retunnel the session to an LNS beyond the LTS, based on the presence or absence of a tunnel switch profile configured on the LTS.

Copyright © 2016, Juniper Networks, Inc.

The tunnel switch profile can be a default profile or it can be applied by the RADIUS server, a domain map configuration, or a tunnel group configuration.

4. If a tunnel switch profile is configured, the LTS creates a second tunnel (if it does not already exist) to the LNS beyond the LTS as specified in the profile and creates the second session in this tunnel.

Figure 7: L2TP Tunnel Switching for Incoming Calls

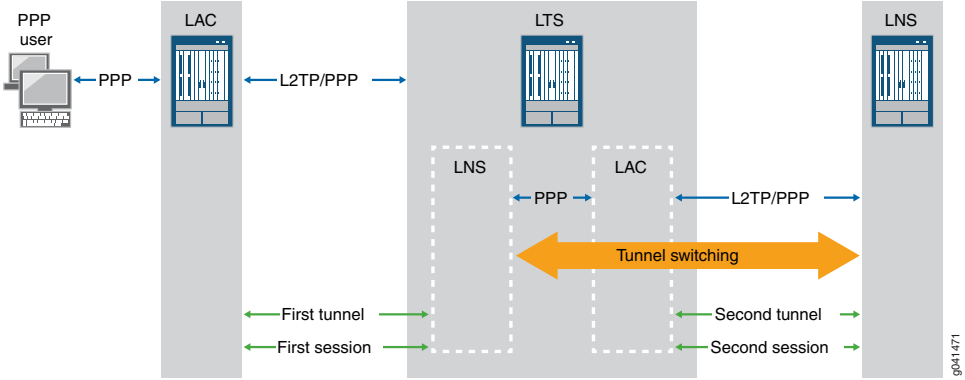

Application of Tunnel Switch Profiles

You can configure a tunnel switch profile to be applied in several ways:

- As a default profile applied globally to traffic received from all LACs
- With a domain map applied to a subscriber session
- With a tunnel group applied to a subscriber session
- In your RADIUS server configuration, returned in the Tunnel Switch-Profile VSA (26-91)

You can configure more than one of these methods of application. When multiple tunnel switch profiles are present, the following order of precedence establishes which profile the LTS uses; the order is from highest (RADIUS) to lowest (default profile):

RADIUS VSA 26-91 > domain map > tunnel group > global tunnel switch profile

The tunnel switch profile must also reference a tunnel profile. This tunnel profile specifies the characteristics of the second tunnel, to which the subscriber packets are switched.

Termination of Tunnel-Switched Sessions on the LTS

Tunnel switched sessions are terminated on the LTS when any of the following happens:

- Either the LAC or LNS interface on the LTS receives a Call-Disconnect-Notify (CDN) message (Table 10 on page 132).

Table 10: Cause of CDN Message

CDN Message Is Received On	When
LAC interface	Either of the following occurs: • The second session cannot be established. • The remote LNS terminates the second session.
LNS interface	Either of the following occurs: • The PPPoE client initiates a logout. • The originating LAC initiates termination of the tunnel

Both the first and second sessions are terminated because the LTS relays the CDN to the interface that did not receive the CDN. The disconnect cause is the same for both sessions.

- Either the LAC or LNS interface on the LTS receives a Stop-Control-Connection-Notification (StopCCN) message (Table 11 on page 132).

Table 11: Cause of StopCCN Message

StopCCN Message Is Received On	When
LAC interface	Either of the following occurs: • The second session cannot be established. • The remote LNS terminates the second tunnel.
LNS interface	The originating LAC initiates termination of the tunnel.

The LTS does not relay the StopCCN message, because a given tunnel can contain both switched and nonswitched sessions. Another reason in a wholesale scenario is that the tunnel ending on the LNS on the LTS can contain sessions from LACs from different providers. Instead, the LTS sends a CDN message to the interface that did not receive the StopCCN to terminate the tunnel-switched session. This CDN relays the error code carried in the StopCCN.

- An administrative **clear** command is issued on the LTS.

Table 12 on page 132 lists the actions taken when an administrative **clear** command is issued on the LTS.

Table 12: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands

Command	LAC or LNS Action	LTS Action
clear services l2tp destination	Clear the destination and all associated tunnels and sessions.	For each switched session in a tunnel to the destination, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.

Copyright © 2016, Juniper Networks, Inc.

Table 12: LAC, LNS, and LTS Actions Taken for Switched Tunnels in Response to Administrative clear Commands *(continued)*

Command	LAC or LNS Action	LTS Action
clear services l2tp destination all	Clear all destinations and all associated tunnels and sessions.	None.
clear services l2tp session	Clear the session.	Clear the corresponding mapped switched session for this session by sending it a CDN message with the cause set to Administrative.
clear services l2tp session all	Clear all sessions.	None.
clear services l2tp tunnel	Clear the tunnel and all its sessions.	For each switched session in the tunnel, clear the corresponding mapped switched session by sending it a CDN message with the cause set to Administrative.
clear services l2tp tunnel all	Clear all tunnels.	None.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 133
- L2TP for Subscriber Access Overview on page 123

Tunnel Switching Actions for L2TP AVPs at the Switching Boundary

When L2TP tunnel switching redirects packets to a different LNS, it performs one of the following default actions at the switching boundary for each AVP carried in the L2TP messages:

- **regenerate**—L2TP regenerates the AVP based on the local policy at the LTS and sent in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—L2TP transparently forwards the AVP in the switched packet with no alteration.

Table 13 on page 134 lists the default action for each AVP. Mandatory AVPs are always included in the L2TP messages from the LAC; optional AVPs might be included in the messages.

You can optionally override the default action taken at the switching boundary for the Bearer Type AVP (18), Calling Number AVP (22), or Cisco NAS Port Info AVP (100). You can configure any of these three AVPs to be dropped from the switched packets or regenerated, or you can restore the default relay action.

NOTE: **L2TP AVPs that have their attribute values hidden are always regenerated at the switching boundary. The value is decoded and sent in clear text when the packet is forwarded to the remote LNS.**

Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary

AVP Name	AVP Type	L2TP Message Type	Default Action
Assigned Session Id	Mandatory	ICRQ, CDN	Regenerate
Assigned Tunnel Id	Mandatory	SCCRQ, CDN	Regenerate
Bearer Capabilities	Optional	SCCRQ	Regenerate
Bearer Type	Optional	ICRQ	Relay
Call Serial Number	Mandatory	ICRQ	Relay
Called Number	Optional	ICRQ	Relay
Calling Number	Optional	ICRQ	Relay
Challenge	Optional	SCCRQ	Regenerate
Challenge Response	Optional	SCCCN	Regenerate
Cisco NAS Port	Optional	ICRQ	Relay
Failover Capability	Optional	SCCRQ	Regenerate
Firmware Revision	Optional	SCCRQ	Regenerate
Framing Capabilities	Mandatory	SCCRQ	Regenerate
Framing Type	Mandatory	ICCN	Relay
Host Name	Mandatory	SCCRQ	Regenerate
Message Type	Mandatory	All	Regenerate
Physical Channel Id	Optional	ICRQ	Regenerate
Private Group Id	Optional	ICCN	Relay
Protocol Version	Mandatory	SCCRQ	Regenerate
Proxy Authentication	Optional	ICCN	Relay if acceptable; otherwise regenerate
Proxy LCP	Optional	ICCN	Relay if acceptable; otherwise regenerate
Receive Window Size	Optional	SCCRQ	Regenerate
Rx Connect Speed	Optional	ICCN	Relay

Copyright © 2016, Juniper Networks, Inc.

Table 13: Default Action for Handling L2TP AVPs at the Switching Boundary *(continued)*

AVP Name	AVP Type	L2TP Message Type	Default Action
Sequencing Required	Optional	ICCN	Regenerate
Sub-Address	Optional	ICRQ	Relay
Tie Breaker	Optional	SCCRQ	Regenerate
Tunnel Recovery	Optional	SCCRQ	Regenerate
Tx Connect Speed	Mandatory	ICCN	Relay
Vendor Name	Optional	SCCRQ	Regenerate

Related Documentation

- L2TP Tunnel Switching Overview on page 129
- Configuring L2TP Tunnel Switching on page 135

Configuring L2TP Tunnel Switching

L2TP tunnel switching enables a router configured as an LTS to forward PPP packets carried on one L2TP session to a second L2TP session terminated on a different LNS. To configure L2TP tunnel switching, you must define a tunnel switch profile and then assign that profile.

You can configure tunnel switch profiles for all sessions globally, all sessions in a tunnel group, all sessions in a domain or in your RADIUS server configuration to be returned in the RADIUS Tunnel Switch-Profile VSA (26-91). The order of precedence for tunnel switch profiles from various sources is as follows:

- RADIUS VSA 26-91 > domain map > tunnel group > global tunnel switch profile

To define an L2TP tunnel switch profile:

1. Create the profile.

   [edit access]
   user@host# **edit** tunnel-switch-profile *profile-name*

2. (Optional) Override the default actions taken for certain L2TP AVPs at the switching boundary.

   [edit access tunnel-switch-profile *profile-name*]
   user@host# **set** avp bearer-type *action*
   user@host# **set** avp calling-number *action*
   user@host# **set** avp cisco-nas-port-info *action*

3. Specify the tunnel profile that defines the tunnel to which the subscriber traffic is switched.

   [edit access tunnel-switch-profile *profile-name*]
   user@host# **set** tunnel-profile *profile-name*

4. (Optional) Apply the profile as a global default profile to switch packets from all incoming sessions from the LAC.

```
[edit services l2tp]
user@host1# set tunnel-switch-profile profile-name
```

5. (Optional) Apply the profile as part of a tunnel group to switch packets from all sessions in the tunnel group.

```
[edit services l2tp tunnel-group name]
user@host1# set tunnel-switch-profile profile-name
```

NOTE: The tunnel group is part of the LTS configuration that enables it to act as the LNS for the original sessions from the LAC.

A tunnel group with a tunnel switch profile must also contain a dynamic profile, because tunnel switching supports only dynamic subscribers.

6. (Optional) Apply the profile as part of a domain map to switch packets from all sessions that are associated with the domain.

```
[edit access domain map domain-map-name]
user@host1# set tunnel-switch-profile profile-name
```

NOTE: A domain map cannot have both a tunnel switch profile and a tunnel profile. You must remove one if you add the other.

7. (Optional) Apply the profile by means of the Tunnel-Switch-Profile VSA [26–91] in the RADIUS Access-Accept message returned when the session from the LAC is authenticated. Refer to the documentation for your RADIUS server to determine how to configure this method.

For example, consider the following configuration. which creates three tunnel switch profiles, l2tp-tunnel-switch-profile, lts-profile-groupA, and lts-profile-example-com:

```
[edit access tunnel-switch-profile l2tp-tunnel-switch-profile]
user@host# set avp bearer-type regenerate
user@host# set avp calling-number regenerate
user@host# set avp cisco-nas-port-info drop
user@host# set tunnel-profile l2tp-tunnel-profile1

[edit access tunnel-switch-profile lts-profile-groupA]
user@host# set tunnel-profile l2tp-tunnel-profile2
[edit access tunnel-switch-profile lts-profile-example.com]
user@host# set tunnel-profile l2tp-tunnel-profile3

[edit services l2tp]
user@host1# set tunnel-switch-profile l2tp-tunnel-switch-profile
user@host1# set tunnel-group groupA tunnel-switch-profile lts-profile-groupA

[edit access domain]
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host1# set map example.com tunnel-switch-profile lts-profile-example.com
```

The profile l2tp-tunnel-switch-profile is applied as the global default. When packets are switched according to this profile, the values for the Bearer Type AVP (18) and Calling Number AVP (22) in the L2TP packets are regenerated based on local policy at the L2TP tunnel switch and then sent with the packets. The Cisco NAS Port Info AVP (100) is simply dropped. Finally, l2tp-tunnel-profile1 provides the configuration characteristics of the tunnel to which the traffic is switched.

Tunnel switch profile lts-profile-groupA is applied by means of a tunnel group, groupA; it specifies a different tunnel profile, l2tp-tunnel-profile2 and it does not override any AVP actions. Tunnel switch profile lts-profile-example.com is applied by means of a domain map for the example.com domain; it specifies a different tunnel profile, l2tp-tunnel-profile3 and it does not override any AVP actions.

Related Documentation

- L2TP Tunnel Switching Overview on page 129
- Tunnel Switching Actions for L2TP AVPs at the Switching Boundary on page 133

Setting the L2TP Receive Window Size

You can configure the L2TP receive window size for an L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

By default, the receive window size is set to four packets. If the receive window size is set to its default value, the router does not send the Receive Window Size AVP, AVP 10, in its first packet sent during tunnel negotiation to its peer.

To configure the receive window size:

```
[edit services l2tp tunnel]
user@host# set rx-window-size packets
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Setting the L2TP Tunnel Idle Timeout

You can configure the LAC or the LNS to specify how long a tunnel without any sessions remains active. The idle timer starts when the last session on the tunnel is terminated. When the timer expires the tunnel is disconnected. This idle timeout frees up resources otherwise consumed by inactive tunnels.

If you set the idle timeout value to zero, the tunnel is forced to remain active indefinitely after the last session is terminated until one of the following occurs:

- You issue the **clear services l2tp tunnel** command.
- The remote peer disconnects the tunnel.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no idle-timeout** statement at the **[edit services l2tp tunnel]** hierarchy level.

To set the tunnel idle timeout:

- Configure the timeout period.

```
[edit services l2tp tunnel]
user@host# set idle-timeout seconds
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Setting the L2TP Destruct Timeout

You can configure the LAC or the LNS to specify how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed. This destruct timeout aids debugging and other analysis by saving underlying memory structures after the destination, tunnel, or session is terminated. Any specific dynamic destination, tunnel, or session may not be maintained for this entire time period if the resources must be reclaimed early to allow new tunnels to be established.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no destruct-timeout** statement at the **[edit services l2tp]** hierarchy level.

To set the L2TP destruct timeout:

- Configure the timeout period.

```
[edit services l2tp]
user@host# set destruct-timeout seconds
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Copyright © 2016, Juniper Networks, Inc.

Configuring the L2TP Destination Lockout Timeout

When multiple sets of tunneling parameters are available, L2TP uses a selection process to choose the best tunnel for subscriber traffic. As part of this selection process, L2TP locks out destinations it cannot connect to when a subscriber tries to reach a domain. L2TP places the destination on the destination lockout list and excludes the destination from consideration for a configurable period called the *destination lockout timeout*.

By default, the destination lockout timeout is 300 seconds (5 minutes). You can configure a value from 60 through 3600 seconds (1 minute through 1 hour). When the lockout timeout expires, L2TP assumes that the destination is now available and includes the destination when performing the tunnel selection process. The destination lockout period is a global value and is not individually configurable for particular destinations, tunnels, or tunnel groups.

BEST PRACTICE: **Configure the lockout timeout to be equal to or shorter than the destruct timeout. Otherwise, the destruct timeout expires before the lockout timeout. In this event, the locked-out destination is destroyed and can be subsequently returned to service before the lockout timeout expires, thus negating the effectiveness of the lockout timeout.**

To configure the destination lockout timeout:

- Specify the period in seconds.

```
[edit services l2tp destination]
user@host# set lockout-timeout seconds
```

The **show services l2tp destination lockout** command displays the destination lockout list and for each destination indicates how much time remains before its timeout expires. The **show services l2tp destination detail** command indicates for each destination whether it is locked and waiting for the timeout to expire or not locked.

Related Documentation

- LAC Tunnel Selection Overview on page 157
- Setting the L2TP Destruct Timeout on page 138
- Removing an L2TP Destination from the Destination Lockout List on page 139
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Removing an L2TP Destination from the Destination Lockout List

When a PPP subscriber tries to log in to a domain, L2TP selects a tunnel associated with a destination in that domain and attempts to access the destination. If the connection attempt fails, L2TP places the destination on the destination lockout list. Destinations on this list are excluded from being considered for subsequent connections for a configurable period called the *destination lockout timeout*.

Copyright © 2016, Juniper Networks, Inc.

You can issue the **request services l2tp destination unlock** command for a particular destination to remove it from the destination lockout list. The result is that this destination is immediately available for consideration when a subscriber logs in to the associated domain.

To remove a destination from the destination lockout list:

- Specify the name of the destination to be unlocked.

  user@host> **request services l2tp destination unlock** ***destination-name***

Related Documentation

- LAC Tunnel Selection Overview on page 157
- Configuring the L2TP Destination Lockout Timeout on page 139

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 17

Configuring L2TP Control Messages for Subscribers

Retransmission of L2TP Control Messages

L2TP peers maintain a queue of control messages that must be sent to the peer device. After the local peer (LAC or LNS) sends a message, it waits for a response from the remote peer. If a response is not received, the local peer retransmits the message. This behavior allows the remote peer more time to respond to the message.

You can control the retransmission behavior in the following two ways:

- Retransmission count—You can configure how many times an unacknowledged message is retransmitted by the local peer. Increasing the count provides more opportunities for the remote peer to respond, but also increases the amount of control traffic. For tunnels that have been established, include the **retransmission-count-established** statement at the **[edit services l2tp tunnel]** hierarchy level. For tunnels that are not yet established, include the **retransmission-count-not-established** statement.
- Retransmission interval—You can configure how long the local peer waits for the first response to a control message. If a response is not received within the first timeout interval, then the retransmission timer doubles the interval between each successive retransmission up to a maximum of 16 seconds. Increasing the interval gives the remote peer more time to respond, but also spends more resources on a potentially unavailable peer. Include the **minimum-retransmission-interval** statement at the **[edit services l2tp tunnel]** hierarchy level.

The local peer continues retransmitting the control message until one of the following occurs:

- A response is received within the current waiting period.
- The maximum retransmission count is reached.

If the maximum count is reached and no response has been received, the tunnel and all its sessions are cleared.

Copyright © 2016, Juniper Networks, Inc.

NOTE: **Reaching the maximum interval of 16 seconds does not halt retransmissions. The local peer continues to wait 16 seconds after each subsequent retransmission.**

The following examples describe the retransmission behavior in different circumstances:

- Example 1—The retransmission count is three and the minimum retransmission interval is 1 second.
  1. The local peer sends a control message.
  2. The local peer waits 1 second, but receives no response.
  3. The local peer retransmits the control message. This is the first retransmission.
  4. The local peer waits 2 seconds, but receives a response before the interval expires.
  5. Retransmission stops because a response is received within the interval.
- Example 2—The retransmission count is two and the minimum retransmission interval is 8 seconds.
  1. The local peer sends a control message.
  2. The local peer waits 8 seconds, but receives no response.
  3. The local peer retransmits the control message. This is the first retransmission.
  4. The local peer waits 16 seconds, but receives no response.
  5. The local peer retransmits the control message. This is the second retransmission.
  6. The local peer again waits 16 seconds, because the interval cannot increase beyond 16, but receives no response.
  7. Retransmission stops because the maximum retransmission count of two was reached.
  8. The tunnel and all its sessions are cleared.

Related Documentation

- Configuring Retransmission Attributes for L2TP Control Messages on page 143
- L2TP for Subscriber Access Overview on page 123

Copyright © 2016, Juniper Networks, Inc.

Configuring Retransmission Attributes for L2TP Control Messages

You can control the retransmission of unacknowledged L2TP control messages by configuring how many times the local peer retransmits the message and how long it waits for a response before retransmission.

L2TP peers maintain a queue of control messages that must be sent to the peer device. After the local peer (LAC or LNS) sends a message, it waits for a response from the remote peer. If a response is not received within the minimum retransmission interval, the local peer retransmits the message and waits for double the retransmission interval. Each time it retransmits the message, the peer doubles how long it waits, up to a maximum of 16 seconds.

If no response is received, the local peer continues to send the message until the number of retransmissions matches the retransmission count. In this case, retransmissions stop and the tunnel and all its sessions are cleared.

BEST PRACTICE: Before you downgrade to a Junos OS Release that does not support these statements, we recommend that you explicitly unconfigure the feature by including the **no retransmission-count-established** statement and the **no retransmission-count-non-established** statement at the **[edit services l2tp tunnel]** hierarchy level.

BEST PRACTICE: During a unified in-service software upgrade (unified ISSU) on an MX Series router configured as the LAC, the LAC does not respond to control messages from the LNS. This can result in dropping LAC L2TP sessions. You can avoid this situation by ensuring that the maximum retransmission count on the LNS is set to 16 or higher.

To set the maximum retransmission count for established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-established count
```

To set the maximum retransmission count for non-established tunnels:

- Configure the count.

```
[edit services l2tp tunnel]
user@host# set retransmission-count-not-established count
```

To set the minimum interval between retransmissions:

- Configure the interval.

```
[edit services l2tp tunnel]
user@host# set minimum-retransmission-timeout seconds
```

For example, the following configuration specifies that established tunnels have a maximum retransmission count of three and a minimum retransmission interval of two seconds:

```
[edit services l2tp tunnel]
user@host# set retransmission-count-established 3
user@host# set minimum-retransmission-timeout 2
```

With this sample configuration, the following sequence applies to each control message sent by the LAC or LNS:

1. The local peer sends the control message and waits for a response from the remote peer.
2. If the response is not received within the minimum interval of 2 seconds, the local peer retransmits the message. This is the first retransmission.
3. If the response is not received within 4 seconds, the local peer retransmits the message. This is the second retransmission.
4. If the response is not received within 8 seconds, the local peer retransmits the message. This is the third and final retransmission, because the maximum count has been reached.
5. If the response is not received within 16 seconds, the tunnel and all its sessions are cleared.

Related Documentation

- Retransmission of L2TP Control Messages on page 141
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 18

Configuring L2TP LAC Subscribers

Subscriber Access Line Information Forwarding by the LAC Overview

L2TP supports a set of AVPs that convey information about subscriber access lines from the LAC to the LNS. The information originates from an ANCP access node (DSLAM) and is distributed to the LAC by means of either DSL Forum VSAs in ANCP messages or PPPoE intermediate agent tags included in the PPPoE PADI and PADR messages.

In the network topology shown in Figure 8 on page 146, when a subscriber initiates a connection through the CPE, the DSLAM relays the subscriber's PPPoE session to the router configured as a LAC. When the router has established the PPPoE session, the LAC initiates an L2TP tunnel to forward the subscriber's encapsulated PPP packets into the provider network.

In parallel to the PPPoE session, an ANCP connection between the DSLAM and the ANCP agent on the router conveys information about the subscriber's local loop as well as the link speeds of the PPPoE sessions on the local loop. The DSLAM sends the router Agent Circuit Id (ACI) and Agent Remote Id (ARI) strings that uniquely identify the DSLAM's receiving interface; this information is encoded in the ANCP Port Up and Port Down messages as Access Line Identifying TLVs. The ANCP messages can also include line attributes such as minimum, maximum, and actual net upstream and downstream data rates in the DSL Line Attributes TLV. The DSLAM can also send the access line attributes in vendor-specific tags that it inserts in the PADI and PADR messages.

Copyright © 2016, Juniper Networks, Inc.

Figure 8: Sample L2TP Network Topology

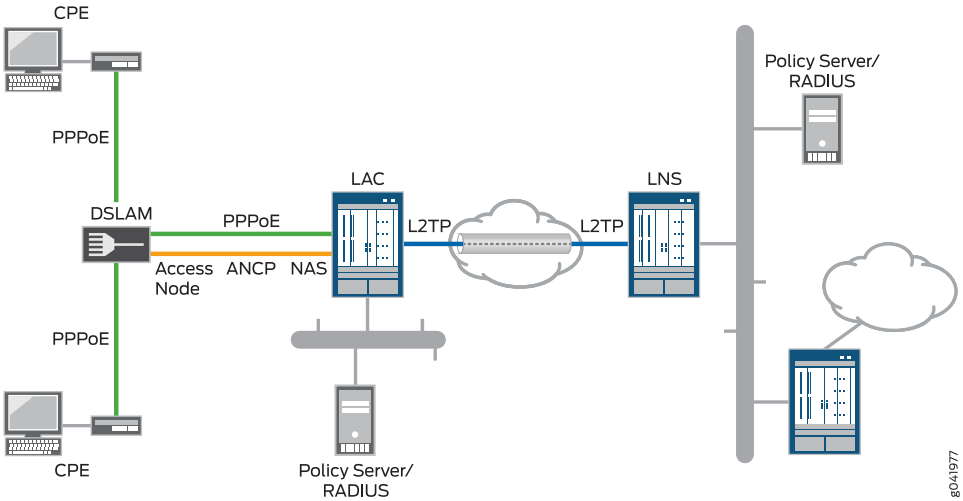

You can include the **access-line-information** statement at the **[edit services l2tp destination *ip-address*]** hierarchy level to configure the LAC to forward this access line information in the ICRQ message that it sends to the LNS. L2TP supports the AVPs listed in Table 14 on page 146 to carry this information. The access line information is not required for the L2TP session to be initiated, and the establishment of that session is not delayed waiting for the values to be sent from the DSLAM.

Table 14: L2TP AVPs That Provide Subscriber Access Line Information

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Actual Data Rate Downstream	130 (26–130)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Actual Data Rate Upstream	129 (26–129)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Actual Interleaving Delay Downstream	142 (26–142)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Actual Interleaving Delay Upstream	140 (26–140)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Access Loop Encapsulation	144 (26–144)	Three one-octet encodings for data link, encapsulation 1, and encapsulation 2.	ICRQ
Agent Circuit ID	1 (26–1)	2-63 octet string; ACI of the logical access loop port on the DSLAM/access node.	ICRQ

Copyright © 2016, Juniper Networks, Inc.

Table 14: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Agent Remote ID	2 (26–2)	2-63 octet statically configured string; uniquely identifies subscriber on the DSLAM (access node).	ICRQ
ANCP Access Line Type	145 (none)	One octet encoding for transmission system type, followed by three MBZ (must be zero) octets (total 4 bytes). This value is not supplied in the ICRQ when the access line parameters are sourced from PPPoE-IA, because the ANCP-sourced information may not be immediately available.	ICRQ
Attainable Data Rate Upstream	133 (26–133)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Attainable Data Rate Downstream	134 (26–134)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Connect Speed Update Enable	98 (none)	Value does not matter: presence indicates support for CSUN, CSURQ message types for this session.	ICRQ
Connect Speed Update	97 (none)	Data structure listing remote session id and the current transmit and receive connection speeds in bits per second.	CSUN, CSURQ
IWF Session	254 (26–254)	Four-octet field indicating whether or not the internetworking function has been performed.	ICRQ
Maximum Data Rate Downstream	136 (26–136)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Maximum Data Rate Upstream	135 (26–135)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Maximum Interleaving Delay Downstream	141 (26–141)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ

Copyright © 2016, Juniper Networks, Inc.

Table 14: L2TP AVPs That Provide Subscriber Access Line Information *(continued)*

Attribute Value Pair	AVP Type (Corresponding DSL Forum VSA)	Description	L2TP Message Type
Maximum Interleaving Delay Upstream	139 (26–139)	32-bit unsigned integer; maximum delay in milliseconds.	ICRQ
Minimum Data Rate Downstream	132 (26–132)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Downstream Low Power	138 (26–138)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Upstream	131 (26–13)	64-bit unsigned integer; data rate in bits per sec.	ICRQ
Minimum Data Rate Upstream Low Power	137 (26–137)	64-bit unsigned integer; data rate in bits per sec.	ICRQ

You can also configure the LAC to notify the LNS when the speed of the subscriber connection changes from the values initially communicated to the LNS by AVP 24 and AVP 38 in ICCN messages. To do so, specify the **connection-speed-update** option with the **access-line-information** statement. When configured, the LAC informs the LNS that it can send these updates by including the Connect Speed Update Enable AVP (98) in the ICRQ message when the L2TP session starts up.

If the connection speed changes, the DSLAM notifies the ANCP agent. The ANCP agent then notifies the LAC, and the LAC in turn relays this information to the LNS by sending a Connect-Speed-Update-Notification (CSUN) message that includes a Connect Speed Update AVP (97) for each session. The LAC collects connection speed updates and sends them in a batch to minimize both the performance overhead on the LAC and the amount of traffic generated as a result of these notifications.

The presence of the Connect Speed Update Enable AVP (98) in the ICRQ message also informs the LNS that the LAC does respond if the LNS sends it a Connect-Speed-Update-Request (CSURQ) message. A third-party LNS can send this message type at any time during the life of a tunnel to request the current transmit and receive connection speed for one or more L2TP sessions. The LNS includes the session IDs in the CSURQ message. The LNS sends the CSURQ as part of failover recovery, for example, when the LAC restarts. If the LAC has previously sent the Connect Speed Update Enable AVP (98) for the requested sessions, then it responds with a CSUN message that includes the Connect Speed Update AVP (97) for each session. If no changes to connection speeds have occurred by this time, the LAC simply includes the initial connection speed values that were reported in AVP 24 and AVP 38.

The absence of the Connect Speed Update Enable AVP (98) in the ICRQ message indicates that the LAC does not send updates for the life of the session. The LNS does not send a CSURQ to a LAC unless it has received AVP 98 from that LAC.

Copyright © 2016, Juniper Networks, Inc.

BEST PRACTICE: When you enable connection speed updates, we recommend that you also specify **ancp** with the **tx-connect-speed** statement as the source for the initial connection speeds reported in AVP 24 and AVP 38. This configuration aligns the initial source of connection speed with the preferred update source. The ANCP agent is the preferred source for both connection speed updates and access line characteristics.

If ANCP is not intended to be used for connection speed updates to the LAC, then we recommend that you do not enable connection speed updates. In the absence of access line information from the ANCP agent, the LAC forwards the speeds from the PPPoE-IA tags in the PPPOE PADR message. However, because these are only initial speed values and do not change, enabling connection speed updates is of no value,

Related Documentation

- Configuring the LAC to Report Access Line Information to the LNS on page 154
- L2TP for Subscriber Access Overview on page 123
- *DSL Forum Vendor-Specific Attributes*
- Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS on page 169

Configuring an L2TP LAC

To configure an L2TP LAC:

1. Configure a tunnel profile to apply to subscribers.

   See "Configuring a Tunnel Profile for Subscriber Access" on page 162.

2. (Optional) Configure the method used for selecting among multiple tunnels.

   - See "Configuring the L2TP LAC Tunnel Selection Parameters" on page 165.
   - See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 166.
   - See "Configuring LAC Tunnel Selection Failover Within a Preference Level" on page 165.

3. (Optional) Configure the LAC to not send Calling Number AVP 22 to the LNS.

   See "Preventing the LAC from Sending Calling Number AVP 22 to the LNS" on page 172.

4. (Optional) Specify the method for setting the transmit and receive connect speeds.

   See "Configuring the Method to Set the LAC Connection Speeds to the LNS" on page 172.

5. (Optional) Disable negotiation of the L2TP failover protocol to force use of only the silent failover resynchronization mechanism.

   See "Preventing the LAC from Negotiating L2TP Failover Protocol" on page 150.

6. (Optional) Specify the format for the tunnel name.

Copyright © 2016, Juniper Networks, Inc.

"Setting the Format for the Tunnel Name" on page 162.

7. (Optional) Specify when and how many times L2TP retransmits unacknowledged control messages.

   See "Configuring Retransmission Attributes for L2TP Control Messages" on page 143.

8. (Optional) Specify how long a tunnel can remain idle before being torn down.

   See "Setting the L2TP Tunnel Idle Timeout" on page 137.

9. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

   See "Setting the L2TP Receive Window Size" on page 137

10. (Optional) Specify how long the L2TP retains information about terminated dynamic tunnels, sessions, and destinations.

    See "Setting the L2TP Destruct Timeout" on page 138.

11. (Optional) Specify how the LAC handles IP address or UDP port change requests.

    See "Configuring the LAC to Ignore Address and Port Changes Requested by the LNS" on page 151

12. (Optional) Configure all tunnels on the LAC for interoperation with Cisco LNS devices.

    See "Globally Configuring the LAC to Interoperate with Cisco LNS Devices" on page 153.

13. (Optional) Specify that the LAC sends information to the LNS about subscriber access lines.

    See "Configuring the LAC to Report Access Line Information to the LNS" on page 154.

14. (Optional) Enable SNMP statistics counters.

    See "Enabling Tunnel and Global Counters for SNMP Statistics Collection" on page 214

15. (Optional) Configure trace options for troubleshooting the configuration.

    See "Tracing L2TP Operations for Subscriber Access" on page 249

Related Documentation

- L2TP for Subscriber Access Overview on page 123

Preventing the LAC from Negotiating L2TP Failover Protocol

The L2TP LAC implementation on MX Series routers supports L2TP failover and peer resynchronization with a failed remote endpoint. The LAC supports both the L2TP failover protocol method and the L2TP silent failover method. By default, L2TP on the LAC attempts to negotiate the L2TP failover protocol with the LNS. When negotiation determines that the LNS supports this method, then the LAC uses L2TP failover protocol if the LNS fails. When the LNS does not support L2TP failover protocol, then the LAC uses silent failover in the event of an LNS failure. The ability to fall back on silent failover prevents the failover from forcing a disconnection of the tunnel to the peer and all the associated sessions.

 Copyright © 2016, Juniper Networks, Inc.

You can disable the default behavior to force the LAC to operate only in silent failover mode. This configuration can be useful when routers that act as the LNS either are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it. However, when you issue this statement and the LNS supports only failover protocol, then the LAC cannot negotiate failover protocol, and recovery (failover protocol recovery initiated by the LNS) always fails.

To disable negotiation of the L2TP failover protocol:

- Configure disabling.

  [edit services l2tp]
  user@host# **set** disable-failover-protocol

Related Documentation

- Configuring an L2TP LAC on page 149

Configuring the LAC to Ignore Address and Port Changes Requested by the LNS

By default, when the LAC receives a request from the LNS in an SCCRP message to change the destination IP address or the UDP port, the LAC accepts the request and makes the change. If this is not the desired behavior, you can use the **tx-address-change** statement to configure how the LAC handles these change requests.

To configure how the LAC handles change requests for the IP address, the UDP port, or both:

- (Optional) Specify the LAC to ignore all change requests.

  [edit services l2tp tunnel]
  user@host# **set** tx-address-change **ignore**

- (Optional) Specify the LAC to ignore only change requests for the IP address.

  [edit services l2tp tunnel]
  user@host# **set** tx-address-change **ignore-ip-address**

- (Optional) Specify the LAC to ignore only change requests for the UDP port.

  [edit services l2tp tunnel]
  user@host# **set** tx-address-change **ignore-udp-port**

- (Optional) Specify the LAC to accept all change requests.

  [edit services l2tp tunnel]
  user@host# **set** tx-address-change **accept**

Use the **show services l2tp summary** command to display the current behavior of the LAC:

show services l2tp summary

```
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Disabled
Tx Connect speed method is static
Rx speed avp when equal is Disabled
Tunnel assignment id format is assignment-id
```

```
  Tunnel Tx Address Change is Ignore IP Address Change
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
Destinations: 1, Tunnels: 0, Sessions: 0
```

Related Documentation

- Configuring an L2TP LAC on page 149

LAC Interoperation with Third-Party LNS Devices

In some network environments, the LAC may need to interoperate with an LNS configured on a device from another vendor that does not run Junos OS. Interoperation with Cisco Systems devices requires the LAC to communicate a NAS port type, but the LAC does not provide this information by default.

You can enable interoperation with Cisco Systems devices by configuring the NAS port method as **cisco-avp**, which causes the LAC to include the Cisco Systems NAS Port Info AVP (100) when it sends an incoming call request (ICRQ) to the LNS. The AVP includes information that identifies the NAS port and indicates whether the port type is ATM or Ethernet.

You can configure the NAS port method globally for all tunnels on the LAC or in a tunnel profile for only the tunnels instantiated by the profile.

You can also include the Tunnel-Nas-Port-Method VSA [26–30] in your RADIUS server configuration with the value set to 1 to indicate Cisco Systems CLID. In this case, RADIUS can override the global value by modifying or creating a tunnel profile. The RADIUS configuration has precedence over the tunnel profile configuration, which in turn has precedence over the global LAC configuration.

If the LNS receiving the AVP is an MX Series router instead of a Cisco Systems device, the LNS simply ignores the AVP, unless the LNS is configured for L2TP tunnel switching. In that case, the LNS preserves the value of the AVP and passes it along when it switches tunnels for the LAC.

Related Documentation

- Globally Configuring the LAC to Interoperate with Cisco LNS Devices on page 153
- Configuring a Tunnel Profile for Subscriber Access on page 162
- Configuring an L2TP LAC on page 149
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- L2TP for Subscriber Access Overview on page 123

Copyright © 2016, Juniper Networks, Inc.

Globally Configuring the LAC to Interoperate with Cisco LNS Devices

Cisco LNS devices require from the LAC both the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM. By default, the LAC does not include this information. You can globally configure the LAC to provide this information by including the NAS Port Info AVP (100) in the ICRQ that it sends to the LNS. This configuration enables the LAC to interoperate with a Cisco LNS.

To globally configure the LAC to include the NAS Port Info AVP:

- Specify the NAS port method.

  [edit services l2tp tunnel]
  user@host# **set nas-port-method cisco-avp**

NOTE: This global configuration for the LAC can be overridden by the configuration in a tunnel profile or RADIUS.

Use the **show services l2tp tunnel extensive** command to display the current behavior of the LAC:

show services l2tp tunnel extensive

```
Tunnel local ID: 51872, Tunnel remote ID: 8660
    Remote IP: 192.168.1.3:1701
    Sessions: 5, State: Established
    Tunnel Name: 1/tunnel-test-2
    Local IP: 10.1.1.2:1701
    Local name: testlac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Thu Jul 25 12:55:41 2013, Up time: 11:18:14
    Idle time: 00:00:00
    Statistics since: Thu Jul 25 12:55:41 2013
                        Packets        Bytes
      Control Tx            702        15.5k
      Control Rx            690         8.5k
      Data Tx            153.3k         6.6M
      Data Rx            126.3k         5.9M
      Errors Tx               0
      Errors Rx               0
```

Related Documentation

- LAC Interoperation with Third-Party LNS Devices on page 152
- Configuring a Tunnel Profile for Subscriber Access on page 162
- Configuring an L2TP LAC on page 149

Copyright © 2016, Juniper Networks, Inc.

Configuring the LAC to Report Access Line Information to the LNS

The L2TP AVP extensions defined in RFC 5515, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extension*, enable the LAC to report to the LNS characteristics of the subscriber's access line, such as identification attributes, line type, connection speed, various data rates, and so on. The LAC receives the access line information when the subscriber's CPE initiates a connection request, and forwards the available information in various AVPs included in ICRQ messages to the LNS. The LAC can also signal to the LNS that it is capable of sending updates about the subscriber connection speeds in the CSUN message.

By default, neither the access line information forwarding or connection speed update capability are enabled on the LAC. You must configure the capabilities for a specific LNS endpoint, which enables them for all tunnels with that destination IP address.

NOTE: **Do not configure the connection-speed-update option when the LNS does not support connection speed changes. Configuring the option for such an LNS generates additional control messages that are ignored.**

To configure the LAC to send information about subscriber access lines to the LNS:

- Configure the capability.

  [edit services l2tp destination *ip-address*]
  user@host# **set** access-line-information

Alternatively, to configure the LAC to also send updates to the LNS about changes in connection speed:

- Include the update option when you configure the capability.

  [edit services l2tp destination *ip-address*]
  user@host# **set** access-line-information **connection-speed-update**

- Typically, when you configure the LAC for updates, you also configure the ANCP agent as the source for the initial connection speed.

  [edit services l2tp]
  user@host# **set** tx-connect-speed-method **ancp**

Consider the following examples:

- In this example, for all tunnels with an endpoint address of 192.168.1.2, the LAC reports access line characteristics sourced from the ANCP agent or the PPPoE intermediate agent (in that order) to the LNS in the ICRQ message. The Connect Speed Update Enable AVP is not included in the ICRQ; consequently no CSUN messages are sent to the LNS to report speed changes in the subscriber access lines reported by the ANCP agent. The LAC ignores any CSURQ messages that it receives from the LNS.

  [edit services l2tp destination *192.168.1.2*]
  user@host# **set** access-line-information

Copyright © 2016, Juniper Networks, Inc.

- In this example, for all tunnels with an endpoint address of 10.2.1.3, the LAC reports access line characteristics sourced from the ANCP agent or the PPPoE intermediate agent (in that order) to the LNS in the ICRQ message. The Connect Speed Update Enable AVP is included in the ICRQ; and CSUN messages are sent to the LNS to report speed changes in the subscriber access lines reported by the ANCP agent. The LAC accepts any CSURQ messages that it receives from the LNS and responds with a CSUN message.

```
[edit services l2tp destination 10.2.1.3]
user@host# set access-line-information connection-speed-update
user@host# up
user@host# set tx-connect-speed-method ancp
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Subscriber Access Line Information Forwarding by the LAC Overview on page 145

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 19

Configuring L2TP LAC Tunneling for Subscribers

LAC Tunnel Selection Overview

L2TP enables you to specify:

- Up to 31 destinations for a domain.
- Up to eight levels of preference. Preference indicates the order in which the LAC attempts to connect to the destinations specified for a domain.

  NOTE: **Zero (0) is the highest level of preference; this is the most-preferred level.**

- Up to 31 destinations for a single preference level.

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from a set of tunnels associated with either the PPP user or the PPP user's domain. The LAC provides the following methods for selecting tunnels:

- Tunnel selection failover between preference levels (the default behavior)
- Tunnel selection failover within a preference level
- Maximum sessions per tunnel
- Weighted load balancing

Copyright © 2016, Juniper Networks, Inc.

Tunnel Selection Failover Between Preference Levels

When a user tries to log in to a domain in a default configuration—that is, when failover within preference and weighted load balance are not configured—the LAC attempts to connect to a destination in that domain by means of the associated tunnel at the highest preference level. If more than one tunnel is available to destinations in the domain at that preference level, the LAC randomly selects one of the tunnels and attempts to connect to the destination of the selected tunnel. If the LAC is unsuccessful, it marks the destination as unreachable—*locks out the destination*—and does not try to connect to that destination for a period of time called the *destination lockout timeout*. The timeout is configurable and has a default value of five minutes.

The LAC makes only one attempt to connect to a destination at any level. If the attempt fails, the LAC drops to the next lower level to attempt a connection through a tunnel at the lower level. If the attempts fail at every level, the LAC rejects the login attempt. If the login attempt is then repeated, the LAC restarts the process at the highest preference level.

If the LAC discovers that all the destinations are locked out as it cycles through the levels, then the LAC behaves differently when it gets to the lowest preference level. The LAC selects the tunnel at this level that has the destination with the shortest lockout time remaining from among all the locked-out destinations at the lowest level. The LAC ignores the lockout for this destination and attempts to connect to it. If this attempt fails, then the LAC rejects the login attempt by the PPP client.

NOTE: **More than one tunnel may be able to reach a destination, and those tunnels can have the same preference level or different preference levels.**

For example, suppose that there are four destinations for a domain and all the destinations are considered reachable (none of them are locked out). Tunnels have been defined for the destinations and assigned preference levels as follows:

- Preference 0, tunnel 1, destination A
- Preference 1, tunnel 2, destination B
- Preference 2, tunnel 3, destination C
- Preference 2, tunnel 4, destination D
- Preference 2, tunnel 5, destination B

When a PPP user tries to connect to the domain, the LAC acts as follows:

1. At the highest preference level, only tunnel 1 has a destination in the domain. The LAC selects tunnel 1 and attempts to reach destination A.
2. This connection attempt fails, so the LAC locks out destination A; that is, it excludes destination A from consideration for the length of the destination lockout timeout.

Copyright © 2016, Juniper Networks, Inc.

3. The LAC drops to the next level, preference level 1, to reach a destination for the domain. The LAC selects tunnel 2, the only tunnel at this preference level that has a destination in the domain, and attempts to connect to that destination, B.
4. The connection attempt to the destination fails, so the LAC locks out destination B.
5. The LAC drops to the final (lowest) level in this example, preference level 2. Destination B is still locked out, so it cannot select tunnel 5. The LAC randomly chooses between tunnel 3 and tunnel 4. It selects tunnel 3 and attempts to connect to destination C.
6. The connection attempt to this destination also fails, so the LAC locks out destination C and rejects the login request from the PPP client.

   NOTE: **Even though tunnel 4 has a reachable destination in the domain, the LAC cannot now select that tunnel to reach destination D, because it can make only one attempt to reach a destination in the domain at each level.**

7. The client tries to log in again and the LAC repeats the tunnel selection process, starting over at preference level 0 to check for an unlocked destination, and cycling through the levels as needed.
8. Destination A is still locked out, so the LAC drops to preference level 1.
9. Destination B is still locked out, so the LAC drops to preference level 2.
10. Destinations B and C are still locked out, but destination D is still reachable. In this case, the LAC selects tunnel 4 and attempts to connect to destination D.
11. This connection attempt fails, so the LAC locks out destination D and rejects the login attempt by the PPP client.
12. The client tries to log in again and the LAC repeats the tunnel selection process, starting over at preference level 0 to check for an unlocked destination.
13. This time through, each destination in the domain at every level is still locked out. Rather than rejecting the client login request, the LAC selects a tunnel at the lowest preference level based on the destination lockout time. The LAC determines that the lockout time for destination B is less than the lockout times for destinations C and D, so it selects tunnel 5. The LAC ignores the lockout on destination B and attempts to connect.
14. This attempt fails and the LAC rejects the login request from the PPP client.

Tunnel Selection Failover Within a Preference Level

When tunnel selection failover within a preference level is configured, if the LAC tries to connect to a destination and is unsuccessful, it locks out that destination and selects a new destination at the same preference level. If all destinations at a preference level are locked out, the LAC does not attempt to connect to a destination at that level. It drops to the next lower preference level to attempt a connection. If all attempts fail at all levels, the LAC rejects the login attempt.

If the login attempt is repeated, the LAC restarts the process at the highest preference level. If all the destinations are locked out as it cycles through the levels, then at the lowest preference level, the LAC selects the tunnel at the lowest preference level that has the destination with the shortest lockout time remaining. The LAC ignores the lockout and attempts to connect to this destination. If this attempt fails, then the LAC rejects the login attempt by the PPP client.

For example, suppose that there are four destinations for a domain and a tunnel has been defined for each destination: A, B, C, and D. All destinations are considered reachable, and the preference levels for the tunnels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

In this example, when a PPP user tries to connect to the domain, the LAC randomly selects between A and B at preference level 0 and chooses destination B.

1. If the connection attempt to destination B fails, the LAC excludes tunnel B from consideration for the length of the destination lockout timeout.
2. The LAC attempts to connect to destination A at preference level 0.
3. If the connection attempt to destination A fails, the LAC excludes destination A from consideration for the length of the destination lockout timeout.
4. The LAC drops to the next level, preference level 1, to reach a destination for the domain. At preference level 1, the LAC randomly selects tunnel C.
5. If the connection attempt to destination C fails, the LAC excludes destination C from consideration for the length of the destination lockout timeout.
6. The LAC attempts to connect to destination D at preference level 1.
7. If the connection attempt to destination D fails, the LAC rejects the login attempt by the PPP client.
8. If the client tries to log in again, the LAC repeats this process, starting over at preference level 0 to check for an unlocked destination, and cycling through the levels as needed.

   If each destination that the LAC attempts this time through is locked out, then the LAC stops at the lowest preference level, 1. It selects the tunnel that has the destination with the shortest lockout time remaining. The LAC ignores the lockout and attempts to connect to this destination. If this attempt fails, then the LAC rejects the login attempt by the PPP client.

Session Failover and Tunnel Selection Behavior

If the destinations in the tunnel selected by the LAC are all unreachable, then the client fails to establish a PPP session over that tunnel. If no other tunnels exist for the domain that also match the tunnel selection criteria, the LAC sends a failure notification to the client, which results in a log-in failure and requires the client to retry the log-in attempt. However, if more tunnels are available, then the LAC attempts to establish an alternate tunnel and client session until either it succeeds or no more tunnels are available for

 Copyright © 2016, Juniper Networks, Inc.

selection. The tunnel with the highest preference that can reach the destination is selected, unless the destination is locked out.

If the destination is unreachable or if the destination is locked out, the tunnel with the next higher preference is selected. The tunnel selection process continues until either a session is successfully created or no more tunnels are available. After attempting to connect to all available destinations, the login is rejected. On the next login attempt, if all destinations are locked out, the LAC selects the tunnel at the lowest preference level that has the destination with the shortest lockout time remaining.

This session failover behavior applies only when the destination is unreachable in the following instances:

- The LNS fails to return an SCCRP message in response to the SCCRQ message from the LAC after the maximum number of retransmission attempts.
- The tunnel is established, but the LNS does not return an ICRP message in response to the ICRQ from the LAC after the maximum number of retransmission attempts.

This session failover behavior does not apply in the following circumstances:

- The client terminates the connection.
- The tunnel is established, but the LNS sends a CDN message while the LAC is attempting to establish the session with the LNS. As a result, the subscriber login attempt fails.

Tunnel Selection and Maximum Sessions per Tunnel

When the maximum number of sessions allowed per tunnel is configured, the LAC takes that setting into consideration during the tunnel selection process. The maximum number of sessions per tunnel can be configured through a RADIUS Tunnel-Max-Sessions VSA [26-33] or by including the **max-sessions** statement in a tunnel profile.

If a randomly selected tunnel has a current session count equal to its maximum session count, the LAC does not attempt to connect to a destination with that tunnel. Instead, it selects an alternate tunnel from the set of reachable tunnels at the same preference level. If no additional reachable tunnels exist at the current preference level, the LAC drops to the next lower preference level to make the selection. This process is consistent, regardless of which failover scheme is currently running on the LAC.

If the maximum number of sessions is not configured for a tunnel, then that tunnel has no upper limit on the number of sessions it can support. By default, the maximum sessions value is 0 (zero), which allows unlimited sessions in the tunnel.

Tunnel Selection with Weighted Load Balancing

The maximum sessions value for a tunnel is used for weighted load balancing to select among multiple tunnels with the same preference level.

The weight of a tunnel is determined by the tunnel's maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the largest maximum session value has the largest weight. The tunnel with the next

largest maximum session value has the next largest weight, and so on. The tunnel with the smallest maximum session value has the smallest weight.

Related Documentation

- Configuring the L2TP LAC Tunnel Selection Parameters on page 165
- Configuring the L2TP Destination Lockout Timeout on page 139

Setting the Format for the Tunnel Name

By default, the name of a tunnel corresponds to the Tunnel-Assignment-Id [82] returned by the AAA server. You can optionally configure the LAC to use more elements in the construction of a tunnel name by including the **assignment-id-format client-server-id** statement at the **[edit services l2tp tunnel]** hierarchy level. This format uses three attributes: Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. These attributes correspond, respectively, to the values configured in the tunnel profile for the LAC (source gateway) name, the tunnel endpoint (remote gateway) address on the LNS, and the tunnel ID.

A consequence of the **client-server-id** format is that the LAC automatically creates a new tunnel when the AAA server returns a different Tunnel-Client-Auth-Id than previously returned.

NOTE: **Before you downgrade to a Junos OS Release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the no assignment-id-format assignment-id statement at the [edit services l2tp tunnel] hierarchy level.**

To change how the tunnel name is formatted:

- Configure the format.

```
[edit services l2tp tunnel]
user@host# set assignment-id-format client-server-id
```

Related Documentation

- Configuring an L2TP LAC on page 149

Configuring a Tunnel Profile for Subscriber Access

The tunnel profile specifies a set of attributes to characterize the tunnel. The profile can be applied by a domain map or automatically when the tunnel is created.

Copyright © 2016, Juniper Networks, Inc.

NOTE: RADIUS attributes and VSAs can override the values you configured by a tunnel profile in a domain map. In the absence of a domain map, RADIUS can supply all the characteristics of a tunnel. The steps in the following procedure list the corresponding standard RADIUS attribute or VSA that you can configure on your RADIUS server to modify or configure the tunnel profile.

RADIUS-supplied attributes are associated with a tunnel by a tag carried in the attribute, which matches the tunnel identifier. A tag of 0 indicates the tag is not used. If L2TP receives a RADIUS attribute with a tag of 0, the attribute cannot be merged with the tunnel profile configuration corresponding to the subscriber domain because a tunnel profile cannot provide a tunnel tag (tunnel identifier) of 0. Only tags in the range of 1 through 31 are supported.

To configure a tunnel definition for a tunnel profile:

1. Specify the tunnel profile for which you are defining a tunnel. (Tunnel-Group [26-64])

   [edit access]
   user@host# set tunnel-profile *profile-name*

2. Specify an identifier (name) for the L2TP control connection for the tunnel.

   [edit access tunnel-profile *profile-name*]
   user@host# set tunnel *tunnel-id*

3. Configure the IP address of the local L2TP tunnel endpoint, the LAC. (Tunnel-Client-Endpoint [66])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set source-gateway address *client-ip-address*

4. Configure the IP address of the remote L2TP tunnel endpoint, the LNS. (Tunnel-Server-Endpoint [67])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set remote-gateway address *server-ip-address*

5. (Optional) Configure the preference level for the tunnel. (Tunnel-Preference [83])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set preference *number*

6. (Optional) Configure the hostname of the local client (LAC). (Tunnel-Client-Auth-Id [90])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set source-gateway gateway-name *client-name*

7. (Optional) Configure the hostname of the remote server (LNS). (Tunnel-Server-Auth-Id [91])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set remote-gateway gateway-name *server-name*

8. (Optional) Specify the medium (network) type for the tunnel. (Tunnel-Medium-Type [65])

[edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
user@host# set medium *type*

9. (Optional) Specify the protocol type for the tunnel. (Tunnel-Type [64])

   [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
   user@host# set type *tunnel-type*

10. (Optional) Configure the assignment ID for the tunnel. (Tunnel-Assignment-Id [82])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set identification *name*

11. (Optional) Configure the maximum number of sessions allowed in the tunnel. (Tunnel-Max-Sessions [26-33])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set max-sessions *number*

12. (Optional) Configure the password for remote server authentication. (Standard RADIUS attribute Tunnel-Password [69] or VSA Tunnel-Password [26-9])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set secret *password*

13. (Optional) Configure the logical system to use for the tunnel.

    If you configure a logical system, you must also configure a routing instance.

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set logical-system *logical-system-name*

14. (Optional) Configure the routing instance to use for the tunnel. (Tunnel-Virtual-Router [26-8])

    If you configure a routing instance, configuring a logical system is optional.

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set routing-instance *routing-instance-name*

15. (Optional) Enable the LAC to interoperate with Cisco LNS devices. (Tunnel-Nas-Port-Method [26-30])

    [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]
    user@host# set nas-port-method cisco-avp

The following example shows a complete configuration for a tunnel profile:

```
tunnel-profile marketing {
  tunnel 1 {
    preference 5;
    remote-gateway {
      address 172.16.98.4;
      gateway-name work;
    }
    source-gateway {
      address 192.168.4.10;
      gateway-name local;
    }
    secret mk5Sn$3k%V;
    logical-system bos-metro-5;
```


Copyright © 2016, Juniper Networks, Inc.

```
        routing-instance rox-12-32;
        medium ipv4;
        type l2tp;
        identification tunnel_to_work;
        max-sessions 32;
        nas-port-method cisco avp;
    }
}
```

Related Documentation

- Configuring an L2TP LAC on page 149
- *Domain Mapping Overview*
- LAC Interoperation with Third-Party LNS Devices on page 152

Configuring the L2TP LAC Tunnel Selection Parameters

When the LAC determines that a PPP session should be tunneled, it selects a tunnel from the set of tunnels associated with either the PPP user or the PPP user's domain. You can configure how a tunnel is selected and whether certain information is sent by the LAC to the LNS.

To configure tunnel selection parameters:

1. (Optional) Configure how a tunnel is selected when a connection attempt fails.

   See "Configuring LAC Tunnel Selection Failover Within a Preference Level" on page 165.

2. (Optional) Configure how sessions are load-balanced among tunnels.

   See "Configuring Weighted Load Balancing for LAC Tunnel Sessions" on page 166.

Related Documentation

- LAC Tunnel Selection Overview on page 157

Configuring LAC Tunnel Selection Failover Within a Preference Level

You can configure how LAC tunnel selection continues in the event of a failure to connect. By default, when the router is unable to connect to a destination at a given preference level, it attempts to connect at the next lower level. You can specify that the router instead attempt to connect to another destination at the same level as the failed attempt.

If all destinations at a preference level are marked as unreachable, the router does not attempt to connect to a destination at that level. It drops to the next lower preference level to select a destination.

If all destinations at all preference levels are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. If the connection fails, the router rejects the PPP user session without attempting to contact the remote router.

For example, suppose there are four tunnels for a domain: A, B, C, and D. All tunnels are considered reachable, and the preference levels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

When the router attempts to connect to the domain, suppose it randomly selects tunnel B from preference 0. If it fails to connect to tunnel B, the router excludes tunnel B for five minutes and attempts to connect to tunnel A. If this attempt also fails, the router drops to preference 1. Then suppose the router selects tunnel C. If it also fails to connect to tunnel C, the router excludes tunnel C for five minutes and attempts to connect to tunnel D.

You configure the preference level used for this tunnel selection method in the tunnel profile or the RADIUS Tunnel-Preference [83] attribute.

To enable tunnel selection failover within a preference level:

- Specify failover within preference.

  ```
  [edit services l2tp]
  user@host# set failover-within-preference
  ```

Related Documentation

- LAC Tunnel Selection Overview on page 157
- Configuring the L2TP LAC Tunnel Selection Parameters on page 165
- Configuring a Tunnel Profile for Subscriber Access on page 162
- *Configuring How RADIUS Attributes Are Used for Subscriber Access*

Configuring Weighted Load Balancing for LAC Tunnel Sessions

You can configure how L2TP LAC sessions are distributed across tunnels. You can specify that the router uses the maximum sessions per tunnel to choose among multiple tunnels that share the same preference level.

The weight of a tunnel is proportional to its maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the largest maximum session value has the highest weight. The tunnel with the next larger maximum session value has the next higher weight, and so on. The tunnel with the smallest maximum session value has the lowest weight.

When you configure weighted load balancing, the tunnel with the highest weight in the preference level is selected until the maximum number of sessions for the tunnel is reached. Then the router selects the tunnel with the next higher weight to establish connections until that tunnel's maximum session limit is reached, and so on.

To configure weighted load balancing:

- Specify load balancing.

  ```
  [edit services l2tp]
  ```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set weighted-load-balancing
```

Related Documentation

- LAC Tunnel Selection Overview on page 157
- Configuring the L2TP LAC Tunnel Selection Parameters on page 165

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 20

Configuring Transmission Connection Speeds to LNS

Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS

An L2TP access concentrator (LAC) uses Incoming-Call-Connected (ICCN) messages to send the attribute-value pair (AVP) 24 (which represents the transmit connect speed) and the AVP 38 (which represents the receive connect speed) to the L2TP network server (LNS).

The value of speed in AVP 24 and AVP 38 is typically not greater than the value that is enforced by CoS on the LAC side of the network. Any difference between the speed reported in these AVPs and that enforced by CoS can occur because of differences between the CoS configuration (of the source that is used to enforce a downstream speed) and the transmit connect speed method used to establish these AVPs.

When you include the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level, the transmit connect speed method selected for the downstream speed, AVP 24, also applies to the selection of the upstream speed , AVP 38. You can configure the transmit and receive connect speeds to be derived from Point-to-Point Protocol over Ethernet (PPPoE) intermediate agent tags that are sent from the digital subscriber line access multiplexer (DSLAM) to the LAC, the Access Node Control Protocol (ANCP) settings of the underlying interface, or the recommended (advisory) shaping rate. When the method you specify as ANCP, PPPoE intermediate agent tags (PPPoE IA tags), or advisory shaping rate does not support an upstream speed, the static or advisory speed is used.

A fallback method is adopted to derive the transmit and receive connect speed values when these values cannot be determined from any of the configured methods, such as from the ANCP settings, the PPPoE IA tags, or the advisory shaping rate, or if the speed value is determined to be 0 from any of the configured methods.

Copyright © 2016, Juniper Networks, Inc.

If you configure the ANCP method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from ANCP.
2. If the values cannot be derived from ANCP, the PPPoE IA tags are used to determine the values. If the PPPoE IA tags are present for either or both transmit and receive connect speeds, these values are used.
3. If the values cannot be derived from the PPPoE IA tags, the recommended (advisory) shaping rate configured on the PPPoE logical interface is used. If the advisory shaping rate is present for either or both transmit and receive connect speeds, these values are used.
4. If the values cannot be derived from the advisory shaping rate, the configured or default port speed is used for transmit and receive connect speeds.

If you configure the PPPoE IA tags method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from PPPoE IA tags.
2. If the values cannot be derived from the PPPoE IA tags, the recommended (advisory) shaping rate configured on the PPPoE logical interface is used. If the advisory shaping rate is present for either or both transmit and receive connect speeds, these values are used.
3. If the values cannot be derived from the advisory shaping rate, the configured or default port speed is used for transmit and receive connect speeds.

If you configure the static or advisory downstream shaping rate method to calculate the connect speed, the following sequence of events takes place:

1. The upstream and downstream connect speed values are derived from the advisory shaping rate.
2. If the values cannot be derived from the advisory shaping rate, the default port speed is used for transmit and receive connect speeds.

The transmit connect speed, AVP 24, is set in the ICCN messages on the basis of the method for determining the transmit connect speed configured using the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level. You can configure the method for determining the transmit connect speed in the following order of precedence:

1. **ancp**—The speed is derived from the configured ANCP value for the underlying interface. You can change this speed after a subscriber has logged in, but those changes do not affect the actual speed used by the LNS.
2. **pppoe-ia-tags**—PPPoE IA tags sent from the DSLAM to the LAC. This speed value is transmitted when a subscriber logs in and it cannot be subsequently changed. This value is used when the **ancp** value is not available. This speed does not apply to the

 Copyright © 2016, Juniper Networks, Inc.

subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved.

3. **static**—The speed is derived from the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory shaping rate is not configured on the underlying interface, then the actual port speed is used.

If you do not configure the transmit connect speed using the CLI interface, and if the advisory speed is also not available, then the actual port speed is used. For ge and xe interfaces, the speed value is set to 10,000,000 and for ae interfaces, the speed value is set to 0 and sent in both AVP 24 and AVP 38.

Related Documentation

- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds Are Equal on page 171
- Configuring an L2TP LAC on page 149

Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds Are Equal

The L2TP Rx Connect Speed (in bits per second) AVP, which is represented by AVP 38, is included in the ICCN message when the receive connect speed is different from the transmit connect speed. By default, when the connection speed is the same in both directions, AVP 38 is not sent; the LNS uses the value in AVP 24 for both transmit and receive connect speeds.

AVP 38 is generated when the receive connect speed of the access interface is set equal to the calculated transmit connect speed by issuing the **rx-connect-speed-when-equal** statement at the **[edit services l2tp]** hierarchy level. In this scenario, the LAC transmits the same value for transmit and receive connect speeds that are sent to the LNS through the AVP 24 and AVP 38 in the ICCN message.

To configure the sending of AVP 38 when the connection speeds are the same in both the downstream and upstream directions:

- Configure the transmission of the receive connect speed, AVP 38, when the receive connect speed is set equal to the calculated transmit connect speed.

```
[edit services l2tp]
user@host# set rx-connect-speed-when-equal
```

Related Documentation

- Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS on page 169
- Configuring an L2TP LAC on page 149
- **rx-connect-speed-when-equal on page 410**

Preventing the LAC from Sending Calling Number AVP 22 to the LNS

Calling Number AVP 22 typically identifies the interface that is connected to the customer in the access network. When RADIUS includes the Calling-Station-Id in the Access-Accept message, that value is used for the Calling Number AVP. Otherwise, the underlying interface (for example, the S-VLAN IFL) on which the PPPoE session is established is used for the Calling Number AVP value.

By default, the LAC includes this AVP in the incoming-call request (ICRQ) packets that it sends to the LNS. However, you may wish to hide your network access interface information. To do so, you can configure the tunnel so that the LAC does not send the Calling Number AVP to the LNS.

To disable sending the Calling Number AVP:

- Configure disabling.

  [edit services l2tp]
  user@host# **set** disable-calling-number-avp

Related Documentation

- LAC Tunnel Selection Overview on page 157

Configuring the Method to Set the LAC Connection Speeds to the LNS

During the establishment of an L2TP tunnel session, the LAC sends the L2TP transmit connect speed in bits per second (BPS) using AVP 24 to the LNS in Incoming-Call-Connected (ICCN) messages. AVP 24 conveys the transmit connect speed of the subscriber's access interface; that is, it represents the speed of the connection from the LAC to the LNS, from the perspective of the LAC. The L2TP receive connect speed, which is represented by AVP 38, is included in the message when the receive connect speed is different from the transmit connect speed. AVP 38 conveys the receive connect speed of the connection from the LNS to the LAC, again from the perspective of the LAC. When AVP 38 is not sent, it means that the connection speed is the same in both directions; the LNS uses the value in AVP 24 for both transmit and receive connect speeds.

When the Tunnel-Tx-Speed-Method RADIUS VSA [26–94] is configured, RADIUS returns it in the Access-Accept message for individual subscribers. The TX speed is determined by the method set in this VSA: actual, CoS, or static layer 2. The dynamic layer 2 value is not supported, but is translated to static layer two rather than being rejected.

If the RADIUS VSA is not returned for the subscriber, then the TX speed is set according to the method that you configure globally with the **tx-connect-speed-method** statement at the **[edit services l2tp]** hierarchy level. The method configured specifies the resource that the LAC uses to set the speed.

You can configure what the LAC uses as a resource for setting these speeds. To use the recommended (advisory) downstream shaping rate for AVP 24 and the recommended upstream shaping rate for AVP 38, include the **tx-connect-speed-method static** statement

 Copyright © 2016, Juniper Networks, Inc.

at the **[edit services l2tp]** hierarchy level. You configure the advisory rates under the PPPoE logical interface underlying the subscriber interface with the **advisory-options** statement at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*****]** hierarchy level. When the advisory speed is not configured on the underlying interface, the **tx-connect-speed-method advisory** statement automatically sets the speed to 1 Gbps and sends this value in both AVP 24 and AVP 38.

To derive the speeds from the PPPoE IA tags, use the **tx-connect-speed-method pppoe-ia-tags** statement. In this case, AVP 24 is the value of Actual-Data-Rate-Downstream (VSA 26-129). AVP 38 is the value of Actual-Data-Rate-Upstream (VSA 26-130), and is sent only when the VSA values differ.

To derive the speeds from the ANCP value configured on the PPPoE interface underlying the subscriber interface, use the **tx-connect-speed-method ancp** statement.

To set the method for calculating the transmit connect speed:

- Configure the ANCP method to use the values derived from the configured ANCP value for the underlying interface.

  ```
  [edit services l2tp]
  user@host# set tx-connect-speed-method ancp
  ```

- Configure the PPPoE IA tags method to use the values provided in the PPPoE IA tags.

  ```
  [edit services l2tp]
  user@host# set tx-connect-speed-method pppoe-ia-tags
  ```

- Configure the static (advisory downstream shaping rate) method to use the underlying interface's recommended shaping rates.

  ```
  [edit services l2tp]
  user@host# set tx-connect-speed-method static
  ```

Related Documentation

- Configuring an L2TP LAC on page 149

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 21

Configuring L2TP LNS Inline Service Interfaces

Configuring an L2TP LNS with Inline Service Interfaces

The L2TP LNS feature license must be installed before you begin the configuration. Otherwise, a warning message is displayed when the configuration is committed.

To configure an L2TP LNS with inline service interfaces:

1. (Optional) Configure a user group profile that defines the PPP configuration for tunnel subscribers.

   See "Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile" on page 179.

2. (Optional) Configure PPP attributes for subscribers on inline service interfaces.

   See "Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface" on page 177.

Copyright © 2016, Juniper Networks, Inc.

3. Configure inline IP reassembly.

   See "Configuring IP Inline Reassembly for L2TP" on page 204.

4. Configure an L2TP access profile that defines the L2TP parameters for each LNS client (LAC).

   See "Configuring an L2TP Access Profile on the LNS" on page 180.

5. (Optional) Configure a AAA access profile to override the access profile configured under the routing instance.

   See "Configuring a AAA Local Access Profile on the LNS" on page 181.

6. Configure a pool of addresses to be dynamically assigned to tunneled PPP subscribers.

   See "Configuring an Address-Assignment Pool for L2TP LNS with Inline Services" on page 182.

7. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address.

   See "Configuring the L2TP LNS Peer Interface" on page 183.

8. Enable inline service interfaces on an MPC.

   See "Enabling Inline Service Interfaces" on page 184.

9. Configure a service interface.

   See "Configuring an Inline Service Interface for L2TP LNS" on page 185.

10. Configure options for each inline service logical interface.

    See "Configuring Options for the LNS Inline Services Logical Interface" on page 185.

11. Configure the L2TP tunnel group.

    See "Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces" on page 198.

12. (Optional) Configure a dynamic profile that dynamically creates L2TP logical interfaces.

    See "Configuring a Dynamic Profile for Dynamic LNS Sessions" on page 200

13. (Optional) Configure a service interface pool for dynamic LNS sessions.

    See "Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions" on page 199.

14. (Optional) Specify how many times L2TP retransmits unacknowledged control messages.

    See "Configuring Retransmission Attributes for L2TP Control Messages" on page 143.

15. (Optional) Specify how long a tunnel can remain idle before being torn down.

    See "Setting the L2TP Tunnel Idle Timeout" on page 137.

Copyright © 2016, Juniper Networks, Inc.

16. (Optional) Specify the L2TP receive window size for the L2TP tunnel. The receive window size specifies the number of packets a peer can send before waiting for an acknowledgment from the router.

    See "Setting the L2TP Receive Window Size" on page 137

17. (Optional) Specify how long the L2TP retains information about terminated dynamic tunnels, sessions, and destinations.

    See "Setting the L2TP Destruct Timeout" on page 138.

18. (Optional) Configure the L2TP destination lockout timeout.

    See "Configuring the L2TP Destination Lockout Timeout" on page 139.

19. (Optional) Configure L2TP tunnel switching.

    See "Configuring L2TP Tunnel Switching" on page 135

20. (Optional) Enable SNMP statistics counters.

    See "Enabling Tunnel and Global Counters for SNMP Statistics Collection" on page 214

21. (Optional) Configure trace options for troubleshooting the configuration.

    See "Tracing L2TP Operations for Subscriber Access" on page 249

You also need to configure CoS for LNS sessions. For more information, see *Configuring Dynamic CoS for an L2TP LNS Inline Service*.

Related Documentation

- L2TP for Subscriber Access Overview on page 123
- *Junos OS Feature Licenses*
- *Software Feature Licenses*

Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface

You can configure PPP attributes that are applied by the LNS on the inline service (si) interface to the PPP subscribers tunneled from the LAC. Because you are configuring the attributes per interface rather than with a user group profile, the attributes for subscribers can be varied with a finer granularity. This configuration matches that used for terminated PPPoE subscribers.

To configure the PPP attributes for dynamically created si interfaces:

1. Specify the predefined dynamic interface and logical interface variables in the dynamic profile.

   [edit dynamic-profiles *profile-name*]
   user@host# edit interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
   user@host# set keepalives interval *seconds*

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
   user@host# set ppp-options chap
   user@host# set ppp-options pap

To configure the PPP attributes for statically created si interfaces:

1. Specify the logical inline service interface.

   [edit interfaces si-*slot/pic/port*]
   user@host# edit unit *logical-unit-number*

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# set keepalives interval *seconds*

3. Configure the number of keepalive packets a destination must fail to receive before the network takes down a link.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# set keepalives down-count *number*

   NOTE: The keepalives up-count option is typically not used for subscriber management.

4. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit interfaces si-*slot/pic/port* unit *logical-unit-number*]
   user@host# set ppp-options chap
   user@host# set ppp-options pap

BEST PRACTICE: Although all other statements subordinate to ppp-options—including those subordinate to chap and pap—are supported, they are typically not used for subscriber management. We recommend that you leave these other statements at their default values.

NOTE: You can also configure PPP attributes with a user group profile that applies the attributes to all subscribers with that profile on a LAC client. See "Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile" on page 179 for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

You can configure a user group profile that enables the LNS to apply PPP attributes to the PPP subscribers tunneled from the LAC. The user group profile is associated with clients (LACs) in the L2TP access profile. Consequently all subscribers handled by a given client share the same PPP attributes.

To configure a user group profile:

1. Create the profile.

   [edit access]
   user@host# **edit group-profile** *profile-name*

2. Configure the interval between PPP keepalive messages for the L2TP tunnel terminating on the LNS.

   [edit access group-profile *profile-name*]
   user@host# **set ppp keepalive** *seconds*

   NOTE: Changes to the keepalive interval in a user group profile affect only new L2TP sessions that come up after the change. Existing sessions are not affected.

3. Configure PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

   [edit access group-profile *profile-name*]
   user@host# **set ppp ppp-options chap**
   user@host# **set ppp ppp-options pap**

4. Configure how long the PPP subscriber session can be idle before it is considered to have timed out.

   [edit access group-profile *profile-name*]
   user@host# **set ppp idle-timeout 200**

   NOTE: You can also configure PPP attributes on a per-interface basis. See "Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface" on page 177 for more information. When you configure the PPP attributes for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the user group profile configuration.

Related Documentation

- Configuring an L2TP Access Profile on the LNS on page 180
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Configuring an L2TP Access Profile on the LNS

Access profiles define how to validate Layer 2 Tunneling Protocol (L2TP) connections and session requests. Within each L2TP access profile, you configure one or more clients (LACs). The client characteristics are used to authenticate LACs with matching passwords, and to establish attributes of the client tunnel and session. You can configure multiple access profiles and multiple clients within each profile.

To configure an L2TP access profile:

1. Create the access profile.

   [edit access]
   user@host# **edit profile** ***access-profile-name***

2. Configure characteristics for one or more clients (LACs).

   [edit access profile *access-profile-name*]
   user@host# client ***client-name***

   NOTE: Except for the special case of the **default** client, the LAC client name that you configure in the access profile must match the hostname of the LAC. In the case of a Juniper Networks router acting as the LAC, the hostname is configured in the LAC tunnel profile with the gateway gateway-name statement at the **[edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]** hierarchy level. Alternatively, the client name can be returned from RADIUS in the attribute, Tunnel-Client-Auth-Id [90].

   NOTE: Use **default** as the client name when you want to define a default tunnel client. The default client enables the authentication of multiple LACs with the same secret and L2TP attributes. This behavior is useful when, for example, many new LACs are added to the network, because it enables the LACs to be used without additional LNS profile configuration.

   Use **default** only on MX Series routers. The equivalent client name on M Series routers is *.

3. (Optional) Specify a local access profile that overrides the global access profile and the tunnel group AAA access profile to configure RADIUS server settings for the client.

   [edit access profile *access-profile-name* client *client-name*]
   user@host# **set l2tp** aaa-access-profile

4. Configure the LNS to renegotiate the link control protocol (LCP) with the PPP client. tunneled from the client.

   [edit access profile *access-profile-name* client *client-name*]
   user@host# **set l2tp** lcp-renegotiation

5. Configure the maximum number of sessions allowed in a tunnel from the client (LAC).

Copyright © 2016, Juniper Networks, Inc.

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp maximum-sessions-per-tunnel number
```

6. Configure the tunnel password used to authenticate the client (LAC).

```
[edit access profile access-profile-name client client-name]
user@host# set l2tp shared-secret shared-secret
```

7. (Optional) Associate a group profile containing PPP attributes to apply for the PPP sessions being tunneled from this LAC client.

```
[edit access profile access-profile-name client client-name]
user@host# set user-group-profile group-profile-name
```

NOTE: If *<user-group-profile>* is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, will go down.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

Configuring a AAA Local Access Profile on the LNS

For some LNS tunnels, you might wish to override the access profile configured at the routing instance that hosts the tunnel with a particular RADIUS server configuration. You can configure a local access profile to do so. You can subsequently use the **aaa-access-profile** statement to apply the local access profile to a tunnel group or LAC client.

A local access profile applied to a client overrides a local access profile applied to a tunnel group, which in turn overrides the access profile for the routing instance.

To configure an AAA local access profile:

1. Create the access profile.

```
[edit access]
user@host# edit profile local-aaa-profile-name
```

2. Configure the order of AAA authentication methods.

```
[edit access profile local-aaa-profile-name]
user@host# set authentication-order radius
```

3. Configure the RADIUS server attributes, such as the authentication password.

```
[edit access profile local-aaa-profile-name]
user@host# set radius-server server-address secret password
```

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175

- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198
- Configuring an L2TP Access Profile on the LNS on page 180

Configuring an Address-Assignment Pool for L2TP LNS with Inline Services

You can configure pools of addresses that can be dynamically assigned to the tunneled PPP subscribers. The pools must be local to the routing instance where the subscriber comes up. The configured pools are supplied in the RADIUS Framed-Pool and Framed-IPv6-Pool attributes. Pools are optional when Framed-IP-Address is sent by RADIUS.

To configure an address-assignment pool, you must specify the name of the pool and configure the addresses for the pool.

You can optionally configure multiple named ranges, or subsets, of addresses within an address-assignment pool. During dynamic address assignment, a client can be assigned an address from a specific named range. To create a named range, you specify a name for the range and define the address range.

NOTE: **Be sure to use the address-assignment pools (address-assignment) statement rather than the address pools (address-pool) statement.**

To configure an IPv4 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv4 family.

   [edit access]
   user@host# **edit address-assignment pool *pool-name* family inet**

2. Configure the network address and the prefix length of the addresses in the pool.

   [edit access address-assignment pool *pool-name* family inet]
   user@host# **set network *ip-prefix/<prefix-length***

3. Configure the name of the range and the lower and upper boundaries of the addresses in the range.

   [edit access address-assignment pool *pool-name* family inet]
   user@host# **set range *range-name* low *lower-limit* high *upper-limit***

For example, to configure an IPv4 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v4-pool family inet
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set network 192.168.1.1/16
[edit access address-assignment pool lns-v4-pool family inet]
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
```


Copyright © 2016, Juniper Networks, Inc.

NOTE: Dual-stack (IPv4/IPv6) LNS is supported, but IPv6-only subscribers are not supported.

To configure an IPv6 address-assignment pool for L2TP LNS:

1. Configure the name of the pool and specify the IPv6 family.

   [edit access]
   user@host# **edit address-assignment pool *pool-name* family inet6**

2. Configure the IPv6 network prefix for the address pool. The prefix specification is required when you configure an IPv6 address-assignment pool.

   [edit access address-assignment pool *pool-name* family inet6]
   user@host# **set prefix *ipv6-prefix***

3. Configure the name of the range and define the range. You can define the range based on the lower and upper boundaries of the prefixes in the range, or based on the length of the prefixes in the range.

   [edit access address-assignment pool *pool-name* family inet6]
   user@host# **set range *range-name* low *lower-limit* high *upper-limit***

For example, to configure an IPv6 address-assignment pool:

```
[edit access]
user@host# edit address-assignment pool lns-v6-pool family inet6
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set prefix 2010:DB8::/32
[edit access address-assignment pool lns-v6-pool family inet6]
user@host# set range lns-v6-pool-range low 2010:DB8:1::/48 high 2010:DB8::ffff::/48
```

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*

Configuring the L2TP LNS Peer Interface

The peer interface connects the LNS to the cloud towards the LACs so that IP packets can be exchanged between the tunnel endpoints. MPLS and aggregated Ethernet can also be used to reach the LACs.

NOTE: On MX Series routers, you must configure the peer interface on an MPC.

To configure the LNS peer interface:

1. Specify the interface name.

   [edit interfaces]

```
user@host# edit interface-name
```

2. Enable VLANs.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

3. Specify the logical interface, bind a VLAN tag ID to the interface, and configure the address family and the IP address for the logical interface.

```
[edit interfaces interface-name]
user@host# edit unit logical-unit-number
[edit interfaces interface-name unit logical-unit-number]
user@host# set vlan-id number
user@host# set family family address ip-address
```

NOTE: The IPv6 address family is not supported as a tunnel endpoint.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Enabling Inline Service Interfaces

The inline service interface is a virtual physical interface that resides on the Packet Forwarding Engine. This **si** interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

NOTE: On MX80 routers, you can configure only four inline services physical interfaces as anchor interfaces for L2TP LNS sessions: si-1/0/0, si-1/1/0, si-1/2/0, and si-1/3/0. You cannot configure si-0/0/0 for this purpose on MX80 routers.

To enable inline service interfaces:

1. Access an MPC-occupied slot and the PIC where the interface is to be enabled.

```
[edit chassis]
user@host# edit fpc slot-number pic number
```

2. Enable the interface and specify the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.

```
[edit chassis fpc slot-number pic number]
user@host# set inline-services bandwidth (1g | 10g)
```

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Copyright © 2016, Juniper Networks, Inc.

Configuring an Inline Service Interface for L2TP LNS

The inline service interface is a virtual physical service interface that resides on the Packet Forwarding Engine. This **si** interface, referred to as an *anchor* interface, makes it possible to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot.

You can maximize the number of sessions that can be shaped in one service interface by setting the maximum number of hierarchy levels to two. In this case, each LNS session consumes one L3 node in the scheduler hierarchy for shaping.

If you do not specify the number of levels (two is the only option), then the number of LNS sessions that can be shaped on the service interface is limited to the number of L2 nodes, or 4096 sessions. Additional sessions still come up, but they are not shaped.

To configure an inline service interface:

1. Access the service interface.

   [edit interfaces]
   user@host# **edit si-*slot/pic/port***

2. (Optional; for per-session shaping only) Enable the inline service interface for hierarchical schedulers and limit the number of scheduler levels to two.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set hierarchical-scheduler maximum-hierarchy-levels 2**

3. (Optional; for per-session shaping only) Configure services encapsulation for inline service interface.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set encapsulation generic-services**

4. Configure the IPv4 family on the reserved unit 0 logical interface.

   [edit interfaces si-*slot/pic/port*]
   user@host# **set unit 0 family inet**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Configuring Options for the LNS Inline Services Logical Interface

You must specify characteristics—**dial-options**—for each of the inline services logical interfaces that you configure for the LNS. LNS on MX Series routers supports only one session per logical interface, so you must configure it as a **dedicated** interface; the **shared** option is not supported. (LNS on M Series routers supports **dedicated** and **shared** options.) You also configure an identifying name for the logical interface that matches the name you specify in the access profile.

You must specify the **inet** address family for each static logical interface. Although the CLI accepts either **inet** or **inet6** for static logical interfaces, the subscriber cannot log in successfully unless the address family is **inet** is configured.

To configure the logical interface options:

1. Access the inline services logical interface.

   [edit]
   user@host# **edit interfaces si-*fpc/pic/port* unit (Interfaces)** ***logical-unit-number***

2. Specify an identifier for the logical interface.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set dial-options l2tp-interface-id** ***name***

3. Configure the logical interface to be used for only one session at a time.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set dial-options dedicated**

4. Configure the address family for each logical interface and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.

   [edit interfaces si-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set family inet unnumbered-address lo0.0**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Access Profile on the LNS on page 180

Example: Configuring an L2TP LNS

This example shows how you can configure an L2TP LNS on an MX Series router to provide tunnel endpoints for an L2TP LAC in your network. This configuration includes a dynamic profile for dual-stack subscribers.

Requirements

L2TP LNS requires the following hardware and software:

- MX Series 3D Universal Edge Router
- One or more MPCs

No special configuration beyond device initialization is required before you can configure this feature.

You must configure certain standard RADIUS attributes and Juniper Networks VSAs in the attribute return list on the AAA server associated with the LNS for this example to

 Copyright © 2016, Juniper Networks, Inc.

work. Table 15 on page 187 lists the attributes with their required order setting and values. We recommend that you use the most current Juniper Networks RADIUS dictionary, available in the *Downloads* box on the *Junos OS Subscriber Management* page at http://www.juniper.net/techpubs/en_US/junos/information-products/pathway-pages/subscriber-access/index.html.

Table 15: VSA and Standard RADIUS Attribute Names, Order, and Values Required for Example

VSA Name [Number]	Order	Value
CoS-Parameter-Type [26–108]	1	T01 Multiplay
CoS-Parameter-Type [26–108]	2	T02 10m
CoS-Parameter-Type [26–108]	3	T08 -36
CoS-Parameter-Type [26–108]	4	T07 cell-mode
Framed-IPv6-Pool [100]	0	jnpr_ipv6_pool
Framed-Pool [88]	0	jnpr_pool
Egress-Policy-Name [26-11]	0	classify
Ingress-Policy-Name [26-10]	0	classify
Virtual-Router [26-1]	0	default

Overview

The LNS employs user group profiles to apply PPP attributes to the PPP subscribers that are tunneled from the LAC. LACs in the network are clients of the LNS. The clients are associated with user group profiles in the L2TP access profile configured on the LNS. In this example, the user group profile **ce-l2tp-group-profile** specifies the following PPP attributes:

- A 30-second interval between PPP keepalive messages for L2TP tunnels from the client LAC terminating on the LNS.
- A 200-second interval that defines how long the PPP subscriber session can be idle before it is considered to have timed out.
- Both PAP and CHAP as the PPP authentication methods that apply to tunneled PPP subscribers at the LNS.

The L2TP access profile **ce-l2tp-profile** defines a set of L2TP parameters for each client LAC. In this example, the user group profile **ce-l2tp-group-profile** is associated with both clients, **lac1** and **lac2**. Both clients are configured to have the LNS renegotiate the link control protocol (LCP) with the PPP client rather than accepting the pre-negotiated LCP parameters that the LACs pass to the LNS. LCP renegotiation also causes authentication to be renegotiated by the LNS; the authentication method is specified in the user group profile. The maximum number of sessions allowed per tunnel is set to 1000 for **lac1** and to 4000 for **lac2**. A different password is configured for each LAC.

Copyright © 2016, Juniper Networks, Inc.

A local AAA access profile, **aaa-profile**, enables you to override the global AAA access profile, so that you can specify an authentication order, a RADIUS server that you want to use for L2TP, and a password for the server.

In this example, an address pool defines a range of IP addresses that the LNS allocates to the tunneled PPP sessions. This example defines ranges of IPv4 and IPv6 addresses.

Two inline service interfaces are enabled on the MPC located in slot 5 of the router. For each interface, 10 Gbps of bandwidth is reserved for tunnel traffic on the interface's associated PFE. These *anchor* interfaces serve as the underlying physical interface. To enable CoS queue support on the individual logical inline service interfaces, you must configure both services encapsulation (**generic-services**) and hierarchical scheduling support on the anchors. The IPv4 address family is configured for both anchor interfaces. Both anchor interfaces are specified in the **lns_p1** service device pool. The LNS can balance traffic loads across the two anchor interfaces when the tunnel group includes the pool.

This example uses the dynamic profile **dyn-lns-profile2** to specify characteristics of the L2TP sessions that are created or assigned dynamically when a subscriber is tunneled to the LNS. For many of the characteristics, a predefined variable is set; the variables are dynamically replaced with the appropriate values when a subscriber is tunneled to the LNS.

The interface to which the tunneled PPP client connects (**$junos-interface-name**) is dynamically created in the routing instance (**$junos-routing-instance**) assigned to the subscriber. Routing options for access routes include the route's next hop address (**$junos-framed-route-nexthop**), metric (**$junos-framed-route-cost**), and preference (**$junos-framed-route-distance**). For access-internal routes, a dynamic IP address variable (**$junos-subscriber-ip-address**) is set.

The logical inline service interfaces are defined by the name of a configured anchor interface (**$junos-interface-ifd-name**) and a logical unit number (**$junos-interface-unit**). The profile assigns **l2tp-encapuslation** as the identifier for the logical interface and specifies that each interface can be used for only a single session at a time.

The IPv4 address is set to a value returned from the AAA server. For IPv4 traffic an input firewall filter **$junos-input-filter** and an output firewall filter **$junos-output-filter** are attached to the interface. The loopback variable (**$junos-loopback-interface**) derives an IP address from a loopback interface (**lo**) configured in the routing instance and uses it in IPCP negotiation as the PPP server address. Because this is a dual-stack configuration, the IPv6 address family is also set, with the addresses provided by the **$junos-ipv6-address** variable.

The **$junos-ipv6-address** variable is used because Router Advertisement Protocol is also configured. This variable enables AAA to allocate the first address in the prefix to be reserved as the local address for the interface. The minimal configuration for the Router Advertisement Protocol in the dynamic profile specifies the **$junos-interface-name** and **$junos-ipv6-ndra-prefix** variables to dynamically assign a prefix value in IPv6 neighbor discovery router advertisements.

The dynamic profile also includes the class of service configuration that is applied to the tunnel traffic. The traffic-control profile (**tc-profile**) includes variables for the scheduler

 Copyright © 2016, Juniper Networks, Inc.

map (**$junos-cos-scheduler-map**), shaping rate (**$junos-cos-shaping-rate**), overhead accounting (**$junos-cos-shaping-mode**), and byte adjustment **$junos-cos-byte-adjust**). The dynamic profile applies the CoS configuration—including the forwarding class, the output traffic-control profile, and the rewrite rules—to the dynamic service interfaces.

The **tg-dynamic** tunnel group configuration specifies the access profile **ce-l2tp-profile**, the local AAA profile **aaa-profile**, and the dynamic profile **dyn-lns-profile2** that are used to dynamically create LNS sessions and define the characteristics of the sessions. The **lns_p1** service device pool associates a pool of service interfaces with the group to enable LNS to balance traffic across the interfaces. The local gateway address **10.1.1.2** corresponds to the remote gateway address that is configured on the LAC. The local gateway name **ce-lns** corresponds to the remote gateway name that is configured on the LAC.

NOTE: **This example does not show all possible configuration choices.**

Configuration

CLI Quick Configuration

To quickly configure an L2TP LNS, copy the following commands, paste them in a text file, remove any line breaks, and then copy and paste the commands into the CLI.

```
[edit]
edit access group-profile ce-l2tp-group-profile
set ppp idle-timeout 200
set ppp ppp-options pap
set ppp ppp-options chap
set ppp keepalive 30
top
edit access profile ce-l2tp-profile
set client lac1 l2tp maximum-sessions-per-tunnel 1000
set client lac1 l2tp interface-id l2tp-encapsulation-1
set client lac1 l2tp lcp-renegotiation
set client lac1 l2tp shared-secret "lac1-secret"
set client lac1 user-group-profile ce-l2tp-group-profile
set client lac2 l2tp maximum-sessions-per-tunnel 4000
set client lac2 l2tp interface-id l2tp-encap-2
set client lac2 l2tp lcp-renegotiation
set client lac2 l2tp shared-secret "lac2-secret"
set client lac2 user-group-profile ce-l2tp-group-profile
top
edit access profile aaa-profile
set authentication-order radius
set radius authentication-server 172.21.146.93
set radius-server 172.21.146.93 secret "aaa-secret”
top
edit access address-assignment pool client-pool1 family inet
set network 192.168.1.1/16
set range lns-v4-pool-range low 192.168.1.1
set range lns-v4-pool-range high 192.168.255.255
top
edit access address-assignment pool client-ipv6-pool2 family inet6
set prefix 2001:DB8::/32
```

```
set range lns-v6-pool-range low 2001:DB8:1::/48
set range lns-v6-pool-range high 2001:DB8:ffff::/48
top
set interfaces ge-5/0/1 unit 11 vlan-id 11
set interfaces ge-5/0/1 unit 11 family inet address 10.1.1.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
top
set chassis fpc 5 pic 0 inline-services bandwidth 10g
set chassis fpc 5 pic 2 inline-services bandwidth 10g
top
edit interfaces si-5/0/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
edit interfaces si-5/2/0
set hierarchical-scheduler maximum-hierarchy-levels 2
set encapsulation generic-services
set unit 0 family inet
top
set services service-device-pools pool lns_p1 interface si-5/0/0
set services service-device-pools pool lns_p1 interface si-5/2/0
top
edit dynamic-profiles dyn-lns-profile2 routing-instances $junos-routing-instance
set interface $junos-interface-name
edit routing-options access route $junos-framed-route-ip-address-prefix
set next-hop $junos-framed-route-nexthop
set metric $junos-framed-route-cost
set preference $junos-framed-route-distance
up 2
edit access-internal route $junos-subscriber-ip-address
set qualified-next-hop $junos-interface-name
up 5
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set dial-options l2tp-interface-id l2tp-encapsulation
set dial-options dedicated
set family inet filter input $junos-input-filter
set family inet filter output $junos-output-filter
set family inet unnumbered-address $junos-loopback-interface
set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
up 3
edit protocols router-advertisement
set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix
top
[edit class-of-service]
edit rewrite-rules dscp rewriteDSCP forwarding-class expedited-forwarding
set loss-priority high code-point af11
set loss-priority high code-point af12
top
edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles tc-profile
set scheduler-map $junos-cos-scheduler-map
set shaping-rate $junos-cos-shaping-rate
set overhead-accounting $junos-cos-shaping-mode
set overhead-accounting bytes $junos-cos-byte-adjust
```


Copyright © 2016, Juniper Networks, Inc.

```
up
edit interfaces $junos-interface-ifd-name unit $junos-interface-unit
set forwarding-class expedited-forwarding
set output-traffic-control-profile tc-profile
set rewrite-rules dscp rewriteDSCP
edit interfaces si-5/0/0
set output-control-profile-remaining tc-profile
top
set services l2tp tunnel-group tg-dynamic l2tp-access-profile ce-l2tp-profile
set services l2tp tunnel-group tg-dynamic aaa-access-profile aaa-profile
set services l2tp tunnel-group tg-dynamic local-gateway address 10.1.1.2
set services l2tp tunnel-group tg-dynamic local-gateway gateway-name ce-lns
set services l2tp tunnel-group tg-dynamic service-device-pool lns_p1
set services l2tp tunnel-group tg-dynamic dynamic-profile dyn-lns-profile2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an L2TP LNS with inline service interfaces:

1. Configure a user group profile that defines the PPP configuration for tunnel subscribers.

```
[edit access]
user@host# edit group-profile ce-l2tp-group-profile
[edit access group-profile ce-l2tp-group-profile]
user@host# set ppp keepalive 30
user@host# set ppp idle-timeout 200
user@host# set ppp ppp-options chap
user@host# set ppp ppp-options pap
```

2. Configure an L2TP access profile that defines the L2TP parameters for each client LAC. This includes associating a user group profile with the client and specifying the identifier for the inline services logical interface that represents an L2TP session on the LNS.

```
[edit access profile ce-l2tp-profile client lac1]
user@host# set l2tp interface-id l2tp-encapsulation
user@host# set l2tp maximum-sessions-per-tunnel 1000
user@host# set l2tp shared-secret "lac1-secret"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile
[edit access profile ce-l2tp-profile client lac2]
user@host# set l2tp interface-id interface-id
user@host# set l2tp maximum-sessions-per-tunnel 4000
user@host# set l2tp shared-secret "lac2-secret"
user@host# set l2tp lcp-renegotiation
user@host# set user-group-profile ce-l2tp-group-profile
```

NOTE: If *<user-group-profile>* is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, will go down.

3. Configure a AAA access profile to override the global access profile for the order of AAA authentication methods and server attributes.

```
[edit access profile aaa-profile]
user@host# set authentication-order radius
user@host# set radius authentication-server 172.21.146.93
user@host# set radius-server 172.21.146.93 secret "aaa-secret"
```

4. Configure IPv4 and IPv6 address-assignment pools to allocate addresses for the clients (LACs).

```
[edit access address-assignment pool client-pool1 family inet]
user@host# set network 192.168.1.1/16
user@host# set range lns-v4-pool-range low 192.168.1.1 high 192.168.255.255
[edit access address-assignment pool client-ipv6-pool2 family inet6]
user@host# set prefix 2010:DB8::/32
user@host# set range lns-v6-pool-range low 2010:DB8:1::/48
user@host# set range lns-v6-pool-range high 2010:DB8:ffff::/48
```

5. Configure the peer interface to terminate the tunnel and the PPP server-side IPCP address (loopback address).

```
[edit interfaces ge-5/0/1
user@host# set vlan-tagging
user@host# set unit 11
[edit interfaces ge-5/0/1.11
user@host# set vlan-id 11
user@host# set family inet address 10.1.1.2/24
[edit interfaces lo0]
user@host# set unit 0 family inet address 127.0.0.1/32
```

6. Enable inline service interfaces on an MPC.

```
[edit chassis fpc 5]
user@host# set pic 0 inline-services bandwidth 10g
user@host# set pic 2 inline-services bandwidth 10g
```

7. Configure the anchor service interfaces with services encapsulation, hierarchical scheduling, and the address family.

```
[edit interfaces si-5/0/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
[edit interfaces si-5/2/0]
user@host# set hierarchical-scheduler maximum hierarchy-levels 2
user@host# set encapsulation generic-services
user@host# set unit 0 family inet
```

8. Configure a pool of service interfaces for dynamic LNS sessions.

```
[edit services service-device-pools pool lns_p1]
user@host# set interface si-5/0/0
user@host# set interface si-5/2/0
```

9. Configure a dynamic profile that dynamically creates L2TP logical interfaces for dual-stack subscribers.

```
[edit dynamic-profiles dyn-lns-profile2]
user@host# edit routing-instances $junos-routing-instance
```


Copyright © 2016, Juniper Networks, Inc.

```
user@host# set interface $junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"]
user@host# edit routing-options access route $junos-framed-route-ip-address-prefix
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
  routing-options access route "$junos-framed-route-ip-address-prefix"]
user@host# set next-hop $junos-framed-route-nexthop
user@host# set metric $junos-framed-route-cost
user@host# set preference $junos-framed-route-distance
[edit dynamic-profiles dyn-lns-profile2 routing-instances "$junos-routing-instance"
  routing-options access-internal]
user@host# set route $junos-subscriber-ip-address qualified-next-hop
  $junos-interface-name
[edit dynamic-profiles dyn-lns-profile2 interfaces "$junos-interface-ifd-name" unit
  "$junos-interface-unit"]
user@host# set dial-options l2tp-interface-id l2tp-encapsulation
user@host# set dial-options dedicated
user@host# set family inet unnumbered-address $junos-loopback-interface
user@host# set family inet filter input $junos-input-filter
user@host# set family inet filter output $junos-output-filter
user@host# set family inet6 address $junos-ipv6-address
set family inet6 filter input $junos-input-ipv6-filter
set family inet6 filter output $junos-output-ipv6-filter
[edit dynamic-profiles dyn-lns-profile2 protocols router-advertisement]
user@host# set interface $junos-interface-name prefix $junos-ipv6-ndra-prefix
```

10. Configure shaping, scheduling, and rewrite rules, and apply in the dynamic profile to tunnel traffic.

```
[edit class-of-service]
user@host# edit rewrite-rules dscp rewriteDSCP forwarding-class
  expedited-forwarding
user@host# set loss-priority high code-point af11
user@host# set loss-priority high code-point af12
[edit dynamic-profiles dyn-lns-profile2 class-of-service traffic-control-profiles
  tc-profile]
user@host# set scheduler-map $junos-cos-scheduler-map
user@host# set shaping-rate $junos-cos-shaping-rate
user@host# set overhead-accounting $junos-cos-shaping-mode
user@host# set overhead-accounting bytes $junos-cos-byte-adjust
[edit dynamic-profiles dyn-lns-profile2 class-of-service interfaces
  "$junos-interface-ifd-name" unit "$junos-interface-unit"]
user@host# set forwarding-class expedited-forwarding
user@host# set output-traffic-control-profile tc-profile
user@host# set rewrite-rules dscp rewriteDSCP
[edit class-of-service interfaces si-5/0/0]
user@host# set output-traffic-control-profile-remaining tc-profile
```

11. Configure the L2TP tunnel group to bring up dynamic LNS sessions using the pool of inline service interfaces to enable load-balancing.

```
[edit services l2tp tunnel-group tg-dynamic]
user@host# set l2tp-access-profile ce-l2tp-profile
user@host# set local-gateway address 10.1.1.2
user@host# set local-gateway gateway-name ce-lns
user@host# set aaa-access-profile aaa-profile
user@host# set dynamic-profile dyn-lns-profile2
user@host# set service-device-pool lns_p1
```

Results From configuration mode, confirm the access profile, group profile, AAA profile, and address-assignment pools configuration by entering the **show access** command. Confirm the inline services configuration by entering the **show chassis** command. Confirm the interface configuration by entering the **show interfaces** command. Confirm the dynamic profile configuration by entering the **show dynamic-profiles** command. Confirm the tunnel group configuration by entering the **show services l2tp** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show access
group-profile ce-l2tp-group-profile {
   ppp {
      idle-timeout 200;
      ppp-options {
         pap;
         chap;
      }
      keepalive 30;
   }
}
profile ce-l2tp-profile {
   client lac1 {
      l2tp {
         maximum-sessions-per-tunnel 1000;
         interface-id l2tp-encapsulation-1;
         lcp-renegotiation;
         shared-secret "$9$ZJGi.Pfz6/tmPtu1IleLxNbwgaZjmPQDi"; ## SECRET-DATA
      }
      user-group-profile ce-l2tp-group-profile;
   }
   client lac2 {
      l2tp {
         maximum-sessions-per-tunnel 4000;
         interface-id l2tp-encap-2;
         lcp-renegotiation;
         shared-secret ""$9$KCjvLNdVYoaUdVDi.m3ntuOREyevLdVY8X"; ## SECRET-DATA
      }
      user-group-profile ce-l2tp-group-profile;
   }
}
profile aaa-profile {
   authentication-order radius;
   radius-server {
      172.21.146.93 secret "$9$41JZjk.5Qz6k."; ## SECRET-DATA
   }
}
address-assignment {
   pool client-pool1 {
      family inet {
         network 192.168.1.1/16;
         range lns-v4-pool-range {
            low 192.168.1.1;
            high 192.168.255.255;
         }
```

 Copyright © 2016, Juniper Networks, Inc.

```
        }
    }
    pool client-ipv6-pool2 {
        family inet6 {
            prefix 2001:DB8::/32;
            range lns-v6-pool-range {
                low 2001:DB8:1::/48;
                high 2001:DB8:ffff::/48;
            }
        }
    }
}

[edit]
user@host# show chassis
fpc 5 {
    pic 0 {
        inline-services {
            bandwidth 10g;
        }
    }
    pic 2 {
        inline-services {
            bandwidth 10g;
        }
    }
}

[edit]
user@host# show interfaces
ge-5/0/1 {
    vlan-tagging;;
    unit 11 {
        vlan-id 11;
        family inet {
            address 10.1.1.2/24;
        }
    }
}
si-5/0/0 {
    hierarchical-scheduler maximum-hierarchy-levels 2;
    encapsulation generic-services;
    unit 0 {
        family inet;
    }
}
si-5/2/0 {
    hierarchical-scheduler maximum-hierarchy-levels 2;
    encapsulation generic-services;
    unit 0 {
        family inet;
    }
}
lo0 {
    unit 0 {
```

```
        family inet {
            address 127.0.0.1/32;
        }
    }
}
[edit]
user@host# show dynamic-profiles
dyn-lns-profile2 {
    routing-instances {
        "$junos-routing-instance" {
            interface "$junos-interface-name";
            routing-options {
                access {
                    route $junos-framed-route-ip-address-prefix {
                        next-hop "$junos-framed-route-nexthop";
                        metric "$junos-framed-route-cost";
                        preference "$junos-framed-route-distance";
                    }
                }
                access-internal {
                    route $junos-subscriber-ip-address {
                        qualified-next-hop "$junos-interface-name";
                    }
                }
            }
        }
    }
    interfaces {
        "$junos-interface-ifd-name" {
            unit "$junos-interface-unit" {
                dial-options {
                    l2tp-interface-id l2tp-encapsulation;
                    dedicated;
                }
                family inet {
                    filter {
                        input "$junos-input-filter";
                        output "$junos-output-filter";
                    }
                    unnumbered-address "$junos-loopback-interface";
                }
                family inet6 {
                    address $junos-ipv6-address;
                    input $junos-input-ipv6-filter;
                    output $junos-output-ipv6-filter;
                }
            }
        }
    }
    protocols {
        router-advertisement {
            interface "$junos-interface-name" {
                prefix $junos-ipv6-ndra-prefix;
            }
        }
    }
```


Copyright © 2016, Juniper Networks, Inc.

```
    class-of-service {
      rewrite-rules {
        dscp rewriteDSCP {
          forwarding-class expedited-forwarding {
            loss-priority high code-point af11
            loss-priority high code-point af12
          }
        }
      }
      traffic-control-profiles {
        tc-profile {
          scheduler-map "$junos-cos-scheduler-map";
          shaping-rate "$junos-cos-shaping-rate";
          overhead-accounting "$junos-cos-shaping-mode" bytes "$junos-cos-byte-adjust";
        }
      }
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-interface-unit" {
            forwarding-class expedited-forwarding;
            output-traffic-control-profile tc-profile;
            rewrite-rules {
              dscp rewriteDSCP;
            }
          }
        }
      }
    }
  }

  [edit]
  user@host# show services l2tp
  tunnel-group tg-dynamic {
    l2tp-access-profile ce-l2tp-profile;
    aaa-access-profile aaa-profile;
    local-gateway {
      address 10.1.1.2;
      gateway-name ce-lns;
    }
    service-device-pool lns_p1;
    dynamic-profile dyn-lns-profile2;
  }
```

When you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- L2TP for Subscriber Access Overview on page 123
- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP LAC on page 149

Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces

The L2TP tunnel group specifies attributes that apply to L2TP tunnels and sessions from a group of LAC clients. These attributes include the access profile used to validate L2TP connection requests made to the LNS on the local gateway address, a local access profile that overrides the global access profile, the keepalive timer, and whether the IP ToS value is reflected.

NOTE: **If you delete a tunnel group, all L2TP sessions in that tunnel group are terminated. If you change the value of the local-gateway-address, service-device-pool, or service-interface statements, all L2TP sessions using those settings are terminated. If you change or delete other statements at the [edit services l2tp tunnel-group *name*] hierarchy level, new tunnels you establish use the updated values but existing tunnels and sessions are not affected.**

To configure the LNS tunnel group:

1. Create the tunnel group.

   [edit services l2tp]
   user@host# **edit** tunnel-group ***name***

2. Specify the service anchor interface responsible for L2TP processing on the LNS.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** service-interface ***interface-name***

   This service anchor interface is required for static LNS sessions, and for dynamic LNS sessions that do not balance traffic across a pool of anchor interfaces. The interface is configured at the **[edit interfaces]** hierarchy level.

3. (Optional; for load-balancing dynamic LNS sessions only) Specify a pool of inline service anchor interfaces to enable load-balancing of L2TP traffic across the interfaces.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** service-device-pool ***pool-name***

   The pool is defined at the **[edit services service-device-pools]** hierarchy level.

4. (For dynamic LNS sessions only) Specify the name of the dynamic profile that defines and instantiates inline service interfaces for L2TP tunnels

   [edit services l2tp tunnel-group *name*]
   user@host# **set** dynamic-profile ***profile-name***

   The profile is defined at the **[edit dynamic-profiles]** hierarchy level.

5. Specify the access profile that validates all L2TP connection requests to the local gateway address.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** l2tp-access-profile ***profile-name***

Copyright © 2016, Juniper Networks, Inc.

6. Configure the local gateway address on the LNS; corresponds to the IP address that is used by LACs to identify the LNS.

   [edit services l2tp tunnel-group *name*]
   user@host# **set local-gateway** address ***address***

7. (Optional) Configure the local gateway name on the LNS, returned in the SCCRP message to the LAC. The name must match the remote gateway name configured on the LAC, or the tunnel cannot be created.

   [edit services l2tp tunnel-group *name*]
   user@host# **set local-gateway** gateway-name ***gateway-name***

8. (Optional) Configure the interval at which the LNS sends hello messages if it has received no messages from the LAC.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** hello-interval ***seconds***

9. (Optional) Specify a local access profile that overrides the global access profile to configure RADIUS server settings for the tunnel group.

   [edit services l2tp tunnel-group *name*]
   user@host# **set** aaa-access-profile ***profile-name***

   This local profile is configured at the **[edit access profile]** hierarchy level.

10. (Optional) Configure the LNS to reflect the IP ToS value from the inner IP header to the outer IP header (applies to CoS configurations).

    [edit services l2tp tunnel-group *name*]
    user@host# **set** tos-reflect

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Access Profile on the LNS on page 180

Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions

You can create a pool of inline service interfaces, also known as a *service device pool*, to enable load-balancing of L2TP traffic across the interfaces. The pool is supported for dynamic LNS configurations, where it provides a set of logical interfaces that can be dynamically created and allocated to L2TP sessions on the LNS. The pool is assigned to an LNS tunnel group. L2TP maintains the state of each inline service interface and uses a round-robin method to evenly distribute the load among available interfaces when new session requests are accepted.

NOTE: **Load balancing is available only for dynamically created subscriber interfaces.**

LNS sessions anchored on an MPC are not affected by a MIC failure as long as some other path to the peer LACs exists. If the MPC hosting the peer interface fails and there

is no path to peer LACs, the failure initiates termination and clean-up of all the sessions on the MPC.

If the MPC anchoring the LNS sessions itself fails, the Routing Engine does not relocate sessions to another slot and all sessions are terminated immediately. New sessions can come up on another available interface when the client retries.

To configure the service device pool:

1. Create the pool.

   [edit services service-device-pools]
   user@host# edit pool *pool-name*

2. Specify the inline service interfaces that make up the pool.

   [edit services service-device-pools pool *pool-name*]
   user@host# set interface *service-interface-name*
   user@host# set interface *service-interface-name*

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

Configuring a Dynamic Profile for Dynamic LNS Sessions

You can configure L2TP to dynamically assign inline service interfaces for L2TP tunnels. You must define one or more dynamic profiles and assign a profile to each tunnel group. Both IPv4-only and dual-stack IPv4/IPv6 interfaces are supported.

To configure the L2TP dynamic profile:

1. Create the dynamic profile.

   [edit]
   user@host# edit dynamic-profiles *profile-name*

2. Configure the interface to be dynamically assigned to the routing instance used by the tunneled PPP clients.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance”]
   user@host# set interface $junos-interface-name

3. Configure the routing options for access routes in the routing instance.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance” routing-options access]
   user@host# set route next-hop $junos-framed-route-nexthop
   user@host# set route metric $junos-framed-route-cost
   user@host# set route preference $junos-framed-route-distance

4. Configure the routing options for access-internal routes in the routing instance.

   [edit dynamic-profiles *profile-name* routing-instances “$junos-routing-instance” routing-options access-internal]
   user@host# set route $junos-subscriber-ip-address

 Copyright © 2016, Juniper Networks, Inc.

5. Define the interfaces used by the dynamic profile. The variable is dynamically replaced by one of the configured inline service interfaces.

   [edit dynamic-profiles *profile-name*]
   user@host# **set interfaces $junos-interface-ifd-name**

6. Configure the inline services logical interfaces to be dynamically instantiated.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name"]
   user@host# **set unit $junos-interface-unit**

7. Specify an identifier for the logical interfaces.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit]
   user@host# **set dial-options l2tp-interface-id *name***

8. Configure each logical interface to be used for only one session at a time.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit]
   user@host# **set dial-options dedicated**

9. Configure **family inet** as the address family for the logical interfaces and enable the local address on the LNS that provides local termination for the L2TP tunnel to be derived from the specified interface name.

   NOTE: For dynamic interfaces, you must specify the inet address family when you configure dial-options. If you configure an IPv6-pnly LNS stack, you can omit only the address portion of the inet address family configuration.

   [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit]
   user@host# **set family inet unnumbered-address $junos-loopback-interface**

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 22

Configuring IP Packet Fragment Reassembly

IP Packet Fragment Reassembly for L2TP Overview

You can configure the service interfaces on the MX Series routers with modular port concentrators (MPCs) to support reassemble fragmented IP packets for an L2TP connection. When packets are transmitted over an L2TP connection, the packets may be fragmented during transmission and need to be reassembled before they are processed further. Efficient reassembly is important for network throughput, scalability, and graceful response to congestion.

Fragmentation of IP packets for transmission and the need to reassemble the IP packets at a destination is a feature of how Layer 2 (the frame layer) and Layer 3 (the packet layer) operate. The maximum size of a frame, set by the Maximum Transmission Unit (MTU) value, and the maximum size of a packet are determined independently. Typically the packet size can far exceed the MTU size defined for the outgoing connection. If the packet size (data plus IP and other headers) exceeds the configured frame size (usually set by the transport medium limits), the packet must be fragmented and split across multiple frames for transmission. Frames are always processed immediately, when they arrive (if error-free), but packet fragments cannot be processed until the whole packet has been reassembled. Each packet fragment inside a frame series, except the last packet fragment, has the more fragments (MF) IP header bit set, indicating that this packet is part of a whole. The last packet fragment inside a frame does not have this MF bit set and therefore ends the fragment sequence. After all of the fragments of a packet have arrived, the entire packet can be reassembled.

In an L2TP connection, packets are transmitted between the L2TP Access Concentrator (LAC) and the L2TP Network Server (LNS). For an IP packet being transmitted over an L2TP connection, the packet is fragmented at the LAC, at an LNS, or at any intermediate router. IP reassembly parameters configured on the service interfaces of the LAC and the LNS determine how the fragments are reassembled at the service interfaces to ensure efficient reassembly over an L2TP connection.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 204
- *Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPCEs*
- ip-reassembly on page 345

Configuring IP Inline Reassembly for L2TP

This procedure shows how to configure a service interface on a LAC or LNS to reassemble fragmented IP packets. This example creates a service set that configures the IP reassembly parameters for L2TP fragments. The service set is then associated with the L2TP service.

Before you configure inline IP reassembly, be sure you have:

- Configured L2TP.
- Configured a valid service interface on the LAC or LNS.

To configure inline IP reassembly:

1. Configure the chassis-level bandwidth used by the inline services interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit chassis]
   user@host# set fpc 2 pic 1 inline-services bandwidth 10g
   ```

2. Configure the interface-level logical unit used by the inline services (**si-**) interface on the FPC and PIC slot for inline IP fragment reassembly.

   ```
   [edit interfaces]
   user@host# set si-2/1/0 unit 0 family inet
   user@host# set si-2/1/0 unit 0 service-domain inside
   ```

   NOTE: This configuration is not unique to L2TP. However, you must configure the family (inet) and service domain (inside) as shown.

3. Configure the service set (**set1**) for IP reassembly in the input match direction. (The **local** option loops the reassembled packets back to the local interface.)

   ```
   [edit services]
   user@host# set service-set set1

   [edit services service set ip-reassembly-set]
   user@host# set ip-reassembly-rules ipr_rule1
   user@host# set next-hop-service inside-service-interface si-2/1/0.0
   user@host# set next-hop-service outside-service-interface-type local
   ```


Copyright © 2016, Juniper Networks, Inc.

NOTE: You must configure both inside (si- interface) and outside type (local) service interfaces statements. The reassembly rule is not formulated outside of the service set; this statement simply initiates the reassembly process.

4. Configure the IP reassembly rule parameter

```
[edit services ip-reassembly]
user@host# set rule ipr_rule1 match-direction input;
```

5. Configure the service set (**set1**) for IP reassembly to bind to the L2TP service.

NOTE:
- The service set must be defined at the **[edit services]** hierarchy level.
- You cannot delete a service set instance if it is associated with an L2TP service.

```
[edit services l2tp]
user@host# set ip-reassembly service-set set1
```

Related Documentation

- IP Packet Fragment Reassembly for L2TP Overview on page 203
- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- *Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPCEs*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 23

Configuring High Availability in the L2TP Access Network

L2TP and Graceful Routing Engine Switchover

Graceful Routing Engine switchover (GRES) is supported on MX Series routers acting as either the L2TP LAC or LNS. In the event that L2TP (the l2tp-universal-edge process) restarts or that the router fails over from the active routing engine (RE) to the standby RE, L2TP graceful Routing Engine switchover ensures that the following occurs:

- The LAC and the LNS recover destinations, tunnels, and sessions that were already established at the time of the failure or restart.
- The LAC and the LNS respond to tunnel keepalive requests received during the switchover for established tunnels, but do not generate any keepalives until the switchover is complete.
- The LAC and the LNS delete all the tunnels and sessions that are not in the Established state.
- The LAC and the LNS reject requests to create new tunnels and sessions.
- The LAC and the LNS send another disconnect notification to the peer for sessions and tunnels that are already in the Disconnecting state at the time of the failure or restart. For sessions and tunnels that were coming up at that time, the LAC and LNS send a disconnect notification to the peer.
- The LAC and the LNS restart timers for the full timeout period for recovered L2TP destinations, tunnels, and sessions.

NOTE: Graceful Routing Engine switchover is supported only by L2TP LAC and LNS on MX Series routers. It is not supported by L2TP LNS on M Series routers.

Related Documentation

- L2TP Failover and Peer Resynchronization on page 208
- L2TP for Subscriber Access Overview on page 123

L2TP Failover and Peer Resynchronization

L2TP failover enables a failed L2TP endpoint to resynchronize with its nonfailed peer during recovery and restart of the L2TP protocol on the failed endpoint. L2TP failover is enabled by default.

The failover and L2TP peer resynchronization process does all of the following:

- Prevents the nonfailed endpoint from prematurely terminating a tunnel while the failed endpoint is recovering.
- Reestablishes the sequence numbers required for the operation of the L2TP control protocol.
- Resolves inconsistencies in the tunnel and session databases of the failed endpoint and the nonfailed endpoint.

The router supports both the L2TP failover protocol method (described in *RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*) and the L2TP silent failover method. The differences between these two methods are as follows:

- With the L2TP failover protocol method, both endpoints must support the method or recovery always fails. The L2TP failover protocol method also requires a nonfailed endpoint to wait an additional recovery time period while the failed endpoint is recovering to prevent the nonfailed endpoint from prematurely disconnecting the tunnel. The additional recovery period delays the detection of tunnel keepalive failures.
- Silent failover operates entirely within the failed endpoint and does not require nonfailed endpoint support—this improves interoperability between peers. Silent failover does not require additional recovery time by the nonfailed endpoint, which also eliminates the potential for degraded responsiveness to the loss of tunnel connectivity.

The default resynchronization method in Junos OS is *failover-protocol-fall-back-to-silent-failover*. The recovery method used depends on the results of the failover capability negotiation that takes place between L2TP peers when they establish a tunnel, which works as follows:

- L2TP on the LAC by default attempts to negotiate the L2TP failover protocol first. When L2TP determines that the remote peer supports the L2TP failover protocol, then the L2TP failover protocol method is used.
- When L2TP determines that the remote peer does not support the L2TP failover protocol, then the L2TP silent failover method is used. Falling back on this secondary method prevents the failover from forcing a disconnection of the tunnel to the peer and all its sessions.

You can change the default behavior by including the **disable-failover-protocol** statement at the **[edit services l2tp]** hierarchy level. This statement forces the LAC to operate only

 Copyright © 2016, Juniper Networks, Inc.

in silent failover mode. This configuration can be useful when routers that act as the LNS either are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it. However, when you issue this statement and the LNS supports only failover protocol, then the LAC cannot negotiate failover protocol, and recovery (failover protocol recovery initiated by the LNS) always fails.

Related Documentation

- L2TP and Graceful Routing Engine Switchover on page 207
- L2TP for Subscriber Access Overview on page 123

High Availability Using Unified ISSU in the L2TP Access Network

The unified in-service software upgrade (unified ISSU) feature supports the L2TP access model used by subscriber management. This support ensures that the router preserves active L2TP subscriber sessions and session services after a unified ISSU has completed.

See *Unified ISSU Concepts* for a description of the supported platforms and modules, CLI statements, and procedures you use to configure and initiate unified ISSU. You can use the **issu** flag with the **traceoptions** statement to trace subscriber management unified ISSU events. You can also use the **show system subscriber-management summary** command to display information about the unified ISSU state.

The LAC and the LNS support unified ISSU. When an upgrade is initiated, the LAC completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade. Subscriber logouts are recorded during the upgrade and are completed after the upgrade has completed.

Related Documentation

- Verifying and Monitoring Subscriber Management Unified ISSU State on page 91
- *Unified ISSU System Requirements*

Verifying and Monitoring Subscriber Management Unified ISSU State

Purpose Display the state of unified ISSU for subscriber management features.

Action The first example indicates that control plane quiescing as part of unified ISSU is not in progress (for example, unified ISSU has not been started, has already completed, or control plane queiscing has not started). The second example shows that unified ISSU is in progress and that a participating subscriber management daemon requires 198 seconds to quiesce the control plane.

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  IDLE
    ISSU State           IDLE
    ISSU Wait            0
```

```
user@host> show system subscriber-management summary
```

```
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

Related Documentation

- High Availability Using Unified ISSU in the PPP Access Network on page 115
- High Availability Using Unified ISSU in the DHCP Access Network on page 86
- High Availability Using Unified ISSU in the L2TP Access Network on page 209
- *Unified ISSU Concepts*

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 24

Monitoring and Managing L2TP for Subscriber Access

Verifying and Managing L2TP for Subscriber Access

Purpose View or clear information about L2TP tunnels and sessions.

Action

- To display a summary of L2TP tunnels, sessions, errors, and control and data packets:

  user@host> show services l2tp summary

- To display the L2TP destinations:

  user@host> show services l2tp destination

- To clear all L2TP destinations:

  user@host> clear services l2tp destination **all**

- To clear statistics for all L2TP tunnels belonging to a destination, tunnels belonging to a specified local-gateway address, and tunnels belonging to a specified peer-gateway address:

  user@host>**clear services l2tp destination statistics all**
  user@host>**clear services l2tp destination local-gateway 10.1.1.2**

- To display the L2TP sessions:

  user@host> show services l2tp session

- To clear all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

  user@host>clear services l2tp session **all**
  user@host>**clear services l2tp session local-session-id 40553**
  user@host>**clear services l2tp session local-gateway 10.1.1.2**
  user@host>**clear services l2tp session local-gateway-name lns-mx960**

- To clear statistics for all L2TP sessions, the session with a specified local session ID, or sessions associated with the local gateway specified by an IP address or name:

```
user@host>clear services l2tp session statistics all
user@host>clear services l2tp session statistics local-session-id 17967
user@host>clear services l2tp session statistics local-gateway 10.1.1.2
user@host>clear services l2tp session statistics local-gateway-name lns-mx960
```

- To display the L2TP tunnels:

```
user@host> show services l2tp tunnel
```

- To clear all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

```
user@host> clear services l2tp tunnel all
user@host>clear services l2tp tunnel local-tunnel-id 40553
user@host>clear services l2tp tunnel local-gateway 10.1.1.2
user@host>clear services l2tp tunnel local-gateway-name lns-mx960
```

- To clear statistics for all L2TP tunnels, the tunnel with a specified local tunnel ID, or tunnels associated with the local gateway specified by an IP address or name:

```
user@host> clear services l2tp tunnel statistics all
user@host>clear services l2tp tunnel statistics local-tunnel-id 40553
user@host>clear services l2tp tunnel statistics local-gateway 10.1.1.2
user@host>clear services l2tp tunnel statistics local-gateway-name lns-mx960
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- CLI Explorer.

Testing L2TP Tunnel Configurations from the LAC

You can test L2TP tunnel configurations on the LAC and successful subscriber authentication and tunneling without bringing up a PPP user and an associated tunnel.

Issue the **test services l2tp tunnel** command from CLI operational mode to map a subscriber to an L2TP tunnel, verify the L2TP tunnel configuration (both locally on the LAC and on a back-end server such as a RADIUS server), and verify that L2TP tunnels from the LAC can be established with the remote LNS.

The Junos OS LAC implementation enables you to configure multiple tunnels from which one tunnel is chosen for tunneling a PPP subscriber. You can use the **test services l2tp tunnel** command to test all possible tunnel configurations to verify that each can be established. Alternatively, you can test only a specific tunnel for the subscriber.

You must specify a configured subscriber username when you issue the command. The test generates a dummy password—*testpass*—for the subscriber, or you can optionally specify the password. The test verifies whether the subscriber identified by that username can be tunneled according to the tunnel configuration. If the subscriber can be tunneled, then the test verifies whether the L2TP tunnel can be established with the LNS according to the L2TP configuration.

Copyright © 2016, Juniper Networks, Inc.

You can optionally specify a tunnel ID, in which case only that tunnel is tested; the tunnel must be already configured for that username. If you omit this option, the test is applied to the full set of tunnel configurations that are returned for the username. The tunnel ID you specify is the same as that used by Tunnel-Assignment-Id (RADIUS attribute 82) and specified by the **identification** statement in the tunnel profile.

To test subscriber authentication and tunnel configuration:

- Specify only the username.

  Example 1:

  ```
  user@host> test services l2tp tunnel user test-user1@example.com
  Subscriber: test-user1@example.com, authentication failed
  ```

  The user failed authentication with the generated password and consequently was not tunneled.

  Example 2:

```
user@host> test services l2tp tunnel user user23@example.com
Subscriber: user23@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  test1tunnel    192.168.2.3      default         default         Up
  test2tunnel    192.168.30.3       default         default        Peer unresponsive
  test3tunnel    192.168.50.1       default           test          Up
```

  This user was authenticated with the generated password and successfully tunneled. A set of tunnels was found to be associated with that username and the entire set was tested.

- Specify the username and the user's configured password.

  ```
  user@host> test services l2tp tunnel user test-user1@example.com password grZ98#jW
  Subscriber: test-user1@example.com, authentication success, locally terminated
  ```

  The subscriber was authenticated. However, the user was terminated locally rather than tunneled; this means that no tunnel was found to be associated with the user.

- Specify the username and a particular tunnel for the subscriber.

  ```
  user@host> test services l2tp tunnel user rx37w@example.com tunnel-name ce-lac
  Subscriber: rx37w@example.com, authentication success, l2tp tunneled
    Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
    ce-lac         192.168.5.10      default         default         Up
  ```

  The subscriber was authenticated and tunneled. The specified tunnel was found for the subscriber and the tunnel was established, confirming the tunnel configuration.

- Specify the username, the user's configured password, and a tunnel.

  ```
  user@host> test services l2tp tunnel user fanta4-mfg-fan@example.com password dieda499
  tunnel-name tunnel5
  Subscriber: fanta4-mfg-fan@example.com, authentication success, l2tp tunneled
  ```

  The subscriber was authenticated and tunneled. The absence of tunnel information in the output indicates that the specified tunnel configuration does not exist.

Related Documentation

- L2TP for Subscriber Access Overview on page 123

Enabling Tunnel and Global Counters for SNMP Statistics Collection

By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in Table 16 on page 214 have a default value of zero.

Table 16: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 60,000 sessions, none of these statistics can be more than 30 minutes old.

BEST PRACTICE: **The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.**

To enable L2TP statistics collection for SNMP:

- Enable statistics collection.

```
[edit services l2tp]
user@host1# set enable-snmp-tunnel-statistics
```

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Copyright © 2016, Juniper Networks, Inc.

PART 4

Configuring Mobile IP Subscribers

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 25

Mobile IP and Subscriber Access Overview

Mobile IP Home Agent Elements and Behavior

Mobile IP is a tunneling-based solution that enhances the utility of Junos routing platforms at the edge of the network between fixed wire and wireless network domains. This tunneling-based solution enables a router on a user's home subnet to intercept and forward IP packets to users who roam beyond traditional network boundaries. Mobile IP is useful in environments where mobility is desired and the traditional land line dial-in model does not provide an adequate solution, and in environments where a wireless technology is used.

You configure Mobile IP home agent parameters in the **[edit services mobile-ip]** hierarchy level, the **[edit logical-systems *logical-system-name*]** hierarchy level, and the **[edit routing-instances *routing-instances-name*]** hierarchy level.

NOTE: **Currently, Junos OS does not support configuration of the Mobile IP foreign agent.**

Traditionally, IP addresses are associated with a fixed network location. To achieve mobility, the mobile node assumes a secondary IP address that matches the new network and redirects the traffic bound to the primary or home address to the mobile node's new network. In the Mobile IP architecture, the two agents that accomplish this task are the home agent and the foreign agent.

When a mobile node roams into a new, foreign network, it negotiates with the foreign agent to get a secondary IP address, which is referred to as the care-of address. The mobile node registers this care-of address with the home agent. The home agent then establishes a tunnel to the care-of address if the tunnel is not established earlier.

Copyright © 2016, Juniper Networks, Inc.

NOTE: You need to establish only one tunnel between the home agent and the care-of address. Demultiplexing of the traffic is done through IP address inspection.

Packets sent to the home address of the mobile node are redirected by the home agent through the tunnel to the care-of address at the foreign agent. The foreign agent routes the packets to the mobile node's home address. Figure 9 on page 218 illustrates this forwarding and routing process behavior. Although the traffic to the correspondent node comes from the foreign agent, to the correspondent node the traffic appears to come from the mobile node's home network.

Figure 9: Mobile IP Network Without Reverse Tunneling

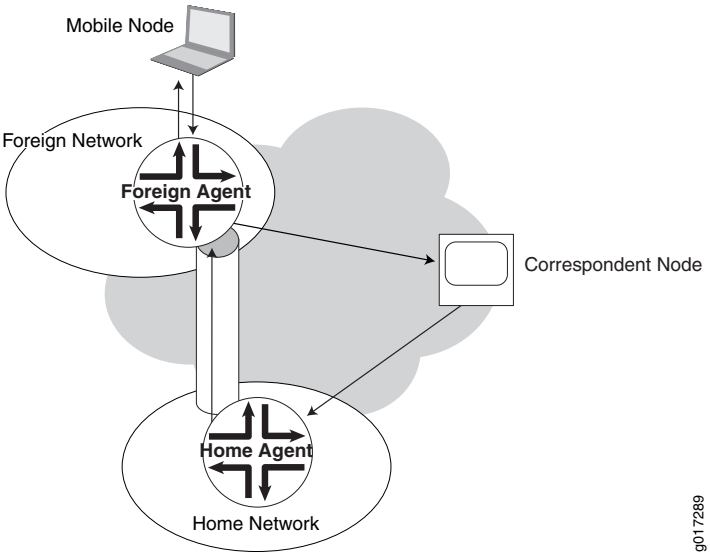

If the mobile node's home address is a private address or if the foreign agent implements ingress filtering, a reverse tunnel from the care-of address to the home agent is required. This reverse tunnel capability is negotiated between the foreign agent and the home agent when the mobile node requests registration. Traffic from a correspondent node to the mobile node is forwarded by the home agent through the foreign agent as in the other scenario. Figure 10 on page 219 shows how traffic from the mobile node to a correspondent node is tunneled from the foreign agent to the home agent and then routed to the correspondent node by the home agent.

Copyright © 2016, Juniper Networks, Inc.

Figure 10: Mobile IP Network With Reverse Tunneling

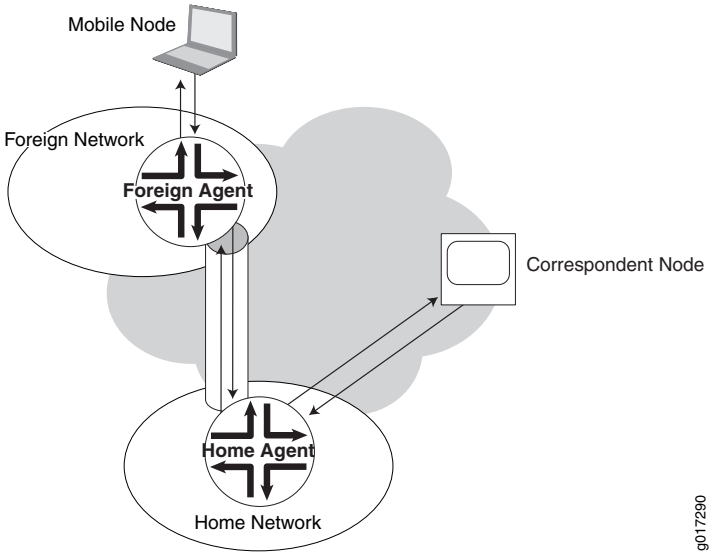

Mobile nodes typically belong to a virtual network, which is an address range or subnet that is not directly served by any physical, routed interface on the home network. These mobile nodes never return home to attach to a physical interface on the home agent. Traffic destined for the mobile node can be forwarded over any interface.

You can use the Mobile IP home agent feature to configure the home agent within the default router context with either local or AAA authentication. When you configure local authentication, you can also configure Mobile IP independently in any named routing instance in any configured logical router. When you configure AAA as the authentication method, you can configure Mobile IP only in the default router context.

The Mobile IP home agent can also receive, process, and send Worldwide Interoperability for Microwave Access (WiMAX) vendor-specific RADIUS attributes (VSAs). This feature enables Mobile IP home agent to work in a WiMAX home connectivity services network (H-CSN), to provide for mobility management at the IP layer.

The home agent handles the following tasks:

- Registration of mobile nodes
- Routing and forwarding of mobile node traffic

Related Documentation

- Mobile IP Registration on page 220
- Mobile IP Routing and Forwarding on page 224
- Mobile IP in the WiMAX Environment on page 225
- Configuring Mobile IP on page 229

Mobile IP Registration

The home agent receives the registration requests (RRQs) on UDP port 434. The registration request contains the home agent IP address. The home agent can support static home address allocation and dynamic home address allocation. The home agent can revoke a mobile node's registration. When this happens, the mobility binding is removed and the foreign agent is informed of the revocation so it can free up its resources. The foreign agent can send a registration revocation request to the home agent when the mobile node roams to another area. The revocation request can include a revocation support extension to indicate that it supports the revocation mechanism.

Home Address Assignment

The mobile node's home address can either be preconfigured, or dynamically allocated by the Mobile IP home agent. If a nonzero home address is preconfigured, the home agent processes the registration request using the home address and NAI (if the NAI is present).

If the home address is dynamically allocated, the mobile node submits a zero home address and requests the home agent to assign an IP address. The mobile node then uses the address provided by the home agent for subsequent registration requests, until the mobile node is rebooted or the registration expires.

Home address allocation is done by one of the existing authentication, authorization, and accounting (AAA) server back-end address mechanisms, such as:

- By RADIUS, in the Framed-IP-Address attribute
- From a local address pool returned by RADIUS in the Framed-Pool attribute

Authentication

The home agent authenticates the requests based on RFC 3344—IP Mobility Support for IPv4 (August 2002). By default, a AAA server is used for authentication; alternatively, you can configure local authentication parameters on the home agent. The mobile node authentication is verified and the authentication algorithm and key are retrieved by checking the security association indexed by the security parameter index (SPI) value. This verification results in the key and the authentication algorithm with which to compute an MD-5 message digest over the registration request. The Mobile IP home agent supports both HMAC-MD5 and keyed-MD5 authentication algorithms. When the result of this computation matches the authenticator, the mobile-home extension is authenticated. For local authentication, the key is limited to a maximum of 128 bits. For AAA authentication, the key can be longer depending on the maximum length configured on the AAA server.

When HA receives the access accept from the AAA, it extracts the MN-HA key from the response. The home agent does the MN-HA authentication extension processing based on the MN-HA key by running authentication algorithm (HMAC-MD5 or Keyed-MD5) on the message to compute a hash (authenticator), which is compared with the hash value in the MN-HA extension. If the hash value matches, the RRQ is considered authenticated.

 Copyright © 2016, Juniper Networks, Inc.

If a security association is configured for the foreign agent, the foreign-home authentication extension is verified; otherwise, authentication success is based only on the mobile-home authenticator.

The home agent checks the identification (ID) field to verify that a registration message has been freshly generated by the mobile node, and is not simply being replayed by an attacker from some previous registration. The ID field represents a 64-bit Network Time Protocol (NTP)-formatted time value. The configured replay timestamp defines the tolerance time window in seconds by which a registration request timestamp and the local time of the HA can differ. By default, the timestamp must be within 7 seconds of the replay tolerance configured for the mobile node or, if that is configured, the timestamp tolerance of the home agent itself.

Reauthentication

Reauthentication is not currently supported by the authentication process. Mobile IP caches a security association for each mobile node, which helps overcome this limitation. When a mobile node requests re-registration or de-registration, Mobile IP refers to the cached security association for that mobile node and performs MD5 message authentication.

When the security association for the mobile node changes after the node is authenticated, the cache entry is not invalidated. Consequently, the mobile node's RRQ is rejected. In this case you must clear the binding with the mobile node so that it can de-register and then log in.

RADIUS server configuration changes relating to the subscriber do not propagate to the cache. In this case you must clear the binding with the mobile node so that it can de-register and then log in.

AAA Authentication

You can store the security associations and configuration information remotely on a RADIUS server. The home agent applies the authentication algorithm and security key to the mobile node's message. The AAA server uses Juniper Networks vendor-specific attributes (VSAs; vendor ID 4874) listed in Table 17 on page 221. These VSAs are mandatory in the reply to provide the appropriate authentication algorithm and the secure key for the authentication request. If the security parameters are not retrieved, then the request for mobility service is rejected, a security violation error is logged, and no registration reply is generated.

Table 17: Juniper Networks VSAs Used by Mobile IP

Attribute Number	Attribute Name	Description	Value
26–84	Mobile-IP-Algorithm	Authentication algorithm used for Mobile-IP registration	integer: 4-octet
26–85	Mobile-IP-SPI	Security parameter index for Mobile IP registration	integer: 4-octet

Copyright © 2016, Juniper Networks, Inc.

Table 17: Juniper Networks VSAs Used by Mobile IP *(continued)*

Attribute Number	Attribute Name	Description	Value
26–86	Mobile-IP-Key	Security association MD5 key for Mobile IP registration	string: key
26–87	Mobile-IP-Replay	Replay timestamp for Mobile IP registration	integer: 4-octet
26–89	Mobile-IP-Lifetime	Registration lifetime for Mobile IP registration	integer: 4-octet

AAA authentication is accomplished by generating a AAA access-request to a AAA server. This is the default authentication mode, but you can include the **authenticate order aaa** statement at the **[edit services mobile-ip]** hierarchy level to explicitly configure AAA authentication. You cannot configure a fallback mechanism for AAA authentication. If the AAA request times out, the home agent does not fall back on the local router to determine the authentication parameters. The registration request is rejected. When the message is authenticated, the AAA server always returns either the Framed-IP-Address or Framed-Pool attribute for the user.

The presence of the mobile node's NAI and home IP address in the authentication request that the home agent sends to the AAA server is determined by their presence in the mobile node RRQ received by the home agent:

- When both the NAI and home IP address of the mobile node are present in the registration request, then the authentication request from Mobile IP to AAA has the NAI as the user name.
- When only the NAI is present in the registration request, then the NAI is used as the user name.
- When only the IP address (home address) is present in the registration request, then the IP address is used as the user name.
- When both the NAI address and the IP address are missing from the registration request, then the registration request is rejected.

Local Authentication

As an alternative to the default authentication by AAA server, you can store the security associations and configuration information locally on the router hosting the home agent. Local authentication is accomplished by querying the locally configured security parameters for the mobile node. The home agent applies the authentication algorithm and security key to the mobile node's message. If the security parameters are not available or do not match the RRQ, then the request for mobility service is rejected, a security violation error is logged, and no registration reply is generated.

For local authentication, include the **authenticate order local** statement at the **[edit services mobile-ip]** hierarchy level. You cannot configure a fallback mechanism for local

 Copyright © 2016, Juniper Networks, Inc.

authentication. If the local authentication fails, the home agent does not fall back on the AAA server to determine the authentication parameters. The registration request is rejected. Include the **peer** statement at the **[edit services mobile-ip]** hierarchy level to configure the authentication attributes on the home agent for a user identified by IP address or network address identifier (NAI). This user can be a mobile node or a foreign agent.

The authentication attributes include a security parameter index (SPI) to identify a particular security context between the home agent and the mobile node or foreign agent among the contexts available in the mobility security association. Associated with each SPI is the MD5 algorithm and key used to authenticate messages from the mobile node or foreign agent. You can also configure the replay timestamp tolerance for the mobile node or foreign agent.

When local authentication is configured, you can configure Mobile IP independently in any named routing instance in any configured logical router. All Mobile IP statements are available in those contexts, except for the **order aaa** statement at the **[edit services mobile-ip authenticate]** hierarchy level.

Accounting

The Junos Mobile IP home agent application supports time-based accounting for Mobile IP subscribers. Include the **statistics time** statement in the subscriber access profile at the **[edit access profile profile-name accounting]** hierarchy level. Time-based accounting for Mobile IP subscribers also requires that you include the **authenticate order aaa** statement at the **[edit services mobile-ip]** hierarchy level. Accounting begins when the Mobile IP home agent registers the mobile node and creates a binding with the mobile node.

Accounting stops when the binding is deleted. Any of the following actions can cause the binding to be deleted:

- The mobile user logs off.
- The binding lifetime expires.
- The mobile node is deregistered for any reason.
- The foreign agent sends a revocation message.

The Acct-Start message the home agent sends to the AAA server includes the network address identifier (NAI) in the User-Name attribute and the home address of the mobile IP node in the Framed-IP-Address attribute. The Acct-Stop message additionally includes the Acct-Session-Id and Acct-Session-Time attributes.

You cannot currently configure time-based accounting for only the Mobile IP service in a given logical router or routing instance. Enabling time-based accounting for Mobile IP also enables time-based accounting for all other services that are configured in that logical router or routing instance. If you do not want time-based accounting to apply to other services, then you must configure those services in a different logical router or routing instance.

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- For information about the specific Juniper Networks VSAs used for Mobile IP RADIUS-based authentication, see *Juniper Networks VSAs Supported by the AAA Service Framework*
- Mobile IP Home Agent Elements and Behavior on page 217
- Mobile IP Routing and Forwarding on page 224
- Mobile IP in the WiMAX Environment on page 225
- Configuring Mobile IP on page 229

Mobile IP Routing and Forwarding

Mobile IP employs a care-of address to process traffic for the mobile node.

The mobile node acquires the a care-of address from the foreign agent. The care-of address is reachable from the mobile node, and routable from the home agent. The mobile node includes the care-of address in its registration request to the home agent. After AAA or local authentication successfully processes and authenticates the RRQ and provides both the authorization parameters for the mobile node and an IP address, the home agent then sets up the data path for the mobile node and sends back a registration reply (RRP) confirming successful registration of the mobile node.

When the foreign agent receives the successful RRP from the home agent, the foreign agent sets up the data path for the mobile node. Then it sends the RRP to the mobile node to acknowledge that the mobile node is now successfully registered and the data path between the home agent and the mobile node is in place.

The home agent supports generic routing encapsulation (GRE) and IP-in-IP tunnel encapsulation for forward and reverse tunneling. The tunnels must be statically configured. When packets destined for the mobile node reach a home agent, the home agent encapsulates the packets and tunnels them to the care-of address. Packets that exceed the maximum transmission unit (MTU) value of the tunnel are dropped and an ICMP error message is sent to the source IP address. Packets without an access route are returned to the source with an ICMP destination unreachable error message. For reverse tunnels, packets are de-tunneled and forwarded towards the next hop to the destination address.

Mobile IP does not support graceful Routing Engine switchover (GRES). It handles the rebooting of processes in the following ways:

- Mobile IP process—After Mobile IP completes a restart, it removes the Mobile IP subscriber entries from AAA and the session database. When that is complete, Mobile IP can process new mobile node registration requests.
- AAA process—After AAA completes a restart, Mobile IP removes all subscriber data held internally by AAA and all corresponding session database entries.
- Routing protocol process—When the connection between the routing protocol process and Mobile IP is lost, Mobile IP responds by clearing the mobile node bindings that are associated with the logical system in which the routing protocol process restarted. The

 Copyright © 2016, Juniper Networks, Inc.

routing protocol process maintains routes to mobile nodes during the restart. The routing protocol process flushes these routes if they are not reinstalled after the restart completes and before the stale route timer expires.

Related Documentation

- Mobile IP Home Agent Elements and Behavior on page 217
- Mobile IP Registration on page 220
- Mobile IP in the WiMAX Environment on page 225
- Configuring Mobile IP on page 229

Mobile IP in the WiMAX Environment

Worldwide Interoperability for Microwave Access (WiMAX) is the international standard for wide area radio access networks. It provides a framework for networks that are implemented in different ways to successfully interoperate with mobile subscribers that roam among the networks. This interoperability enables the subscribers to be authenticated by their home network wherever they roam, and to receive the services for which they are authorized.

The Mobile IP home agent can operate in either of two access modes, generic and WiMAX. The generic access type is appropriate when the home agent is deployed in a generic Mobile IP home network. When deployed as a home agent in a WiMAX home connectivity services network (H-CSN), you must configure the WiMAX access type. The WiMAX access type enables the Mobile IP home agent to receive, process, and send WiMAX vendor-specific attributes (VSAs) that are used by AAA and the RADIUS server to authenticate the mobile subscriber. When the access type is generic, the Mobile IP home agent cannot handle these VSAs.

NOTE: **The Mobile IP configuration for WiMAX requires that AAA be used for the authentication method. For that reason, WiMAX is available only in the default router context.**

A WiMAX H-CSN is analogous to the Mobile IP home network for non-WiMAX implementations. When WiMAX is enabled for the Mobile IP home agent in an H-CSN, the Mobile IP home agent triggers subscriber authentication when the agent receives the registration request. The home agent stores WiMAX Forum (vendor ID 24757) vendor-specific attributes (VSAs) listed in Table 18 on page 226 in the session database based on the registration request.

Table 18: WiMAX Forum VSAs used by Mobile IP

Attribute Number	Attribute Name	Description	Value
26–1	WiMAX-Capability	Identifies the WiMAX capabilities supported by the home agent (sent in the Access-Request message). In an Access-Accept message, Identifies the capabilities selected by the RADIUS server (returned in the Access-Accept message).	string or integer
26–6	hHA-IP-MIP4	IP address of the home home agent (hHA) making the request	octet string: IP address
26–10	MN-HA-MIP4-KEY	MN-hHA key sent by the RADIUS server for validation by the home agent	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key
26–11	MN-HA-MIP4-SPI	Security parameter index (SPI) associated with the MN-HA-MIP4 key	integer: 4-octet
26–15	hHA-RK-KEY	Key used by the NAS to generate FA-HA keys	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key
26–16	hHA-RK-SPI	SPI associated with the hHA-RK key	integer: 4-octet
26–17	HA-RK-Lifetime	Lifetime of the hHA-RK key and derived keys	integer: 4-octet
26–18	RRQ-HA-IP	IP address of the home agent contained in the Mobile IP registration request or the binding update	octet string: IP address

Copyright © 2016, Juniper Networks, Inc.

Table 18: WiMAX Forum VSAs used by Mobile IP *(continued)*

Attribute Number	Attribute Name	Description	Value
26–19	RRQ-MN-HA-KEY	The MN-HA key bound to the home agent IP address as reported by the RRQ-HA-IP attribute. Used to validate the MN-HA-AE of the Mobile IP registration request.	integer: 2-octet salt followed by 16-octet encrypted MN-hHA hash key

The home agent requests AAA to fetch the corresponding WiMAX-related information from the RADIUS server. The AAA client sends an Access-Request message to the server. The RADIUS server responds with the necessary WiMAX information, such as the MN-HA key and the HA-RK key, and then the AAA client passes the response to the home agent. The Mobile IP home agent verifies the response received from AAA, processes the registration request, and then grants, extends, or denies subscriber registration.

Figure 11 on page 227 shows the elements of a sample WiMAX topology.

Figure 11: Sample Mobile IP WiMAX Topology

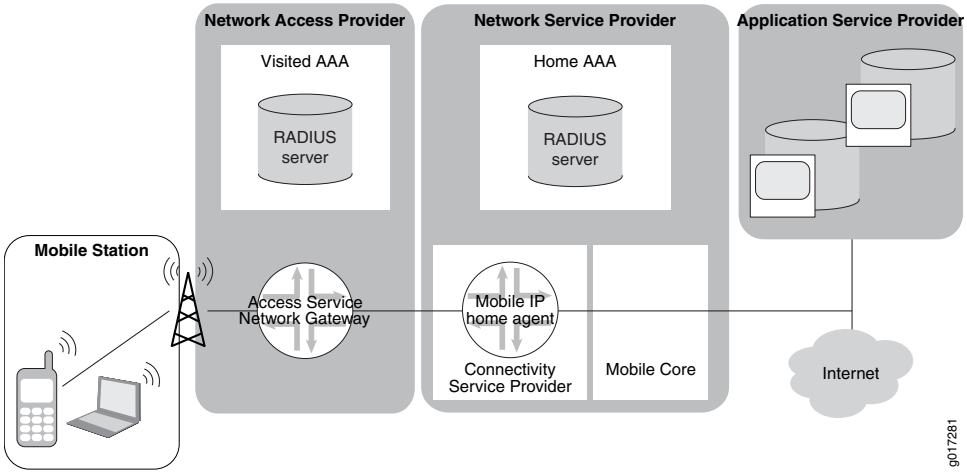

The Mobile IP subscriber registration flow is a four-step process.

1. The access service network gateway (ASN-GW) sends the subscriber registration request from the mobile node to the Mobile IP home agent. The registration request is protected by the MN-HA authentication extension and the FA-HA authentication extension.
2. The home agent requests that the RADIUS server send the cryptographic keys for the Mobile IP session identified by user@realm. The home agent notifies the RADIUS server that it seeks to source IP session-based accounting messages.

Copyright © 2016, Juniper Networks, Inc.

3. The RADIUS server agrees to use IP session-based accounting, provides the requested cryptographic keys, and sends the AAA-Session-ID for this session.
4. The home agent replies to the Mobile IP registration request.

Reauthentication of WiMAX subscribers is not currently supported.

You can configure the Mobile IP home agent for WiMAX access by including the **wimax** statement at the **[edit services mobile-ip access-type**] hierarchy level. You can prevent the Mobile IP home agent from being able to process WiMAX VSAs by either removing the **wimax** statement at the **[edit services mobile-ip access-type**] hierarchy level or by including the **generic** statement at the **[edit services mobile-ip access-type**] hierarchy level. The default access type for Mobile IP home agent is generic.

Related Documentation

- For information about the specific Juniper Networks VSAs used for Mobile IP RADIUS-based authentication, see *Juniper Networks VSAs Supported by the AAA Service Framework*
- Mobile IP Home Agent Elements and Behavior on page 217
- Mobile IP Registration on page 220
- Mobile IP Routing and Forwarding on page 224
- Configuring Mobile IP on page 229

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 26

Configuring Mobile IP for Mobile Subscriber Access

Configuring Mobile IP

You can configure Mobile IP to provide mobility for subscribers in IP networks. The Mobile IP home agent authenticates registration requests from mobile users and forward traffic to them at their care-of address without having to advertise that address to the wider network.

To configure Mobile IP for mobile subscriber access:

1. Configure the authentication method for registration requests, local or AAA.

   See "Configuring the Mobile IP Authentication Method" on page 230.

2. Configure the Mobile IP home agent.

   See "Configuring the Mobile IP Home Agent" on page 230.

3. Configure the authentication attributes for the mobile node.

   See "Configuring the Local Authentication Attributes for the Mobile Node" on page 231.

4. Configure accounting for Mobile IP subscribers.

   See "Configuring Accounting for Mobile IP Subscribers" on page 231

5. Configure the dynamic reassignment of the mobile node to another home agent.

   See "Configuring Dynamic Home Assignment for the Mobile Node" on page 232.

6. Configure the access type for Mobile IP.

Copyright © 2016, Juniper Networks, Inc.

See "Configuring the Access Type for Mobile IP" on page 232.

7. Configure trace options for troubleshooting the configuration.

   See "Tracing Mobile IP Operations for Subscriber Access" on page 257.

Configuring the Mobile IP Authentication Method

You can configure Mobile IP to authenticate registration requests from mobile nodes by either the locally configured attributes or a AAA server. AAA server authentication is the default method.

NOTE: AAA server authentication is available only in the default router context. Local authentication is available in both default and nondefault router contexts.

To configure the Mobile IP authentication method:

- Specify either local or AAA authentication.

```
[edit services mobile-ip]
user@host# set authenticate order local
```

Related Documentation

- Configuring Mobile IP on page 229

Configuring the Mobile IP Home Agent

To configure the home agent for a Mobile IP virtual network:

1. Configure the loopback IP address that is used as the home agent IP address.

```
[edit services mobile-ip home-agent virtual-network]
user@host# set home-agent-address 10.5.5.0
```

2. (Optional) Configure the maximum lifetime that the home agent accepts in any registration request from a mobile node.

```
[edit services mobile-ip home-agent virtual-network]
user@host# set home-agent-address 10.5.5.0 registration-lifetime 100
```

3. (Optional) Configure a timestamp tolerance for registration replay protection.

```
[edit services mobile-ip home-agent virtual-network]
user@host# set home-agent-address 10.5.5.0 timestamp-tolerance 200
```

4. Configure whether the home agent can revoke a mobile node's registration to deactivate the node.

```
[edit services mobile-ip home-agent virtual-network]
user@host# set home-agent-address 10.5.5.0 revocation-required
```

5. Specify the interfaces on which the home agent accepts registration requests.

```
[edit services mobile-ip home-agent]
```

 Copyright © 2016, Juniper Networks, Inc.

```
user@host# set enable-service ge-0/0/1.0
user@host# set enable-service ge-0/0/2.0
user@host# set enable-service ge-0/0/3.0
user@host# set enable-service ge-0/0/4.0
```

Related Documentation

- Configuring Mobile IP on page 229

Configuring the Local Authentication Attributes for the Mobile Node

You specify for each mobile node several attributes that enable authentication of registration requests from the node. These attributes include security association context for the peering relationship, the entity type of the node, the encryption algorithm and key used to authenticate the request, and replay protection.

To configure authentication attributes for the mobile node:

1. Configure the peer entity for the security parameter.

   ```
   [edit services mobile-ip]
   user@host# set peer ip-address 10.4.2.20 spi 500 entity-type mobility-agent
   ```

2. Configure the algorithm used for authenticating Mobile IP messages. By default, the hmac-md5 algorithm is used.

   ```
   [edit services mobile-ip]
   user@host# set peer ip-address 10.4.2.20 spi 500 algorithm md5
   ```

3. Configure the authentication key for the security association, in either HEX or ASCII format.

   ```
   [edit services mobile-ip]
   user@host# set peer ip-address 10.4.2.20 spi 500 key ascii xf125j9m
   ```

4. Configure a timestamp tolerance for registration replay protection or specify that the timestamp tolerance be taken from the value configured on the home agent.

   ```
   [edit services mobile-ip]
   user@host# set peer ip-address 10.4.2.20 spi 500 replay-method timestamp tolerance
     250
   ```

Related Documentation

- Configuring Mobile IP on page 229

Configuring Accounting for Mobile IP Subscribers

You can configure time-based accounting to track the subscriber sessions of Mobile IP subscribers.

To configure Mobile IP accounting:

1. Configure the IP address for the RADIUS accounting server.

   ```
   [edit access profile mip-win4]
   user@host# set radius accounting-server 192.168.20.5
   ```

2. Specify RADIUS as the accounting method for Mobile IP subscribers.

```
[edit access profile mip-win4 accounting]
user@host# set order radius
```

3. Specify time-based accounting for the access profile used for the subscriber.

```
[edit access profile mip-win4 accounting]
user@host# set statistics time
```

Related Documentation

- Configuring Mobile IP on page 229
- *Specifying the Authentication and Accounting Methods for Subscriber Access*
- *Configuring Per-Subscriber Session Accounting*
- *Configuring RADIUS Server Parameters for Subscriber Access*

Configuring Dynamic Home Assignment for the Mobile Node

The mobile node can request that the home agent dynamically assign an IP address for the home agent. The mobile node uses this address for the home agent in all subsequent registration requests until the registration expires or the mobile node is rebooted.

To configure the IP address to be used by the mobile node for the home agent:

- Configure the IP address for the specified mobile node.

```
[edit services mobile–ip]
user@host# set dynamic-home-assignment home-agent nai bws@example.com
  home-agent 192.168.4.5
```

Related Documentation

- Configuring Mobile IP on page 229

Configuring the Access Type for Mobile IP

You can configure the Mobile IP home agent to operate in a Worldwide Interoperability for Microwave Access (WiMAX) home connectivity services network (H-CSN). This configuration enables the home agent to receive, process, and send WiMAX VSAs for subscriber authentication and registration. By default, Mobile IP cannot process the WiMAX VSAs. For operation in non-WiMAX environments, you can return it to this mode by configuring the **generic** access type.

NOTE: **The Mobile IP configuration for WiMAX requires that AAA be used for the authentication method. For that reason, WiMAX is available only in the default router context.**

To configure the access type, do one of the following:

- Configure generic operation.

Copyright © 2016, Juniper Networks, Inc.

```
[edit services mobile-ip]
user@host# set access-type generic
```

- Configure WiMAX operation.

```
[edit services mobile-ip]
user@host# set access-type wimax
```

Related Documentation

- Configuring Mobile IP on page 229

Copyright © 2016, Juniper Networks, Inc.

PART 5

Troubleshooting

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 27

Configuring PPP Log Files

Configuring the Number and Size of PPP Service Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

  [edit protocols ppp-service traceoptions]
  user@host# **set file ppp-service_1 _logfile_1 files 20 size 2097152**

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Configuring Access to the PPP Service Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

  [edit protocols ppp-service traceoptions]
  user@host# **set file ppp-service_1 _logfile_1 world-readable**

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

  [edit protocols ppp-service traceoptions]
  user@host# **set file ppp-service_1 _logfile_1 no-world-readable**

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Configuring the Severity Level to Filter Which PPP Service Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the

 Copyright © 2016, Juniper Networks, Inc.

messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit protocols ppp-service traceoptions]
  user@host# set level *severity*

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 28

Configuring PPP Trace Flags and Operations

Tracing PPP Service Operations for Subscriber Access

The Junos OS trace feature tracks PPP service operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **jpppd**. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file ***filename*** reaches 128 kilobytes (KB), it is compressed and renamed ***filename*.0.gz**. Subsequent events are logged in a new file called ***filename***, until it reaches capacity again. At this point, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the System Log Explorer.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPP service tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the PPP Service Trace Log Filename" on page 242.

2. (Optional) Configure the number and size of trace logs.

   See "Configuring the Number and Size of PPP Service Log Files" on page 237.

3. (Optional) Configure user access to trace logs.

   See "Configuring Access to the PPP Service Log File" on page 238.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See *Configuring a Regular Expression for PPP Service Messages to Be Logged*.

5. (Optional) Configure flags to specify which events are logged.

   See "Configuring the PPP Service Tracing Flags" on page 242.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which PPP Service Messages Are Logged" on page 238.

Configuring the PPP Service Trace Log Filename

By default, the name of the file that records trace output for PPP service is **jpppd**. You can specify a different name with the **file** option.

To configure the filename for PPP service tracing operations:

- Specify the name of the file used for the trace output.

  [edit protocols ppp-service traceoptions]
  user@host# **set file ppp-service_logfile_1**

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Configuring the PPP Service Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit protocols ppp-service traceoptions]
  user@host# **set flag *flag***

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Configuring Subscriber Filtering for PPP Service Trace Operations

You can apply filters to the PPP service to limit tracing to particular subscribers or domains. Subscriber filtering simplifies troubleshooting in a scaled environment by enabling you to focus on a reduced set of trace results.

For subscriber usernames that have the expected form of ***user@domain***, you can filter on the user, the domain, or both. You can use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms to match a greater number of subscribers.

NOTE: **You cannot filter results using a wildcard in the middle of the user or domain terms. For example, the following uses of the wildcard are not supported: tom*25@example.com, tom125@ex*.com.**

When you enable filtering by username, traces that have insufficient information to determine the username are automatically excluded.

To configure subscriber filtering:

- Specify the filter.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user *user@domain***

Consider the following examples of using the wildcard for filtering:

- Filter results for the specific subscriber with the username, tom@example.com.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user tom@example.com**

- Filter results for all subscribers whose username begins with tom.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user tom***

- Filter results for all subscribers whose username ends with tom.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user *tom**

- Filter results for subscribers with the username tom at all domains beginning with ex.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user tom@ex***

- Filter results for all subscribers at all domains that end with ample.com.

  [edit protocols ppp-service traceoptions]
  user@host# **set filter user *ample.com**

- Filter results for all subscribers whose username begins with tom at domains that end with example.com.

  ```
  [edit protocols ppp-service traceoptions]
  user@host# set filter user tom*@*example.com
  ```

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 29

Configuring L2TP Log Files

Configuring the Number and Size of L2TP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 files 20 size 2097152**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Configuring Access to the L2TP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 world-readable**

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 no-world-readable**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Configuring a Regular Expression for L2TP Messages to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_1 _logfile_1 match *regex***

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Configuring the Severity Level to Filter Which L2TP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

Copyright © 2016, Juniper Networks, Inc.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit services l2tp traceoptions]
  user@host# **set level** *severity*

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 30

Configuring L2TP Trace Flags and Operations

Tracing L2TP Operations for Subscriber Access

The Junos OS trace feature tracks L2TP operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

NOTE: This topic refers to tracing L2TP operations on MX Series routers. To trace L2TP operations on M Series routers, see *Tracing L2TP Operations*.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **jl2tpd**. You can specify a different filename, but you cannot change the directory in which trace files are located.
2. When the trace log file ***filename*** reaches 128 kilobytes (KB), it is compressed and renamed ***filename*.0.gz**. Subsequent events are logged in a new file called ***filename***, until it reaches capacity again. At this point, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the System Log Explorer.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure L2TP tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the L2TP Trace Log Filename" on page 250.

2. (Optional) Configure the number and size of trace logs.

   See "Configuring the Number and Size of L2TP Log Files" on page 245.

3. (Optional) Configure user access to trace logs.

   See "Configuring Access to the L2TP Log File" on page 246.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for L2TP Messages to Be Logged" on page 246.

5. (Optional) Configure flags to specify which events are logged.

   See "Configuring the L2TP Tracing Flags" on page 250.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which L2TP Messages Are Logged" on page 246.

Configuring the L2TP Trace Log Filename

By default, the name of the file that records trace output for L2TP is **jl2tpd**. You can specify a different name with the **file** option.

To configure the filename for L2TP tracing operations:

- Specify the name of the file used for the trace output.

  [edit services l2tp traceoptions]
  user@host# **set file l2tp_logfile_1**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Configuring the L2TP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit services l2tp traceoptions]
  user@host# **set flag *flag***

Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Configuring Subscriber Filtering for L2TP Trace Operations

You can apply filters to L2TP to limit tracing to particular subscribers or domains. Subscriber filtering simplifies troubleshooting in a scaled environment by enabling you to focus on a reduced set of trace results.

For subscriber usernames that have the expected form of ***user@domain***, you can filter on the user, the domain, or both. You can use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms to match a greater number of subscribers.

NOTE: **You cannot filter results using a wildcard in the middle of the user or domain terms. For example, the following uses of the wildcard are not supported: tom*25@example.com, tom125@ex*.com.**

When you enable filtering by username, traces that have insufficient information to determine the username are automatically excluded.

To configure subscriber filtering:

- Specify the filter.

  [edit services l2tp traceoptions]
  user@host# **set filter user *user@domain***

  NOTE: **This syntax is different than the syntax used to filter subscribers on M Series routers.**

Consider the following examples of using the wildcard for filtering:

- Filter results for the specific subscriber with the username, tom@example.com.

  [edit services l2tp traceoptions]
  user@host# **set filter user tom@example.com**

- Filter results for all subscribers whose username begins with tom.

  [edit services l2tp traceoptions]
  user@host# **set filter user tom***

- Filter results for all subscribers whose username ends with tom.

  [edit services l2tp traceoptions]
  user@host# **set filter user *tom**

- Filter results for subscribers with the username tom at all domains beginning with ex.

  [edit services l2tp traceoptions]
  user@host# **set filter user tom@ex***

- Filter results for all subscribers at all domains that end with ample.com.

  [edit services l2tp traceoptions]
  user@host# **set filter user *ample.com**

- Filter results for all subscribers whose username begins with tom at domains that end with example.com.

  [edit services l2tp traceoptions]
  user@host# **set filter user tom*@*example.com**

Related Documentation

- Tracing L2TP Operations for Subscriber Access on page 249

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 31

Configuring Mobile IP Log Files

Configuring the Number and Size of Mobile IP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output. (Mobile IP supports the **files** and **size** options for the **traceoptions** statement.)

  [edit services mobile-ip traceoptions]
  user@host# **set file mip_1 _logfile_1 files 20 size 2097152**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Configuring Access to the Mobile IP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

To specify that all users can read the log file:

- Configure the log file to be world-readable.

  [edit services mobile-ip traceoptions]
  user@host# **set file mip_1 _logfile_1 world-readable**

To explicitly set the default behavior, only the user who configured tracing can read the log file:

- Configure the log file to be no-world-readable.

  [edit services mobile-ip traceoptions]
  user@host# **set file mip_1 _logfile_1 no-world-readable**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Configuring a Regular Expression for Mobile IP Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events.

You can refine the output by including regular expressions that will be matched.

To configure regular expressions to be matched:

- Configure the regular expression.

  [edit services mobile-ip traceoptions]
  user@host# **set file mip_1 _logfile_1 match** ***regex***

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

Copyright © 2016, Juniper Networks, Inc.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit services mobile-ip traceoptions]
  user@host# set level *severity*

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 32

Configuring Mobile IP Trace Flags and Operations

Tracing Mobile IP Operations for Subscriber Access

The Junos OS trace feature tracks Mobile IP operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

Trace-related configurations are independent for each logical system and routing instance in which Mobile IP is configured. Mobile IP can generate two types of log messages:

- Trace messages common to all logical systems and routing instances in which Mobile IP is configured. Examples of this global message type are the messages generated by Mobile IP during initialization after it starts up. These trace messages are stored in the default trace file, **/var/log/mipd**. You cannot configure Mobile IP to save global messages in a different file. Mobile IP traces global messages by default.
- Trace messages specific to a logical system or routing instance in which Mobile IP is configured. An example of this message type is the message generated by Mobile IP when it receives a registration request. These trace messages are stored in the trace file configured for that logical system or routing instance. These messages cannot be saved in **/var/log/mipd**.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file located in the **/var/log** directory. By default, the router uses the filename **mipd** for global tracing. You can specify a different filename, but you cannot change the directory in which trace files are located. Logical system and routing instance messages are logged in a file that you must configure separately from **mipd** in the **/var/log** directory.
2. When the trace log file ***filename*** reaches 128 kilobytes (KB), it is compressed and renamed ***filename*.0.gz**. Subsequent events are logged in a new file called ***filename***,

until it reaches capacity again. At this point, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the System Log Explorer.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure Mobile IP tracing operations:

1. (Optional) Configure a trace log filename.

   See "Configuring the Mobile IP Trace Log Filename" on page 259.

2. (Optional) Configure the number and size of trace logs.

   See "Configuring the Number and Size of Mobile IP Log Files" on page 253.

3. (Optional) Configure user access to trace logs.

   See "Configuring Access to the Mobile IP Log File" on page 254.

4. (Optional) Configure a regular expression to filter the information to be included in the trace log.

   See "Configuring a Regular Expression for Mobile IP Messages to Be Logged" on page 254.

5. (Optional) Configure flags to specify which events are logged.

   See "Configuring the Mobile IP Tracing Flags" on page 259.

6. (Optional) Configure a severity level for messages to specify which event messages are logged.

   See "Configuring the Severity Level to Filter Which Mobile IP Messages Are Logged" on page 254.

Copyright © 2016, Juniper Networks, Inc.

Configuring the Mobile IP Trace Log Filename

Global messages common to all Mobile IP logical systems and routing instances are recorded only in **/var/log/mipd**. Mobile IP automatically creates this file if it is not present when Mobile IP starts. You cannot configure global messages to be recorded in any other file.

You must specify a different name with the **file** option for messages that are specific to a logical system or routing instance in which Mobile IP is configured. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

To configure the filename for Mobile IP tracing operations for a logical system or routing instance:

- Specify the name of the file used for the trace output.

  [edit logical-systems lr1 services mobile-ip traceoptions]
  user@host# **set file mip-lr1_1**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Configuring the Mobile IP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged:

- Configure the flags.

  [edit services mobile-ip traceoptions]
  user@host# **set flag home-agent**

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 33

Contacting Juniper Networks Technical Support

Collecting Subscriber Access Logs Before Contacting Juniper Networks Technical Support

Problem **Description:** When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper Networks Technical Support. This topic shows you the most useful logs for a variety of network implementations. In addition to the relevant log information, you must also collect standard troubleshooting information and send it to Juniper Networks Technical Support in your request for assistance.

Solution To collect standard troubleshooting information:

- Redirect the command output to a file.

```
user@host> request support information | save rsi-1
```

To configure logging to assist Juniper Networks Technical Support:

1. Review the following blocks of statements to determine which apply to your configuration.

   ```
   [edit]
   set system syslog archive size 100m files 25
   set system auto-configuration traceoptions file filename
   set system auto-configuration traceoptions file filename size 100m files 25
   set protocols ppp-service traceoptions file filename size 100m files 25
   set protocols ppp-service traceoptions level all
   set protocols ppp-service traceoptions flag all
   set protocols ppp traceoptions file filename size 100m files 25
   set protocols ppp traceoptions level all
   set protocols ppp traceoptions flag all
   set protocols ppp monitor-session all
   set interfaces pp0 traceoptions flag all
   set demux traceoptions file filename size 100m files 25
   set demux traceoptions level all
   set demux traceoptions flag all
   set system processes dhcp-service traceoptions file filename
   set system processes dhcp-service traceoptions file size 100m
   set system processes dhcp-service traceoptions file files 25
   set system processes dhcp-service traceoptions flag all
   set class-of-service traceoptions file filename
   set class-of-service traceoptions file size 100m
   set class-of-service traceoptions flag all
   set class-of-service traceoptions file files 25
   set routing-options traceoptions file filename
   set routing-options traceoptions file size 100m
   set routing-options traceoptions flag all
   set routing-options traceoptions file files 25
   set interfaces traceoptions file filename
   set interfaces traceoptions file size 100m
   set interfaces traceoptions flag all
   set interfaces traceoptions file files 25
   set system processes general-authentication-service traceoptions file filename
   set system processes general-authentication-service traceoptions file size 100m
   set system processes general-authentication-service traceoptions flag all
   set system processes general-authentication-service traceoptions file files 25
   ```

2. Copy the relevant statements into a text file and modify the log filenames as you want.

3. Copy the statements from the text file and paste them into the CLI on your router to configure logging.

4. Commit the logging configuration to begin collecting information.

NOTE: The maximum file size for DHCP local server and DHCP relay log files is 1 GB. The maximum number of log files for DHCP local server and DHCP relay is 1000.

Copyright © 2016, Juniper Networks, Inc.

BEST PRACTICE: Enable these logs only to collect information when troubleshooting specific problems. Enabling these logs during normal operations can result in reduced system performance.

Related Documentation

- *Compressing Troubleshooting Logs from /var/logs to Send to Juniper Networks Technical Support*

Copyright © 2016, Juniper Networks, Inc.

PART 6

Configuration Statements and Operational Commands

Copyright © 2016, Juniper Networks, Inc.

CHAPTER 34

Configuration Statements

Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

 Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

- weighted-load-balancing (L2TP LAC) on page 447
- wimax on page 447
- virtual-network on page 448
- vlan-id (Dynamic Profiles) on page 449
- vlan-tags on page 450

[edit access tunnel-profile] Hierarchy Level

```
access {
  tunnel-profile profile-name {
    tunnel tunnel-id {
      identification name;
      logical-system logical-system-name;
      max-sessions number;
      medium type;
      nas-port-method cisco-avp;
      preference number;
      remote-gateway {
        address server-ip-address;
        gateway-name server-name;
      }
      routing-instance routing-instance-name;
      secret password;
      source-gateway {
        address client-ip-address;
        gateway-name client-name;
      }
      type tunnel-type;
    }
  }
}
```

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

[edit access tunnel-switch-profile] Hierarchy Level

```
access {
  tunnel-switch-profile profile-name {
    avp {
      bearer-type action;
      calling-number action;
      cisco-nas-port-info action;
    }
    tunnel-profile profile-name;
  }
}
```

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

Copyright © 2016, Juniper Networks, Inc.

[edit dynamic-profiles] Hierarchy Level

```
dynamic-profiles {
  profile-name {
    class-of-service {
      interfaces {
        interface-name {
          unit logical-unit-number {
            classifiers {
              type (classifier-name | default);
            }
            output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
            rewrite-rules {
              dscp (rewrite-name | default);
              dscp-ipv6 (rewrite-name | default);
              ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
              inet-precedence (rewrite-name | default);
              }
            }
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        (scheduler-name) {
          buffer-size (percent percentage | remainder | temporal microseconds |
            $junos-cos-scheduler-bs);
          drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
            protocol (any | non-tcp | tcp) drop-profile (profile-name | predefined-variable);
          excess-priority (low | high | $junos-cos-scheduler-excess-priority);
          excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
          overhead-accounting (shaping-mode) <bytes (byte-value>;
          priority (priority-level | $junos-cos-scheduler-priority);
          shaping-rate (rate | predefined-variable);
          transmit-rate (rate | percent percentage | remainder | percent percentage
            $junos-cos-scheduler-tx) <exact | rate-limit>;
        }
      }
      traffic-control-profiles profile-name {
        delay-buffer-rate (percent percentage | rate);
        excess-rate (percent percentage | proportion value | percent
          $junos-cos-excess-rate);
        guaranteed-rate (percent percentage | rate);
        overhead-accounting (shaping-mode) <bytes (byte-value>;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate | predefined-variable);
      }
    }
    firewall {
      family family {
```

```
    fast-update-filter filter-name {
      interface-specific;
    match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      only-at-create;
    filter filter-name {
    interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
  policer policer-name {
    filter-specific;
    if-exceeding {
      (bandwidth-limit bps | bandwidth-percent percentage);
      burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
      policer-action;
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit-limit bps;
        burst-size-limit bytes;
      }
      then {
        policer-action;
      }
    }
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        policer-action;
      }
    }
  }
```


Copyright © 2016, Juniper Networks, Inc.

```
    three-color-policer policer-name {
      action {
        loss-priority high then discard;
      }
      logical-interface-policer;
      single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
      }
      two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
        }
      }
    }
  }
  policy-options {
    prefix-listname {
      ip-addresses;
      dynamic-db;
    }
  }
  interfaces {
    interface-name {
      unit logical-unit-number {
        family family {
          access-concentrator name;
          address address;
          direct-connect;
          duplicate-protection;
          dynamic-profile profile-name;
          filter {
            adf {
              counter;
              input-precedence precedence;
              not-mandatory;
              output-precedence precedence;
              rule rule-value;
            }
            input filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
            output filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
          }
          max-sessions number;
          max-sessions-vsa-ignore;
          rpf-check {
```

```
                fail-filter filter-name;
                mode loose;
            }
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                output {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
            }
            service-name-table table-name;
            short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                maximum-seconds>;
            unnumbered-address interface-name <preferred-source-address address>;
        }
        ppp-options {
            chap;
            pap;
        }
        vlan-id number;
    }
    vlan-tagging;
}
interface-set interface-set-name {
    interface interface-name {
        unit logical-unit-number;
    }
}
demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        demux-source {
            source-prefix;
        }
        family family {
            access-concentrator name;
            address address;
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
                input filter-name;
                output filter-name;
            }
            mac-validate (loose | strict):
            max-sessions number;
            max-sessions-vsa-ignore;
            service-name-table table-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                  maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
              }
            }
          }
          pp0 {
            unit logical-unit-number {
              keepalives interval seconds;
              no-keepalives;
              pppoe-options {
                underlying-interface interface-name;
                server;
              }
              ppp-options {
                authentication [ authentication-protocols ];
                chap {
                  challenge-length minimum minimum-length maximum maximum-length;
                }
                pap;
              }
              family inet {
                unnumbered-address interface-name;
                address address;
                service {
                  input {
                    service-set service-set-name {
                      service-filter filter-name;
                    }
                    post-service-filter filter-name;
                  }
                  output {
                    service-set service-set-name {
                      service-filter filter-name;
                    }
                  }
                }
                filter {
                  input filter-name {
                    precedence precedence;
                  }
                  output filter-name {
                    precedence precedence;
                  }
                }
              }
            }
          }
        }
        protocols {
          igmp {
            interface interface-name {
              accounting;
              disable;
              group-policy;
              immediate-leave
```

```
                no-accounting;
                promiscuous-mode;
                ssm-map ssm-map-name;
                static {
                    group group {
                        source source;
                    }
                }
                version version;
            }
            mld {
                interface interface-name {
                    disable;
                    (accounting | no-accounting);
                    group-policy;
                    immediate-leave;
                    oif-map;
                    passive;
                    ssm-map ssm-map-name;
                    static {
                        group multicast-group-address {
                            exclude;
                            group-count number;
                            group-increment increment;
                            source ip-address {
                                source-count number;
                                source-increment increment;
                            }
                        }
                    }
                    version version;
                }
            }
            router-advertisement {
                interface interface-name {
                    current-hop-limit number;
                    default-lifetime seconds;
                    (managed-configuration | no-managed-configuration);
                    max-advertisement-interval seconds;
                    min-advertisement-interval seconds;
                    (other-stateful-configuration | no-other-stateful-configuration);
                    prefix prefix {
                        (autonomous | no-autonomous);
                        (on-link | no-on-link);
                        preferred-lifetime seconds;
                        valid-lifetime seconds;
                    }
                    reachable-time milliseconds;
                    retransmit-timer milliseconds;
                }
            }
        }
    }
}
routing-instances routing-instance-name {
    interface interface-name;
```


Copyright © 2016, Juniper Networks, Inc.

```
    routing-options {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
      multicast {
        interface interface-name {
          no-qos-adjust;
        }
      }
    }
    rib routing-table-name {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
    }
  }
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
```

```
        multicast {
          interface interface-name {
            no-qos-adjust;
          }
        }
      }
      variables {
        variable-name {
          default-value default-value;
          equals expression;
          mandatory;
          uid;
          uid-reference;
        }
      }
    }
```

Related Documentation

- *Dynamic Profiles Overview*
- *CoS for Subscriber Access Overview*
- *Configuring a Basic Dynamic Profile*
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Two-Color Policer Configuration Overview*
- *Three-Color Policer Configuration Overview*
- *Hierarchical Policer Configuration Overview*
- *Guidelines for Applying Traffic Policers*

[edit protocols ppp-service] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  ppp-service {
    on-demand-ip-address;
    reject-unauthorized-ipv6cp;
    traceoptions {
      file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
      filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
        user user@domain;
      }
      flag flag;
      level severity;
      no-remote-trace;
```


Copyright © 2016, Juniper Networks, Inc.

```
        }
      }
    }
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit services l2tp] Hierarchy Level

NOTE: The **tunnel-group** *group-name* stanza is not supported for L2TP LAC. It applies only to L2TP LNS. Similarly, some of the options for the **traceoptions** statement apply only to L2TP LNS; for more information, see traceoptions.

```
services {
  l2tp {
    destination [ip-address] {
      access-line-information <connection-speed-update>;
      lockout-timeout seconds;
    }
    destruct-timeout seconds;
    disable-calling-number-avp;
    disable-failover-protocol;
    enable-snmp-tunnel-statistics;
    failover-within-preference;
    ip-reassembly;
    rx-connect-speed-when-equal;
    traceoptions {
      debug-level level;
      file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
      filter {
        protocol name;
        user user@domain;
        user-name username;
      }
      flag flag;
      interfaces interface-name {
        debug-level severity;
        flag flag;
      }
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
    tunnel {
      assignment-id-format (assignment-id |client-server-id);
      idle-timeout seconds;
      minimum-retransmission-timeout;
      nas-port-method;
      retransmission-count-established count;
      retransmission-count-not-established count;
      rx-window-size
```

```
        tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
      }
      tunnel-group group-name {
        aaa-access-profile profile-name;
        dynamic-profile;
        hello-interval seconds;
        hide-avps;
        l2tp-access-profile profile-name;
        local-gateway address {
          address address;
          gateway-name gateway-name;
        }
        maximum-send-window packets;
        ppp-access-profile profile-name;
        receive-window packets;
        retransmit-interval seconds;
        service-device-pool;
        service-interface interface-name;
        syslog {
          host hostname {
            facility-override facility-name;
            log-prefix prefix-value;
            services severity-level;
          }
        }
        tos-reflect;
        tunnel-switch-profile profile-name;
        tunnel-timeout seconds;
      }
      tunnel-switch-profile profile-name;
      tx-connect-speed-method method;
      weighted-load-balancing;
    }
  }
```

Related Documentation

- L2TP for Subscriber Access Overview on page 123
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

[edit services mobile-ip] Hierarchy Level

```
services {
  mobile-ip {
    access-type {
      (generic | wimax);
    }
    authenticate {
      order (aaa | local);
    }
    dynamic-home-assignment {
      home-agent {
        nai (name@domain | @domain) {
          home-agent ip-address;
```


Copyright © 2016, Juniper Networks, Inc.

```
            }
          }
        }
        home-agent {
          enable-service interface-name;
          virtual-network {
            home-agent-address ip-address {
              registration-lifetime seconds;
              revocation-required;
              timestamp-tolerance seconds;
            }
          }
        }
        peer {
          (ip-address address | nai name@domain) {
            spi hexadecimal-value {
              algorithm (hmac-md5 | md5);
              entity-type (host | mobility-agent);
              key (hex | ascii) string;
              replay-method (none | timestamp seconds);
            }
          }
        }
        traceoptions {
          file filename <files number> <match regular-expression > <size maximum-file-size>
            <world-readable | no-world-readable>;
          flag flag;
          level <all | error | info | notice | verbose | warning>;
          no-remote-trace;
        }
      }
    }
```

Related Documentation

- Mobile IP Home Agent Elements and Behavior on page 217
- Configuring Mobile IP on page 229

aaa-access-profile (L2TP LNS)

Syntax aaa-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced in Junos OS Release 11.4.
Support at the **[edit access profile *profile-name* client *client-name* l2tp]** hierarchy level introduced in Junos OS Release 12.1.

Description Specify a AAA access profile that overrides the AAA access profile configured for the routing instance with the **access-profile** statement. You can configure a profile to specify the RADIUS server settings for a tunnel group or for a LAC client, or both. The AAA access profile configured for the client takes precedence over the AAA access profile configured for the tunnel group, which takes precedence over the access profile configured for the routing instance.

Options ***profile-name***—Name of the local access profile for the tunnel group or client.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

 Copyright © 2016, Juniper Networks, Inc.

access (Dynamic Access Routes)

Syntax

```
access {
  route prefix {
    next-hop next-hop;
    metric route-cost;
    preference route-distance;
    tag route-tag;
  }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*],
[edit dynamic-profiles *profile-name* routing-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure access routes.

BEST PRACTICE: We recommend that you always include the **access-internal** stanza in the dynamic-profile when the **access** stanza is present for framed-route support.

Options

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45

access-internal (Dynamic Access-Internal Routes)

Syntax

```
access-internal {
  route subscriber-ip-address {
    qualified-next-hop underlying-interface {
      mac-address address;
    }
  }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*],
[edit dynamic-profiles routing-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure access-internal routes. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.

BEST PRACTICE: We recommend that you always include the access-internal stanza in the dynamic-profile when the access stanza is present for framed-route support.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101

Copyright © 2016, Juniper Networks, Inc.

access-line-information (L2TP LAC)

Syntax access-line-information <connection-speed-update>;

Hierarchy Level [edit services l2tp tunnel destination *ip-address*]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the LAC to forward subscriber line identification and other DSL attributes to the LNS by means of L2TP AVPs for tunnels with the specified endpoint address.

Options **connection-speed-update**—(Optional) Include the Connect Speed Update Enable AVP (98) in ICRQ messages to alert the LNS that the LAC may send CSUN messages that report speed changes originating with the ANCP agent.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the LAC to Report Access Line Information to the LNS on page 154

access-type

Syntax
```
access-type {
   (generic | wimax);
}
```

Hierarchy Level [edit services mobile-ip],
[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip],
[edit routing-instances *routing-instance-name* services mobile-ip]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure the access type for Mobile IP.

The remaining statements are explained separately.

Default The generic access type is used by default.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 229
- Configuring the Access Type for Mobile IP on page 232

address (LNS Local Gateway)

Syntax address *address*;

Hierarchy Level [edit services l2tp tunnel-group *name* local-gateway]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the local (LNS) IP address for L2TP tunnel.

Options ***address***—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Local Gateway Address and PIC*.
- *Configuring L2TP Tunnel Groups*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

address (Tunnel Profile Remote Gateway)

Syntax address *server-ip-address*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* remote-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address of the remote gateway device at the L2TP tunnel endpoint, the LNS.

Options ***server-ip-address***—IP address of the remote gateway device.

Default: 0.0.0.0.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

 Copyright © 2016, Juniper Networks, Inc.

address (Tunnel Profile Source Gateway)

Syntax address *client-ip-address*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address of the source gateway device at the local L2TP tunnel endpoint, the LAC. This value overrides the default address for the logical system or routing instance.

Options ***client-ip-address***—IP address of the source gateway device.
Default: 0.0.0.0.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 162

address-change-immediate-update

Syntax address-change-immediate-update;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the router to send an Address-Change-Update message to the RADIUS accounting server. Any change to this setting takes effect for all new subscriber logins. Existing subscribers are not impacted by this change except when the AAA daemon restarts.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

algorithm

Syntax algorithm (hmac-md5 | md5);

Hierarchy Level
[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer nai *user@domain* spi *hexadecimal-value*]

Release Information
Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the algorithm used for authenticating Mobile IP messages.

Default hmac-md5

Options
hmac-md5—Specifies algorithm hmac-md5

md5—Specifies algorithm md5

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

 Copyright © 2016, Juniper Networks, Inc.

allow-snooped-clients

Syntax allow-snooped-clients;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information
Statement introduced in Junos OS Release 10.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 12.1.

Description
Explicitly enable DHCP snooping support on the router.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.

NOTE: DHCP snooping is *disabled* by default.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Extended DHCP Relay Agent Overview*
- *Overriding the Default DHCP Relay Configuration Settings*
- DHCP Snooping Support on page 51

always-write-option-82

Syntax

always-write-option-82;

Hierarchy Level

[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.
- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Relay Agent Overview*

Copyright © 2016, Juniper Networks, Inc.

assignment-id-format (L2TP LAC)

Syntax assignment-id-format (assignment-id | client-server-id);

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 11.4.

Description Set the format for the name used for a tunnel, the tunnel assignment ID.

NOTE: Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing **no services l2tp tunnel assignment-id-format**.

Default assignment-Id

Options

assignment-Id—The tunnel name corresponds to RADIUS attribute Tunnel-Assignment-Id [82].

client-server-id—The tunnel name is a combination of RADIUS attributes Tunnel-Client-Auth-Id [90], Tunnel-Server-Auth-Id [91], and Tunnel-Assignment-Id [82].

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the Format for the Tunnel Name on page 162

authenticate

Syntax

```
authenticate {
  order (aaa | local);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication method performed for Mobile IP.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Access Type for Mobile IP on page 232

Copyright © 2016, Juniper Networks, Inc.

authentication (Static and Dynamic PPP)

Syntax authentication [*authentication-protocols*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options], [edit interfaces pp0 unit *unit-number* ppp-options]

Release Information Statement introduced in Junos OS Release 12.2.

Description Specify the order in which the router tries to negotiate PPP authentication protocols when verifying that a PPP client can access the network. By default, the router tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication first, and then tries Password Authentication Protocol (PAP) authentication if the attempt to negotiate CHAP authentication is unsuccessful.

You can specify one or both authentication protocols. If you specify both CHAP and PAP in either order, you must enclose the set of protocol names within square brackets ([ ]).

Options ***authentication-protocols***—One or both of the following PPP authentication protocols:

- **chap**—Challenge Handshake Authentication Protocol
- **pap**—Password Authentication Protocol

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Controlling the Negotiation Order of PPP Authentication Protocols on page 110

avp (L2TP Tunnel Switching)

Syntax

```
avp {
  bearer-type;
  calling-number;
  cisco-nas-port-info;
}
```

Hierarchy Level [edit access tunnel-switch-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on L2TP AVPs that are negotiated when the first session is created; these AVPs are contained in the L2TP packets that are switched by the tunnel switch profile.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

bandwidth (Inline Services)

Syntax bandwidth (1g | 10g);

Hierarchy Level [edit chassis fpc *slot-number* pic *number* inline-services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services.

Options

1g—Reserves 1 Gbps of bandwidth for tunnel traffic. Configuring a bandwidth of 1 Gbps creates a virtual tunnel interface that is represented as **vt-<*slot-number/pic-number/*10>**.

10g—Reserves 10 Gbps of bandwidth for tunnel traffic. Configuring a bandwidth of 10 Gbps creates a virtual tunnel interface that is represented as **vt-<*slot-number/pic-number/*0>** .

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Inline Service Interfaces on page 184
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

bearer-type (L2TP Tunnel Switching)

Syntax bearer-type *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on the Bearer Type AVP (18) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

bfd

Syntax

```
bfd {
  version (0 | 1 | automatic);
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  detection-time {
    threshold milliseconds;
  }
  session-mode (automatic | multihop | singlehop);
  holddown-interval milliseconds;
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit system services dhcp-local-server group *group-name* liveness-detection method],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

Copyright © 2016, Juniper Networks, Inc.

calling-number (L2TP Tunnel Switching)

Syntax calling-number *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the action taken on the Calling Number AVP (22) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

challenge-length (Static and Dynamic PPP)

Syntax challenge-length minimum *minimum-length* maximum *maximum-length*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options chap],
[edit interfaces pp0 unit *unit-number* ppp-options chap]

Release Information Statement introduced in Junos OS Release 12.2.

Description Modify the length of the Challenge Handshake Authentication Protocol (CHAP) challenge by specifying the minimum and maximum allowable length, in bytes.

BEST PRACTICE: **We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.**

Options ***minimum-length***—Minimum length, in bytes, of the CHAP challenge.

Range: 8 through 63

Default: 16

maximum-length—Maximum length, in bytes, of the CHAP challenge. The *maximum-length* must be equal to or greater than the *minimum-length*.

Range: 8 through 63

Default: 32

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Modifying the CHAP Challenge Length on page 104

 Copyright © 2016, Juniper Networks, Inc.

chap

Syntax

```
chap {
   access-profile name;
   challenge-length minimum minimum-length maximum maximum-length;
   default-chap-secret name;
   local-name name;
   passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* ppp-options],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Allow each side of a link to challenge its peer, using a "secret" known only to the authenticator and that peer. The secret is not sent over the link.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: **On inline service (si) interfaces for L2TP, only the chap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.**

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

chap (Dynamic PPP)

Syntax

```
chap {
   challenge-length minimum minimum-length maximum maximum-length;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]** hierarchy level introduced in Junos OS Release 12.2.

Description

Specify CHAP authentication in a PPP dynamic profile.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

chap (L2TP)

Syntax

```
chap;
```

Hierarchy Level

[edit access group-profile *profile-name* ppp ppp-options]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

(MX Series routers only) Specify CHAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level

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

Related Documentation

- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

Copyright © 2016, Juniper Networks, Inc.

cisco-nas-port-info (L2TP Tunnel Switching)

Syntax cisco-nas-port-info *action*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name* avp]

Release Information Statement introduced in Junos OS Release 13.2.

Description Define a tunnel profile for subscriber access.

Specify the action taken on the Cisco NAS Port Info AVP (100) in the L2TP packets during tunnel switching if the AVP is negotiated when the first session is created.

Options ***action***—One of the following actions:

- **drop**—Drop the AVP.
- **regenerate**—Regenerate the AVP based on the local policy at the LTS and send it in the switched packet. The local policy may or may not use the value for the AVP received during negotiation for the first session.
- **relay**—Forward the AVP transparently as is and send it in the switched packet.

Default: relay

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

client

Syntax

```
client client-name {
  chap-secret chap-secret;
  group-profile profile-name;
  ike {
    allowed-proxy-pair {
      remote remote-proxy-address local local-proxy-address;
    }
    pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
    ike-policy policy-name;
    interface-id string-value;
  }
  l2tp {
    aaa-access-profile profile-name;
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
    multilink {
      drop-timeout milliseconds;
      fragment-threshold bytes;
    }
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    shared-secret shared-secret;
  }
  pap-password pap-password;
  ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-ip-address ip-address;
    framed-pool framed-pool;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
  user-group-profile profile-name;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the peer identity.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Copyright © 2016, Juniper Networks, Inc.

Options ***client-name***—A peer identity. For L2TP clients, you can use a special name to configure a default client. This client enables the LNS to accept any LAC to establish the session. On M Series routers, use * for the default client configuration. On MX Series routers, use **default**.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the L2TP Client*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- Configuring an L2TP Access Profile on the LNS on page 180

destination (L2TP)

Syntax

```
destination [ip-address] {
  access-line-information <connection-speed-update>;
  lockout-timeout seconds;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.2.
Option *ip-address* introduced in Junos OS Release 14.1.

Description Configure attributes for all L2TP destinations or a specified L2TP destination.

Options ***ip-address***—(Optional) Endpoint address of an L2TP tunnel. This option applies only to the **access-line-information** statement.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the L2TP Destination Lockout Timeout on page 139
- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- Configuring the LAC to Report Access Line Information to the LNS on page 154

destruct-timeout (L2TP)

Syntax destruct-timeout *seconds*;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp destruct-timeout.**

Options *seconds*—Length of the destruct timeout.

Range: 10 through 3600

Default: 300

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the L2TP Destruct Timeout on page 138
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

detection-time

Syntax

```
detection-time {
  threshold milliseconds;
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

dhcp-relay

Syntax

```
dhcp-relay {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    active-server-group server-group-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
      use-primary primary-profile-name;
    }
    group group-name {
      active-server-group server-group-name;
      authentication {
        ...
      }
      dynamic-profile profile-name {
        ...
      }
      interface interface-name {
        exclude;
        liveness-detection {
          failure-action (clear-binding | clear-binding-if-interface-up | log-only);
          method {
```


Copyright © 2016, Juniper Networks, Inc.

```
                bfd {
                  version (0 | 1 | automatic);
                  minimum-interval milliseconds;
                  minimum-receive-interval milliseconds;
                  multiplier number;
                  no-adaptation;
                  transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                  }
                  detection-time {
                    threshold milliseconds;
                  }
                  session-mode(automatic | multihop | singlehop);
                  holddown-interval milliseconds;
                }
              }
            }
            overrides {
              ...
            }
            relay-option {
              ...
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
          }
          route-suppression:
          service-profile dynamic-profile-name;
          overrides {
            ...
          }
          relay-agent-interface-id {
            ...
          }
          relay-agent-remote-id {
            ...
          }
          relay-option {
            ...
          }
          route-suppression;
          server-response-time seconds;
          service-profile dynamic-profile-name;
        }
        liveness-detection {
          failure-action (clear-binding | clear-binding-if-interface-up | log-only);
          method {
            bfd {
              version (0 | 1 | automatic);
              minimum-interval milliseconds;
              minimum-receive-interval milliseconds;
              multiplier number;
              no-adaptation;
              transmit-interval {
```

```
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    delay-authentication;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
```

 Copyright © 2016, Juniper Networks, Inc.

```
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                  minimum-interval milliseconds;
                  threshold milliseconds;
                }
                detection-time {
                  threshold milliseconds;
                }
                session-mode(automatic | multihop | singlehop);
                holddown-interval milliseconds;
              }
            }
          }
          overrides {
            ...
          }
          service-profile dynamic-profile-name;
          trace;
          upto upto-interface-name;
        }
        overrides {
          ...
        }
        relay-option {
          ...
        }
        relay-option-82 {
          ...
        }
        route-suppression:
        service-profile dynamic-profile-name;
      }
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
```

```
        allow-snooped-clients;
        always-write-giaddr;
        always-write-option-82;
        client-discover-match (option60-and-option82 | incoming-interface);
        delay-authentication;
        disable-relay;
        interface-client-limit number;
        layer2-unicast-replies;
        no-allow-snooped-clients;
        no-bind-on-request;
        proxy-mode;
        replace-ip-source-with;
        send-release-on-delete;
        trust-option-82;
    }
    relay-option {
        option-number option-number;
        default-action {
            drop;
            forward-only;
            relay-server-group group-name;
        }
        equals (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            relay-server-group relay-server-group;
        }
        starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
            drop;
            forward-only;
            local-server-group local-server-group;
            relay-server-group relay-server-group;
        }
    }
    relay-option-82 {
        circuit-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
        remote-id {
            prefix prefix;
            use-interface-description (logical | device);
        }
    }
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    route-suppression:
    server-response-time seconds;
    service-profile dynamic-profile-name;
}
```


Copyright © 2016, Juniper Networks, Inc.

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or 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 client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

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

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Relay Agent Overview*
- *DHCPv6 Relay Agent Overview*
- *DHCP Relay Proxy Overview*
- *Using External AAA Authentication Services with DHCP*

dial-options

Syntax

```
dial-options {
   ipsec-interface-id name;
   l2tp-interface-id name;
   (shared | dedicated);
}
```

Hierarchy Level

[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces si-*fpc/pic/port* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
The **[edit ...si-...]** hierarchy levels introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id ***name***—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile** ***name*** **client * ike]** hierarchy level.

l2tp-interface-id ***name***—Interface identifier that must be replicated at the **[edit access profile** ***name*****]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

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

Related Documentation

- *Configuring the Identifier for Logical Interfaces that Provide L2TP Services*
- *Configuring Dynamic Endpoints for IPsec Tunnels*
- Configuring Options for the LNS Inline Services Logical Interface on page 185

Copyright © 2016, Juniper Networks, Inc.

dial-options (Dynamic Profiles)

Syntax

```
dial-options {
   ipsec-interface-id name;
   l2tp-interface-id name;
   (shared | dedicated);
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the options for configuring logical interfaces in dynamic profiles for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level. This options is not currently supported.

l2tp-interface-id *name*—(MX Series routers only) L2TP interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Dynamic Profile for Dynamic LNS Sessions on page 200

disable-calling-number-avp (L2TP LAC)

Syntax disable-calling-number-avp;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Prevent the LAC from sending L2TP Calling Number AVP 22 in incoming-call request (ICRQ) packets to the LNS. By default, the LAC in an L2TP network generates this AVP from the Calling-Station-Id and sends it to the LNS.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Preventing the LAC from Sending Calling Number AVP 22 to the LNS on page 172

disable-failover-protocol (L2TP LAC)

Syntax disable-failover-protocol;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure the LAC to use only the silent failover method when resynchronizing with the peer LNS in the event of LAC failover. This command prevents the default behavior, wherein the LAC first attempts to use the failover protocol and then falls back on the silent failover method. This configuration can be useful when routers that act as the LNS are configured for silent failover or incorrectly negotiate use of the failover protocol even though they do not support it.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Preventing the LAC from Negotiating L2TP Failover Protocol on page 150

 Copyright © 2016, Juniper Networks, Inc.

duplicate-clients-in-subnet (DHCP Local Server and DHCP Relay Agent)

Syntax

duplicate-clients-in-subnet (incoming-interface | option-82);

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Configure how the router distinguishes between duplicate clients in the same subnet. Duplicate clients are defined as clients that have the same hardware address or client ID.

NOTE: You must configure the **duplicate-clients-in-subnet** statement identically for both **[forwarding-options dhcp-relay]** and **[system services dhcp-local-server]** hierarchy levels.

Options

incoming-interface—Use the incoming interface information in packets to differentiate between duplicate clients.

option-82—Use the option 82 information to differentiate between duplicate clients.

Required Privilege Level

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

Related Documentation

- DHCP Duplicate Client In Subnet Overview on page 69
- Guidelines for Configuring Support for DHCP Duplicate Clients on page 70
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Their Incoming Interfaces on page 72
- Configuring the Router to Distinguish Between DHCP Duplicate Clients Based on Option 82 Information on page 70

Copyright © 2016, Juniper Networks, Inc.

dynamic-home-assignment

Syntax

```
dynamic-home-assignment {
  home-agent {
    nai (name@domain.com | @domain.com) {
      home-agent ip-address;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the dynamic assignment rule for the home agent.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring Dynamic Home Assignment for the Mobile Node on page 232

dynamic-profile (L2TP)

Syntax

dynamic-profile *profile-name*;

Hierarchy Level

[edit services l2tp tunnel-group *name*]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Assign a dynamic profile to the tunnel group for dynamic LNS sessions.

Options

profile-name—Name of the dynamic profile for the tunnel group.

Required Privilege Level

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

Related Documentation

- Configuring a Dynamic Profile for Dynamic LNS Sessions on page 200

Copyright © 2016, Juniper Networks, Inc.

dynamic-profile (PPP)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information Statement introduced in Junos OS Release 9.5.
Support for MLPPP on LSQ interfaces introduced in Junos OS Release 10.2.

Description Specify the dynamic profile that is attached to the interface. On the MX Series routers, this statement is currently supported on PPPoE interfaces only. On the M120 and M320 routers, this statement is supported for MLPPP bundles only on LSQ interfaces on Adaptive Services PICs and Multiservices PICs.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- *Attaching Dynamic Profiles to MLPPP Bundles*
- For hardware requirements, see *Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces*

enable-service

Syntax enable-service *interface-name*;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit services mobile-ip home-agent]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Define the list of interfaces on which the home agent service can be enabled. The system accepts registration requests only if it is on one of these interfaces. Include the statement once for each interface to be enabled.

Options ***interface-name***—Interface on which the home agent can be enabled.

Required Privilege Level view—To view this statement in the configuration.
view-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

 Copyright © 2016, Juniper Networks, Inc.

enable-snmp-tunnel-statistics (L2TP)

Syntax enable-snmp-tunnel-statistics;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 12.1R4 and supported in later 12.1Rx releases.
Statement supported in Junos OS Release 12.2R2 and later 12.2Rx releases. (Not supported in Junos OS Release 12.2R1.)
Statement supported in Junos OS Release 12.3 and later releases.

Description Enable collection of L2TP tunnel and global counters for SNMP statistics.

NOTE: **The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.**

Default Disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Enabling Tunnel and Global Counters for SNMP Statistics Collection on page 214

enforce-strict-scale-limit-license (Subscriber Management)

Syntax enforce-strict-scale-limit-license;

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure the router to strictly enforce the subscriber scaling license, and to not allow the normal grace period. No additional subscribers are allowed to log in after the number of subscribers reaches the maximum allowed for the license.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- *Configuring the Router to Strictly Enforce the Subscriber Scaling License*

entity-type

Syntax entity-type (host | mobility-agent);

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip peer spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer spi *hexadecimal-value*],
[edit services mobile-ip peer spi *hexadecimal-value*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the security parameter for the peer entity—, either a mobile node, home agent, or foreign agent.

Options **host**—Use the mobile node in home agent

mobility-agent—Use the home agent or foreign agent

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

equals (Dynamic Profile)

Syntax equals *expression*;

Hierarchy Level [edit dynamic-profiles *profile-name* variables *variable-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure an expression for a user-defined variable that is evaluated at run time and returned as the variable value.

Options ***expression***—Expression evaluated to return a value for the user-defined variable.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *User-Defined Variables*
- *Using Variable Expressions in User-Defined Variables*
- *Configuring User-Defined Dynamic Variables in Dynamic Profiles*

 Copyright © 2016, Juniper Networks, Inc.

failover-within-preference (L2TP LAC)

Syntax failover-within-preference;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Enable L2TP LAC tunnel selection within a preference level. When the router is unable to connect to a destination at a given preference level, it attempts to connect to another destination at the same level. By default, when a connection attempt fails at one preference level, the next attempt is made at the next lower level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring LAC Tunnel Selection Failover Within a Preference Level on page 165
- Configuring the L2TP LAC Tunnel Selection Parameters on page 165

failure-action

Syntax failure-action (clear-binding | clear-binding-if-interface-up | log-only);

Hierarchy Level [edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group *group-name* liveness-detection],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the action the router (or switch) takes when a liveness detection failure occurs.

Options

clear-binding—The client session is cleared when a liveness detection failure occurs.

clear-binding-if-interface-up—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.

log-only—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- DHCP Liveness Detection Overview on page 75
- Configuring Detection of DHCP Local Server Client Connectivity on page 81
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 77
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

 Copyright © 2016, Juniper Networks, Inc.

forward-snooped-clients (DHCP Local Server)

Syntax

forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure how the DHCP local server handles DHCP snooped packets on specific interfaces.

Options

all-interfaces—Perform the action on all interfaces.

configured-interfaces—Perform the action only on configured interfaces.

non-configured-interfaces—Perform the action only on nonconfigured interfaces.

Required Privilege Level

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

Related Documentation

- DHCP Snooping Support on page 51
- Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 52

forward-snooped-clients (DHCP Relay Agent)

Syntax

forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]

Release Information

Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure how DHCP relay agent handles DHCP snooped packets on specific interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the **forward-snooped-clients** configuration and the configuration of either the **allow-snooped-clients** statement or the **no-allow-snooped-clients** statement.

The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on nonconfigured interfaces.

Options

all-interfaces—Perform the action on all interfaces.

configured-interfaces—Perform the action only on configured interfaces.

non-configured-interfaces—Perform the action only on nonconfigured interfaces.

Required Privilege Level

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

Related Documentation

- DHCP Snooping Support on page 51
- *Configuring DHCP Snooping for DHCP Relay Agent*

Copyright © 2016, Juniper Networks, Inc.

fpc (MX Series 3D Universal Edge Routers)

Syntax

```
fpc slot-number {
  inline-services {
    flow-table-size {
      ipv4-flow-table-size units;
      ipv4-flow-table-size  units;
    }
  }
  pic number {
    inline-services {
      bandwidth (1g | 10g);
    }
    port-mirror-instance port-mirroring-instance-name-pic-level;
    tunnel-services {
      bandwidth (1g | 10g)
    }
  }
  port-mirror-instance port-mirroring-instance-name-fpc-level;
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 8.2.
Option **port-mirror-instance** introduced in Junos OS Release 9.3.

Description

Configure properties for the DPC or MPC and corresponding Packet Forwarding Engines to create tunnel interfaces.

(MX Series Virtual Chassis only) To configure properties for MPCs in a member router in an MX Series Virtual Chassis configuration, you must specify the router's Virtual Chassis member number *before* the **fpc** statement. Specify the member number in the form **member** ***member-id***, where ***member-id*** is 0 or 1. If you do not specify the member number before the **fpc** statement, the commit operation fails and the software displays an error message indicating that the **fpc** statement must include the member number for routers in Virtual Chassis mode.

Options

fpc ***slot-number***—Specify the slot number of the DPC.

Range: 0 through 11

pic ***number***—Specify the number of the Packet Forwarding Engine. Each DPC includes four Packet Forwarding Engines.

Range: 0 through 4

port-mirror-instance ***port-mirroring-instance-name-fpc-level***—Associate a port-mirroring instance with the DPC and its corresponding PICs. The port-mirroring instance is configured under the **[edit forwarding-options port-mirroring]** hierarchy level.

The remaining statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Port-Mirroring Instances on MX Series 3D Universal Edge Routers*
- Enabling Inline Service Interfaces on page 184

gateway-name (LNS Local Gateway)

Syntax
gateway-name *gateway-name*;

Hierarchy Level
[edit services l2tp tunnel-group *group-name* local-gateway]

Release Information
Statement introduced in Junos OS Release 12.2.

Description
Specify the gateway name for the LNS, which the LNS returns to the LAC in response to the LAC's SCCRQ message. This name must match the remote gateway name configured on the LAC, or the tunnel cannot be established.

Options
gateway-name—Name of the LNS.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

gateway-name (Tunnel Profile Remote Gateway)

Syntax
gateway-name *server-name*;

Hierarchy Level
[edit access tunnel-profile *profile-name* tunnel *tunnel-id* remote-gateway]

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Specify the hostname expected by the remote gateway—the LNS—from the source gateway—the LAC—when you set up a tunnel.

Options
server-name—Name of the LNS.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 162

Copyright © 2016, Juniper Networks, Inc.

gateway-name (Tunnel Profile Source Gateway)

Syntax gateway-name *client-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id* source-gateway]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the hostname provided by the source gateway—the LAC—to the remote gateway—the LNS—when you set up a tunnel.

Options ***client-name***—Name of the LAC.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

generic

Syntax generic;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip access-type],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip access-type],
[edit routing-instances *routing-instance-name* services mobile-ip access-type],
[edit services mobile-ip access-type]

Release Information Statement introduced in Junos OS Release 9.5.

Description Disable WiMAX features for Mobile IP home agent, preventing interoperation in a WiMAX environment.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Access Type for Mobile IP on page 232

gres-route-flush-delay (Subscriber Management)

Syntax gres-route-flush-delay;

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.2.

Description For a subscriber network configured with either nonstop active routing (NSR) or graceful restart, configure the router to wait 180 seconds (3 minutes) before removing (flushing) static or dynamic access routes and access-internal routes from the forwarding table after a graceful Routing Engine switchover (GRES) has taken place.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Delaying Removal of Access Routes and Access-Internal Routes After Graceful Routing Engine Switchover on page 88
- Overview of Access Routes and Access-Internal Routes Removal After Graceful Routing Engine Switchover on page 87

 Copyright © 2016, Juniper Networks, Inc.

group-profile (Group Profile)

Syntax

```
group-profile profile-name {
  l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
  }
  ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-pool pool-id;
    idle-timeout seconds;
    interface-id interface-id;
    keepalive seconds;
    ppp-options {
      chap;
      pap;
    }
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the group profile.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *profile-name*—Name assigned to the group profile.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

holddown-interval

Syntax holddown-interval *milliseconds*;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.

Options ***milliseconds***—Interval specifying how long a BFD session must remain up before a state change notification is sent.

Range: 0 through 255,000

Default: 0

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

 Copyright © 2016, Juniper Networks, Inc.

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the keepalive timer for L2TP tunnels.

NOTE: **Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.**

Options ***seconds***—Interval, in seconds, after which the server sends a hello message if no messages are received. A value of **0** means that no hello messages are sent.

Default: 60 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Timers for L2TP Tunnels*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

home-agent (Mobile IP Dynamic Assignment)

Syntax

```
home-agent {
  nai (name@domain | @domain) {
    home-agent ip-address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment],
[edit routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment],
[edit services mobile-ip dynamic-home-assignment]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip** dynamic-home-assignment**]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip** dynamic-home-assignment**]**, and **[edit routing-instances *routing-instances-name* services mobile-ip** dynamic-home-assignment**]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the IP address to which registration requests are sent as part of the home agent's dynamic assignment rule.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring Dynamic Home Assignment for the Mobile Node on page 232

 Copyright © 2016, Juniper Networks, Inc.

home-agent (Mobile IP Network Address Identifier)

Syntax

home-agent *ip-address*;

Hierarchy Level

[edit services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai *name*@*domain*],
[edit services mobile-ip dynamic-home-assignment home-agent nai @*domain*],
[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain]*,
[edit routing-instances *routing-instance-name* services mobile-ip dynamic-home-assignment home-agent nai @*domain*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the IP address to which registration requests are sent as part of the home agent's dynamic assignment rule.

Options

ip-address—IP address of the home agent

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring Dynamic Home Assignment for the Mobile Node on page 232

home-agent (Mobile IP Networks)

Syntax

```
home-agent {
  enable-service interface-name;
  virtual-network {
    home-agent-address ip-address {
      registration-lifetime seconds;
      revocation-required;
      timestamp-tolerance seconds;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip]**, and **[edit routing-instances *routing-instances-name* services mobile-ip]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the virtual networks and non-virtual networks for the Mobile IP home agent.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229

Copyright © 2016, Juniper Networks, Inc.

home-agent-address

Syntax

```
home-agent-address ip-address {
   registration-lifetime seconds;
   revocation-required;
   timestamp-tolerance seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network],
[edit services mobile-ip home-agent virtual-network]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **edit logical-systems *logical-system-name* services mobile-ip home-agent** virtual-network], **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent** virtual-network], and **[edit routing-instances *routing-instances-name* services mobile-ip home-agent** virtual-network] hierarchy levels introduced in Junos OS Release 9.5.

Description

Defines addressing for the virtual network of the Mobile IP home agent.

Options

ip-address—For virtual networks, the loopback IP address for the virtual network. For non-virtual networks, a public address.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

identification (Tunnel Profile)

Syntax identification *name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the assignment ID of an L2TP tunnel. L2TP sessions with the same tunnel assignment identification and destination are grouped into the same tunnel.

Options ***name***—Tunnel assignment ID; string of up to 32 alphanumeric characters.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

 Copyright © 2016, Juniper Networks, Inc.

idle-timeout (Access)

Syntax idle-timeout *seconds*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.

Options ***seconds***—Number of seconds a user can remain idle before the session is terminated.

Range: 0 through 4,294,967,295 seconds

Default: 0

NOTE: **The [edit access] hierarchy is not available on QFabric systems.**

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring PPP Properties for a Client-Specific Profile*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

idle-timeout (L2TP)

Syntax idle-timeout *seconds*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify how long a tunnel is active after its last session is terminated. The timer starts when the session is terminated and the tunnel is disconnected when the timer expires.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel idle-timeout.**

Options *seconds*—Length of the idle timeout. A value of **0** creates a persistent tunnel; that is, the tunnel remains active indefinitely until the remote peer disconnects it or you issue the **clear services l2tp tunnel** command.

Range: 0 through 86,400

Default: 60

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the L2TP Tunnel Idle Timeout on page 137
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

inline-services (FPC Level)

Syntax

```
inline-services {
  flow-table-size {
    ipv4-flow-table-size units;
    ipv6-flow-table-size units;
  }
}
```

Hierarchy Level [edit chassis fpc *slot-number*]

Release Information Statement introduced in Junos OS Release 12.1.

Description Enable inline services on MPCs, configured at the FPC level. To enable inline services that are specified at the PIC level, see configuration statement **inline-services (PIC level)**.

Options The remaining statements are defined separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Inline Service Interfaces on page 184
- Configuring an L2TP LNS with Inline Service Interfaces on page 175
- *Configuring Inline Sampling*

inline-services (PIC level)

Syntax

```
inline-services {
  bandwidth (1g | 10g);
}
```

Hierarchy Level [edit chassis fpc *slot-number* pic *number*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Enable inline services on PICs residing on MPCs. To enable inline services that are specified at the fpc level, see configuration statement **inline-services (FPC Level)**

The remaining statement is explained separately.

Options The option is described separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Inline Service Interfaces on page 184
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

interface (Dynamic Routing Instances)

Syntax interface *interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name*routing-instances *routing-instance-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Assign the specified interface to the dynamically created routing instance.

Options ***interface-name***—The interface name variable (*$junos-interface-name*). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *[edit routing-instances] Hierarchy Level*

interface (L2TP Service Interfaces)

Syntax interface *service-interface-name*;

Hierarchy Level [edit services service-device-pools pool *pool-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify a service interface assigned to a service interface pool. You specify more than one interface for each pool; the interfaces are used by an L2TP tunnel group to balance traffic loads.

Options ***service-interface-name***—Name of the service interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 199
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

 Copyright © 2016, Juniper Networks, Inc.

interface-id

Syntax interface-id *interface-id*;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ike],
[edit access profile *profile-name* client *client-name* l2tp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface identifier.

Options ***interface-id***—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring an IKE Access Profile*
- Configuring an L2TP Access Profile on the LNS on page 180

ip-address-change-notify

Syntax ip-address-change-notify *message*;

Hierarchy Level [edit access profile *profile-name* radius options]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the Unisphere-IPv4-release-control VSA in RADIUS messages. When enabled, the BNG includes Unisphere-Ipv4-release-control VSA in the Access-Request that is sent during on-demand IP address allocation and in the immediate Interim-Accounting messages that are sent to report an address change. Disabled by default, there is no effect when on-demand IP address allocation or deallocation is not configured. An change takes effect immediately. It is optional to specify the message, but if specified, the message is inserted into Unisphere-Ipv4-release-control VSA. Otherwise, a default value (NO MESSAGE) is be inserted into the VSA.

Options ***message***—VSA message.

Range: 1 through 32 characters,

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

 Copyright © 2016, Juniper Networks, Inc.

ip-reassembly

Syntax

```
ip-reassembly {
   profile profile-name
   rule rule-name{
      match-direction direction
   };
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the IP reassembly parameters to be applied to the L2TP server.

NOTE: Inline IP reassembly configuration does not require you to configure the profile statement. The profile configuration is used when IP reassembly is configured on services PICs.

Options **profile *profile-name***—Name of the IP reassembly profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 204
- IP Packet Fragment Reassembly for L2TP Overview on page 203

ip-reassembly (L2TP)

Syntax

```
ip-reassembly {
   service-set service-set-name;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.1.

Description Associate the reassembly service-set with the L2TP service.

NOTE: **The service set must be defined at the [edit services] hierarchy level.**

Options

service-set *service-set-name*—Identifies the service set to be associated with the L2TP service.

Required Privilege Level

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

Related Documentation

- IP Packet Fragment Reassembly for L2TP Overview on page 203
- Configuring IP Inline Reassembly for L2TP on page 204

Copyright © 2016, Juniper Networks, Inc.

ip-reassembly-rules (Service Set)

Syntax ip-reassembly-rules *rule-name*;

Hierarchy Level [edit services service-set*service-set-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Specify one or more previously configured IP reassembly rules to associate with the service set.

NOTE: **The IP reassembly rule must be defined at the [edit services ip-reassembly rule] hierarchy level.**

Options ***rule-name***—Name of an IP reassembly rule.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring IP Inline Reassembly for L2TP on page 204
- IP Packet Fragment Reassembly for L2TP Overview on page 203

Copyright © 2016, Juniper Networks, Inc.

initiate-ncp (Dynamic and Static PPP)

Syntax

initiate-ncp (ip | ipv6 | dual-stack-passive);

Hierarchy Level

[edit access group-profile *profile-name* ppp ppp-options],
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit *logical-unit-number* ppp-options],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number* ppp-options]

Release Information

Statement introduced in Junos OS Release 14.1.

Description

Configure PPP Network Control Protocol (NCP) negotiation mode (active or passive) for dynamic and static IPv4 and IPv6 PPP subscriber interfaces. You can also configure PPP NCP negotiation mode for the PPP server in an IPv4/IPv6 dual-stack configuration.

Options

dual-stack-passive—Enable passive PPP NCP negotiation for the PPP server in an IPv4/IPv6 dual-stack configuration. The **initiate-ncp dual-stack-passive** statement overrides the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured in an IPv4/IPv6 dual-stack configuration.

ip—Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv4 (**inet**) protocol address family, and for which IPv4 address attributes are assigned during authorization. By default, dynamic and static IPv4 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the **initiate-ncp ip** statement to enable active PPP NCP negotiation for the IPv4 subscriber interface.

ipv6—Enable active PPP NCP negotiation for dynamic and static PPP subscriber interfaces configured with the IPv6 (**inet6**) protocol address family, and for which IPv6 address attributes are assigned during authorization. By default, dynamic and static IPv6 subscriber interfaces use passive PPP NCP negotiation. In an IPv4/IPv6 dual-stack configuration, use the **initiate-ncp ipv6** statement to enable active PPP NCP negotiation for the IPv6 subscriber interface.

Required Privilege Level

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

Related Documentation

- Configuring the PPP Network Control Protocol Negotiation Mode on page 112
- PPP Network Control Protocol Negotiation Mode Overview on page 107

Copyright © 2016, Juniper Networks, Inc.

keepalive

Syntax keepalive *seconds*;

Hierarchy Level [edit access group-profile *profile-name* ppp],
[edit access profile *profile-name* client *client-name* ppp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the keepalive interval for an L2TP tunnel.

Options ***seconds***—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer.

For L2TP on MX Series routers, the minimum recommended interval is 30 seconds. A value of 0 disables generation of keepalive messages from the LNS.

Range: 0 through 32,767 seconds

Default: 30 seconds

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring PPP Properties for a Client-Specific Profile*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

keepalives

Syntax keepalives <interval *seconds*> <down-count *number*> <up-count *number*>;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.

For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Default Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.

Options **down-count *number***—The number of keepalive packets a destination must fail to receive before the network takes down a link.

Range: 1 through 255

Default: 3

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 10 seconds

up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up.

Range: 1 through 255

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Keepalives*
- *Configuring Frame Relay Keepalives*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

 Copyright © 2016, Juniper Networks, Inc.

keepalives (Dynamic Profiles)

Syntax

```
keepalives {
   interval seconds;
   }
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit *logical-unit-number*]
[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 10.1.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.

Description

Specify the keepalive interval in a PPP dynamic profile.

Default

Sending of keepalives is enabled by default.

Options

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

key

Syntax

key (hex | ascii) *string*;

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *name@domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer nai *name@domain* spi *hexadecimal-value*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the authentication key for the security association, in either HEX or ASCII format. The resulting 128–bit key is specified as a hexadecimal number with each character in the range 0x0–0xF.

Options

hex *string*—Key specified in HEX format

ascii *string*—Key specified in ASCII format

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

Copyright © 2016, Juniper Networks, Inc.

l2tp

Syntax

```
l2tp {
  destination [ip-address] {
    access-line-information <connection-speed-update>;
    lockout-timeout seconds;
  }
  destruct-timeout seconds;
  disable-calling-number-avp;
  disable-failover-protocol;
  enable-snmp-tunnel-statistics;
  failover-within-preference;
  ip-reassembly;
  rx-connect-speed-when-equal;
  traceoptions {
    debug-level level;
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    filter {
      protocol name;
      user user@domain;
      user-name username;
    }
    flag flag;
    interfaces interface-name {
      debug-level severity;
      flag flag;
    }
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
  tunnel {
    assignment-id-format (assignment-id |client-server-id);
    idle-timeout seconds;
    minimum-retransmission-timeout;
    nas-port-method;
    retransmission-count-established count;
    retransmission-count-not-established count;
    rx-window-size packets;
    tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
  }
  tunnel-group group-name {
    aaa-access-profile profile-name;
    dynamic-profile profile-name;
    hello-interval seconds;
    hide-avps;
    l2tp-access-profile profile-name;
    local-gateway address {
      address address;
      gateway-name gateway-name;
    }
    maximum-send-window packets;
    ppp-access-profile profile-name;
    receive-window packets;
```

```
    retransmit-interval seconds;
    service-device-pool pool-name;
    service-interface interface-name;
    syslog {
      host hostname {
        facility-override facility-name;
        log-prefix prefix-value;
        services severity-level;
      }
    }
    tos-reflect;
    tunnel-switch-profile profile-name;
    tunnel-timeout seconds;
  }
  tunnel-switch-profile profile-name;
  tx-connect-speed-method method;
  weighted-load-balancing;
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description Configure L2TP services to establish PPP tunnels across a network.

The remaining statements are explained separately.

NOTE: **Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Layer 2 Tunneling Protocol Overview*
- L2TP for Subscriber Access Overview on page 123

 Copyright © 2016, Juniper Networks, Inc.

l2tp-access-profile

Syntax l2tp-access-profile *profile-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the profile used to validate all L2TP connection requests to the local gateway address.

Options ***profile-name***—Identifier for the L2TP connection profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Access Profiles for L2TP Tunnel Groups*
- Configuring an L2TP Access Profile on the LNS on page 180

lcp-renegotiation

Syntax lcp-renegotiation;

Hierarchy Level [edit access group-profile *profile-name* l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the L2TP network server (LNS) so it renegotiates the link control protocol (LCP) with the PPP client. When LCP renegotiation is disabled, LNS uses the pre-negotiated LCP parameters between the L2TP access concentrator (LAC) and PPP client to set up the session. When LCP renegotiation is enabled, authentication is also renegotiated.

NOTE: **This statement is not supported at the [edit access group-profile l2tp] hierarchy level for L2TP LNS on MX Series routers.**

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*

liveness-detection

Syntax

```
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode(automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name*],
[edit forwarding-options dhcp-relay group *group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure bidirectional failure detection timers and authentication criteria for static routes.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- DHCP Liveness Detection Overview on page 75
- Configuring Detection of DHCP Local Server Client Connectivity on page 81
- Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 77
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83

Copyright © 2016, Juniper Networks, Inc.

- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

local-gateway (L2TP LNS)

Syntax

```
local-gateway {
  address address;
  gateway-name gateway-name;
}
```

Hierarchy Level

[edit services l2tp tunnel-group *name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the IP address or name for the local (LNS) gateway for L2TP tunnel.

The remaining statements are explained separately.

Options

address—Local IP address; corresponds to the IP address that is used by LACs to identify the LNS. When the LAC is an MX Series router, this address matches the remote gateway address configured in the LAC tunnel profile.

Required Privilege Level

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

Related Documentation

- *Configuring the Local Gateway Address and PIC*.
- *Configuring L2TP Tunnel Groups*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

lockout-timeout (L2TP Destination Lockout)

Syntax lockout-timeout *seconds*;

Hierarchy Level [edit services l2tp destination]

Release Information Statement introduced in Junos OS Release 13.2.

Description Set the duration of the timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created. Destinations are locked out when L2TP cannot connect to the destination during the tunnel selection process. This statement does not affect destinations that are currently locked out.

NOTE: The *ip-address* option for the **destination** statement does not apply to the **lockout-timeout** statement.

Options *seconds*—Length of the period during which the destination is locked out.

Range: 60 through 3600 seconds

Default: 300 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the L2TP Destination Lockout Timeout on page 139
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

logical-system (Tunnel Profile)

Syntax logical-system *logical-system-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify a logical system for a tunnel. When you specify a logical system, you must also specify a routing instance.

Options ***logical-system-name***— Name of the logical system.

Default: Logical system *default*

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

mac-address (Dynamic Access-Internal Routes)

Syntax mac-address *address*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*],
[edit dynamic-profiles routing-options access-internal route *subscriber-ip-address* qualified-next-hop *underlying-interface*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address* qualified-next-hop *underlying-interface*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the MAC address variable for an access-internal route for unnumbered interfaces such as DHCP subscriber interfaces.

Options *address*—Either the specific MAC address you want to assign to the access-internal route or the MAC address variable ($junos-subscriber-mac-address). The MAC address variable is dynamically replaced with the value supplied by DHCP when a subscriber logs in.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44

 Copyright © 2016, Juniper Networks, Inc.

match-direction (IP Reassembly Rule)

Syntax match-direction *direction*

Hierarchy Level [edit services ip-reassemblyrule*rule-name*]

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the direction in which the IP reassembly rule matching is applied. The match direction is used with respect to the traffic flow through the inline services interface. You must configure a match direction for an IP reassembly rule.

Options ***direction***—Match direction. For inline IP reassembly, **input** is the only match direction supported.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring IP Inline Reassembly for L2TP on page 204
- IP Packet Fragment Reassembly for L2TP Overview on page 203

maximum-sessions-per-tunnel

Syntax maximum-sessions-per-tunnel *number*;

Hierarchy Level [edit access group-profile l2tp],
[edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the maximum sessions for a Layer 2 tunnel.

NOTE: **This statement is not supported at the [edit access group-profile l2tp] hierarchy level for L2TP LNS on MX Series routers.**

Options ***number***—Maximum number of sessions for a Layer 2 tunnel.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Configuring the Group Profile for Defining L2TP Attributes*
- *Configuring L2TP Properties for a Client-Specific Profile*

max-sessions (Tunnel Profile)

Syntax max-sessions *number*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the maximum number of sessions allowed in the tunnel.

Options *number*—Maximum number of sessions allowed in the tunnel.
Range: 0 through 60,000
Default: 0

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 162

medium (Tunnel Profile)

Syntax medium *type*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the medium type for the tunnel.

Default ipv4

Options *type*—Medium type for the tunnel. The only value currently available is **ipv4**.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 162

 Copyright © 2016, Juniper Networks, Inc.

method

Syntax

```
method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
```

Hierarchy Level

[edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit forwarding-options dhcp-relay liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group *group-name* liveness-detection],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the liveness detection method.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

metric (Dynamic Access-Internal Routes)

Syntax metric *route-cost*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the cost for an access route.

Options *route-cost*—Either the specific cost you want to assign to the access route or either of the following cost variables:

- **$junos-framed-route-cost**—Cost of an IPv4 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-cost**—Cost of an IPv6 access route; the variable is dynamically replaced with the metric value (Subattribute 3) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 45

 Copyright © 2016, Juniper Networks, Inc.

minimum-interval

Syntax minimum-interval *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd transmit-interval]

Release Information
Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimal-interval** and **minimum-receive-interval** statements.

Options *milliseconds* — Specify the minimum interval value for BFD liveliness detection.

Range: 1 through 255,000

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

minimum-receive-interval

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.

Options ***milliseconds*** — Specify the minimum receive interval value.

Range: 1 through 255,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

 Copyright © 2016, Juniper Networks, Inc.

minimum-retransmission-timeout (L2TP Tunnel)

Syntax minimum-retransmission-timeout *seconds*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the minimum (initial) interval that the LAC or the LNS waits for a response after transmitting an L2TP control message to a peer. If no response has been received by the time the period expires, the message is retransmitted. The timeout period is doubled for each retransmission until the maximum of 16 seconds is reached.

Options *seconds*—Minimum interval before initial retransmission.

Range: 1, 2, 4, 8, or 16 seconds

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Retransmission Attributes for L2TP Control Messages on page 143
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

mobile-ip

Syntax

```
mobile-ip {
  access-type {
    (generic | wimax);
  }
  authenticate {
    order (aaa | local);
  }
  dynamic-home-assignment {
    home-agent {
      nai (name@domain | @domain) {
        home-agent ip-address;
      }
    }
  }
  home-agent {
    enable-service interface-name;
    virtual-network {
      home-agent-address ip-address {
        registration-lifetime seconds;
        revocation-required;
        timestamp-tolerance seconds;
      }
    }
  }
  peer {
    (ip-address address | nai name@domain) {
      spi hexadecimal-value {
        algorithm (hmac-md5 | md5);
        entity-type (host | mobility-agent);
        key (hex | ascii) string;
        replay-method (none | timestamp seconds);
      }
    }
  }
  traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services]**, and **[edit**

Copyright © 2016, Juniper Networks, Inc.

routing-instances *routing-instances-name* services] hierarchy levels introduced in Junos OS Release 9.5.

Description Configure Junos Mobile IP features.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Mobile IP Home Agent on page 230

multiplier

Syntax multiplier *number*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.

Options **number**—Maximum allowable number of hello packets missed by the neighbor.

Range: 1 through 255

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

no-adaptation

Syntax

no-adaptation;

Hierarchy Level

[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.

Required Privilege Level

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

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

Copyright © 2016, Juniper Networks, Inc.

nai

Syntax

```
nai (name@domain | @domain) {
  home-agent ip-address;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip dynamic-home-assignment home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip dynamic-home-assignment home-agent],
[edit services mobile-ip dynamic-home-assignment home-agent]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the network address identifiers (NAI) to which registration requests are sent as part of the home agent's dynamic assignment rule .

Options

name*@*domain—User at a specified domain

@***domain***—All users at a specified domain

NOTE: The *name* can include only alphanumeric characters, dots, hyphens, or underscores. The *name* cannot end in @; @ must be used to separate *name* and *domain*. The *domain* can include only alphanumeric characters, dots, or hyphens. The *domain* must be in the format *domain.suffix*, where the *suffix* is com, org, net, and so on. The *suffix* must consist of at least two alphanumeric characters.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring Dynamic Home Assignment for the Mobile Node on page 232

nas-port-method (L2TP LAC)

Syntax nas-port-method cisco-avp;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 14.1.

Description Globally configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.

NOTE: This global configuration can be overridden by the configuration in a tunnel profile or by the RADIUS configuration.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Globally Configuring the LAC to Interoperate with Cisco LNS Devices on page 153
- Configuring a Tunnel Profile for Subscriber Access on page 162
- LAC Interoperation with Third-Party LNS Devices on page 152

nas-port-method (Tunnel Profile)

Syntax nas-port-method cisco-avp;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure the LAC to interoperate with Cisco LNS devices by including the Cisco NAS Port Info AVP (100) in the ICRQ to the LNS. This AVP conveys the physical NAS port number identifier and the type of the physical port, such as Ethernet or ATM.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

 Copyright © 2016, Juniper Networks, Inc.

next-hop (Dynamic Access-Internal Routes)

Syntax

next-hop *next-hop*;

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.

Options

next-hop—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.

- For IPv4 access routes, use the variable, **$junos-framed-route-nexthop**. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, **$junos-framed-route-ipv6-nexthop**. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45

next-hop-service

Syntax

```
next-hop-service {
   inside-service-interface interface-name.unit-number;
   outside-service-interface interface-name.unit-number;
   outside-service-interface-type local;
   service-interface-pool name;
}
```

Hierarchy Level

[edit services service-set *service-set-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
service-interface-pool option added in Junos OS Release 9.3.

Description

Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.

Options

inside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied inside the network.

outside-service-interface ***interface-name.unit-number***—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type ***interface-type***—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.

service-interface-pool ***name***—Name of the pool of logical interfaces configured at the **[edit services service-interface-pools pool** ***pool-name*****]** hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

NOTE: **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Required Privilege Level

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

Related Documentation

- *Configuring Service Sets to be Applied to Services Interfaces*

Copyright © 2016, Juniper Networks, Inc.

no-allow-snooped-clients

Syntax no-allow-snooped-clients;

Hierarchy Level
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides],
[edit forwarding-options dhcp-relay dhcpv6 overrides],
[edit forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay ...],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information
Statement introduced in Junos OS Release 10.2.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description
Explicitly disable DHCP snooping support on the router or switch.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to explicitly disable snooping support on the router or switch for DHCPv6 relay agent.

NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is *enabled* by default. In Release 10.1 and later, DHCP snooping is *disabled* by default.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Extended DHCP Relay Agent Overview*
- *Overriding the Default DHCP Relay Configuration Settings*
- DHCP Snooping Support on page 51

on-demand-ip-address

Syntax on-demand-ip-address;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit interfaces pp0 unit *unit-number* ppp-options],
[[edit protocols ppp-service] on page 280]

Release Information Statement introduced in Junos OS Release 13.1.

Description Allocates and de-allocates an IPv4 address after initial PPP authentication for a subscriber who does not have an existing IPv4 address and can be configured at either the interface level or at the system level. Disabled by default. When configured at the interface level, dynamic profile changes take effect only for any new subscriber logins. Changes for static PPP IFLs logs out the subscriber. When configured at the system level, globally enables an on-demand-ip-address for PPP subscribers. If configured at both the interface level and the system level, the system level configuration takes precedence and changes take effect only for new subscriber logins.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Saving IPv4 Addresses for Dual-Stack PPP Subscribers*

 Copyright © 2016, Juniper Networks, Inc.

order (Mobile IP)

Syntax

order (aaa | local);

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip authenticate],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip authenticate],
[edit routing-instances *routing-instances-name* services mobile-ip authenticate],
[edit services mobile-ip authenticate]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication method performed for Mobile IP.

Default

AAA is the default authentication method.

Options

aaa—Authentication is performed by AAA. This option is available only in the default router and default routing instance, and therefore only in the **[edit services mobile-ip]** hierarchy level.

local—Authentication is performed using parameters defined in the local database.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Access Type for Mobile IP on page 232

overrides (DHCP Relay Agent)

Syntax

```
overrides {
  allow-snooped-clients;
  allow-no-end-options;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match (option60-and-option82 | incoming-interface);
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.3.
Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Support for the **delete-binding-on-renegotiation** statement introduced in Junos OS Release 13.2 for EX Series switches.
Support for the **allow-no-end-options** statement introduced in Junos OS Release 14.1X53 for EX Series switches.

Description

Override the default configuration settings for the extended DHCP relay agent. Specifying the **overrides** statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.

The following statements are supported at both the **[edit ... dhcp-relay]** and **[edit ... dhcpv6]** hierarchy levels. All other statements are supported at the **dhcp-relay** hierarchy levels only.

- **allow-snooped-clients**

Copyright © 2016, Juniper Networks, Inc.

- **interface-client-limit**
- **no-allow-snooped-clients**
- **no-bind-on-request**
- **send-release-on-delete**

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Relay Agent Overview*
- *Overriding the Default DHCP Relay Configuration Settings*
- *Deleting DHCP Local Server and DHCP Relay Override Settings*
- dhcp-relay on page 308

pap

Syntax

```
pap {
   access-profile name;
   default-pap-password password;
   local-name name;
   local-password password;
   passive;
}
```

Hierarchy Level

[edit interfaces *interface-name* ppp-options],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

BEST PRACTICE: **On inline service (si) interfaces for L2TP, only the pap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.**

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Configuring PPP PAP Authentication*
- *Tracing Operations of the pppd Process*
- *traceoptions (PPP Process)*
- *Example: Configuring PAP for an L2TP Profile*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

Copyright © 2016, Juniper Networks, Inc.

pap (Dynamic PPP)

Syntax pap;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options], [edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]** hierarchy level introduced in Junos OS Release 12.2.

Description Specify PAP authentication in a PPP dynamic profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Dynamic Profiles Overview*
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

pap (L2TP)

Syntax pap;

Hierarchy Level [edit access group-profile *profile-name* ppp ppp-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description (MX Series routers only) Specify PAP authentication for PPP subscribers in an L2TP LNS user group profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

peer

Syntax

```
peer {
  (ip-address address | nai name@domain) {
    spi hexadecimal-value {
      algorithm (hmac-md5 | md5);
      entity-type (host | mobility-agent);
      key (hex | ascii) string;
      replay-method (timestamp seconds | none);
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the authentication configurations for a home agent mobile node. An authentication enables the registration message as acceptable to the final recipient of the registration message.

Options

ip-address *address*—IP address of the peer.

nai *name@domain*—Network address identifier (NAI) of the peer. The *name* can include only alphanumeric characters, dots, hyphens, or underscores. The *name* cannot end in @; @ must be used to separate *name* and *domain*. The *domain* can include only alphanumeric characters, dots, or hyphens The *domain* must be in the format *domain.suffix*, where the *suffix* is com, org, net, and so on. The *suffix* must consist of at least two alphanumeric characters.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

Copyright © 2016, Juniper Networks, Inc.

pic (M Series and T Series Routers)

Syntax

```
pic pic-number {
  ce1 {
    e1 port-number {
      channel-group group-number timeslots slot-number;
    }
  }
  ct3 {
    port port-number {
      t1 link-number {
        channel-group group-number timeslots slot-number;
      }
    }
  }
  framing (sdh | sonet);
  idle-cell format {
    itu-t;
    payload-pattern payload-pattern-byte;
  }
  inline-services {
    bandwidth (1g | 10g);
  }
  max-queues-per-interface (8 | 4);
  no-concatenate;
}
```

Hierarchy Level

[edit chassis fpc *slot-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure properties for an individual PIC.

Options

pic-number—Slot number in which the PIC is installed.

Range: 0 through 3

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the Junos OS to Enable SONET/SDH Framing for SONET/SDH PICs*
- *Configuring the Junos OS to Enable a SONET PIC to Operate in Channelized (Multiplexed) Mode*
- *Configuring the Junos OS to Support Channelized DS3-to-DS0 Naming for Channel Groups and Time Slots*
- *Configuring the Junos OS to Support Channel Groups and Time Slots for Channelized E1 PICs*

pool (L2TP Service Interfaces)

Syntax

```
pool pool-name {
   interface service-interface-name;
}
```

Hierarchy Level

[edit services service-device-pools]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Define a pool of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.

Options

pool-name—Name of the service interface pool.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 199

Copyright © 2016, Juniper Networks, Inc.

ppp (Group Profile)

Syntax

```
ppp {
   cell-overhead;
   encapsulation-overhead bytes;
   framed-pool framed-pool;
   idle-timeout seconds;
   interface-id interface-id;
   keepalive seconds;
   ppp-options {
      chap;
      pap;
   }
   primary-dns primary-dns;
   primary-wins primary-wins;
   secondary-dns secondary-dns;
   secondary-wins secondary-wins;
}
```

Hierarchy Level

[edit access group-profile *profile-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure PPP properties for a group profile.

The remaining statements are explained separately.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Required Privilege Level

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

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

ppp-options

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  initiate-ncp (ip | ipv6 | dual-stack-passive)
  lcp-max-conf-req number
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-max-conf-req number
  ncp-restart-timer milliseconds;
  on-demand-ip-address
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to

Copyright © 2016, Juniper Networks, Inc.

ppp-options—including those subordinate to chap and pap—at their default values.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

ppp-options (Dynamic PPP)

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
  }
  initiate-ncp (ip | ipv6 | dual-stack-passive)
  on-demand-ip-address;
  pap;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.

Description

Configure PPP-specific interface properties in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- Configuring Dynamic Authentication for PPP Subscribers on page 103
- Attaching Dynamic Profiles to Static PPP Subscriber Interfaces on page 98
- Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface on page 177

ppp-options (L2TP)

Syntax

```
ppp-options {
  chap;
  initiate-ncp (ip | ipv6 | dual-stack-passive)
  pap;
}
```

Hierarchy Level

[edit access group-profile *profile-name* ppp]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Configure PPP-specific properties in a group profile that applies to tunneled PPP subscribers at the LNS.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the Group Profile for Defining L2TP Attributes*
- Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile on page 179

Copyright © 2016, Juniper Networks, Inc.

preference (Subscriber Management)

Syntax preference *route-distance*

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 9.5.
Support at **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the distance for an access route.

Options ***route-distance***—Either the specific distance you want to assign to the access route or either of the following distance variables:

- **$junos-framed-route-distance**—Distance of an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-distance**—Distance of an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 5) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 45

preference (Tunnel Profile)

Syntax preference *number*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the preference for a tunnel. You can specify up to 8 levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.

This value can be overridden by RADIUS attribute Tunnel-Preference [83].

Options *number*—Number that indicates the order in which the router attempts to connect to the destination. Zero is the highest level of preference.

Range: 0 through 2000

Default: 2000

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

 Copyright © 2016, Juniper Networks, Inc.

proxy-mode

Syntax

proxy-mode;

Hierarchy Level

[edit forwarding-options dhcp-relay overrides],
[edit forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* overrides],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* interface *interface-name* overrides]

Release Information

Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.

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

Required Privilege Level

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

Related Documentation

- *DHCP Relay Proxy Overview*
- *Extended DHCP Relay Agent Overview*
- *Enabling DHCP Relay Proxy Mode*

qualified-next-hop (Subscriber Management)

Syntax

```
qualified-next-hop interface-name {
    mac-address address;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal route *subscriber-ip-address*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal route *subscriber-ip-address*],
[edit dynamic-profiles *profile-name* routing-options access-internal route *subscriber-ip-address*]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options route *subscriber-ip-address*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* route *subscriber-ip-address*]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the qualified next-hop and the MAC address for an access-internal route for DHCP and PPP subscriber interfaces.

Options

interface-name—Either the specific interface you want to assign to the access route or the variable, or the **$junos-interface-name** variable. The variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44

Copyright © 2016, Juniper Networks, Inc.

registration-lifetime

Syntax registration-lifetime *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure maximum period for registration lifetime that is accepted by the Mobile IP home agent.

Options **registration-lifetime *seconds***—Maximum lifetime that the home agent accepts in any registration request. The registration lifetime is not affected if you change the system clock.

Range: 7 through 65535 seconds

Default: 3600 seconds

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

reject-unauthorized-ipv6cp

Syntax reject-unauthorized-ipv6cp;

Hierarchy Level [edit protocols ppp-service]

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure the router to reject any IPv6 Control Protocol (IPv6CP) negotiation messages on dynamic interfaces when no appropriate IPv6 address or prefix has been received from AAA. IPv6CP negotiation attempts are also rejected when only a Framed-IPv6-Prefix attribute is received but router advertisement is not enabled in the dynamic profile.

NOTE: **IPv6CP negotiation messages are not rejected for static interfaces.**

Default IPv6CP negotiation is allowed regardless of the presence of IPv6 attributes received from AAA.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Avoiding Negotiation of IPv6CP in the Absence of an Authorized Address*

 Copyright © 2016, Juniper Networks, Inc.

relay-option-82

Syntax

```
relay-option-82 {
  circuit-id {
    include-irb-and-l2;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
  }
  remote-id {
    include-irb-and-l2;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-vlan-id;
  }
}
```

Hierarchy Level

[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay group *group-name*],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name*],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay],
[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.

When you configure **relay-option-82** without configuring the **circuit-id** or **remote-id** option, the Agent Circuit ID is added by default.

You can use the **relay-option-82** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level to control insertion of option 82 information globally, or at the **[edit forwarding-options dhcp-relay group *group-name*]** hierarchy level to control insertion of option 82 information for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), use the **delete relay-option-82** statement.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Using DHCP Relay Agent Option 82 Information*
- dhcp-relay on page 308

remote-gateway (Tunnel Profile)

Syntax

```
remote-gateway {
   address server-ip-address;
   gateway-name server-name;
}
```

Hierarchy Level

[edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the IP address and hostname of the remote gateway at the L2TP tunnel endpoint, the LNS.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

Copyright © 2016, Juniper Networks, Inc.

replay-method

Syntax

replay-method (none | timestamp *seconds*);

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* services mobile-ip peer *nai@domain* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer *nai@domain* spi *hexadecimal-value*].
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address* spi *hexadecimal-value*].
[edit routing-instances *routing-instances-name* services mobile-ip peer *nai@domain* spi *hexadecimal-value*],
[edit services mobile-ip peer ip-address *address* spi *hexadecimal-value*],
[edit services mobile-ip peer *nai@domain* spi *hexadecimal-value*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Configure the replay protection method. The Identification field enables the home agent to verify that a registration message has been recently generated by the mobile node, rather than replayed by an attacker from a previous registration. You can specify a timestamp tolerance for the mobile node, which causes the request to be rejected if the tolerance is exceeded, or you can specify that the tolerance be taken from the value configured on the home agent.

Default

If you do not configure the replay protection method, then the timestamp tolerance is taken from the home agent by default.

Options

none—Timestamp tolerance is obtained from the setting configured for the home agent

timestamp *seconds*—Tolerance time in which a registration request timestamp and the local time of the home agent can differ.

Range: 1 through 255 seconds

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

request services l2tp destination unlock

Syntax **request services l2tp destination unlock *destination-name***

Release Information Command introduced in Junos OS Release 13.2.

Description Remove the specified destination from the destination lockout list immediately, before its lockout period expires. As a result, the L2TP process can again consider the destination during the selection of new tunnels.

Options ***destination-name***—Name of the destination to be unlocked.

Required Privilege Level view

Related Documentation

- Removing an L2TP Destination from the Destination Lockout List on page 139
- Configuring the L2TP Destination Lockout Timeout on page 139
- **show services l2tp destination lockout on page 511**

List of Sample Output **request services l2tp destination unlock on page 398**

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

Sample Output

request services l2tp destination unlock

```
user@host> request services l2tp destination unlock dest-a
Destination dest-a unlocked
```

 Copyright © 2016, Juniper Networks, Inc.

retransmission-count-established (L2TP)

Syntax retransmission-count-established *count*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the maximum number of times control messages are retransmitted for established tunnels.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel retransmission-count-established.**

Options *count*—Number of retransmissions.

Range: 2 through 30

Default: 7

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Retransmission Attributes for L2TP Control Messages on page 143
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

retransmission-count-not-established (L2TP)

Syntax retransmission-count-not-established *count*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the maximum number of times control messages are retransmitted for tunnels that are not established.

BEST PRACTICE: **Before you downgrade to a Junos OS Release that does not support this statement, unconfigure the statement by issuing no services l2tp tunnel retransmission-count-not-established.**

Options *count*—Number or retransmissions.

Range: 2 through 30

Default: 5

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Retransmission Attributes for L2TP Control Messages on page 143
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

revocation-required

Syntax revocation-required;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-networkhome-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the Mobile IP home agent to accept registration revocation requests only when the request includes the revocation extension.

Default The Mobile IP home agent supports registration revocation requests that include the revocation extension, but it does not require the extension.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

route (Access)

Syntax

```
route prefix {
  next-hop next-hop;
  metric route-cost;
  preference route-distance;
  tag route-tag;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access],
[edit dynamic-profiles *profile-name* routing-options access]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure the parameters for access routes.

Options

prefix—Either the specific route prefix that you want to assign to the access route or one of the following route prefix variables.

- For IPv4 access routes, use the variable, **$junos-framed-route-ip-address-prefix**. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22].
- For IPv6 access routes, use the variable, **$junos-framed-route-ipv6-address-prefix**. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access Routes for Subscriber Management on page 45

Copyright © 2016, Juniper Networks, Inc.

route (Access Internal)

Syntax

```
route subscriber-ip-address {
   next-hop next-hop;
   qualified-next-hop underlying-interface {
      mac-address address;
   }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal],
[edit dynamic-profiles *profile-name* routing-options access-internal]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access-internal]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access-internal]** hierarchy levels introduced in Junos OS Release 10.1.

Description

Dynamically configure parameters for an access-internal route.

Options

subscriber-ip-address—Either the specific IP address you want to assign to the access-internal route or the subscriber IP address variable ($junos-subscriber-ip-address). The subscriber IP address variable is dynamically replaced with the value supplied by DHCP or PPP when a subscriber logs in.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44
- Configuring Dynamic Access-Internal Routes for PPP Subscriber Management on page 101

route-suppression (DHCP Local Server and Relay Agent)

Syntax route-suppression (access | access-internal | destination);

Hierarchy Level
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay group *group-name*],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name*],
[edit logical-systems *logical-system-name* ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*...],
[edit routing-instances *routing-instance-name* ...],
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure the jdhcpd process to suppress the installation of access, access-internal, or destination routes during client binding.

NOTE: You cannot suppress access-internal routes when the subscriber is configured with both IA_NA and IA_PD addresses over IP demux interfaces—the IA_PD route relies on the IA_NA route for next hop connectivity.

Options

access—(DHCPv6 only) Suppress installation of access routes. You can use the **access** and **access-internal** options in the same statement for DHCPv6.

access-internal—In a DHCPv4 hierarchy, suppress installation of both access-internal and destination routes. In a DHCPv6 hierarchy, suppress access-internal routes only. Can be configured in the same statement with the **access** option.

destination—(DHCPv4 only) Suppress installation of destination routes. This option and the **access-internal** option are mutually exclusive; however, the **access-internal** option also suppresses destination routes.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Preventing DHCP from Installing Access, Access-Internal, and Destination Routes by Default on page 48

 Copyright © 2016, Juniper Networks, Inc.

routing-instance (Tunnel Profile)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify a routing instance for a tunnel.

Options ***routing-instance-name***—Name of the routing instance.

Default: Routing instance *default*

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

routing-instances (Dynamic Profiles)

Syntax

```
routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
    rib routing-table-name {
      access {
        route prefix {
          metric route-cost;
          next-hop next-hop;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles]

Release Information Statement introduced in Junos OS Release 9.6.
The **routing-options** statement introduced in Junos OS Release 10.1.

Description Dynamically configure an additional routing entity for a router.

 Copyright © 2016, Juniper Networks, Inc.

Options

routing-instance-name—The routing instance variable (*$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution*

routing-options (Dynamic Profiles)

Syntax

```
routing-options {
  access {
    route prefix {
      metric route-cost;
      next-hop next-hop;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]

Release Information

Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]** hierarchy level introduced in Junos OS Release 10.1.

Description

Configure protocol-independent routing properties in a dynamic profile.

The remaining statements are explained separately.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 45
- Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management on page 44

rule (IP Reassembly)

Syntax

```
rule rule-name {
   match-direction direction;
}
```

Hierarchy Level
[edit services ip-reassembly]

Release Information
Statement introduced in Junos OS Release 13.1.

Description
Configure an IP reassembly rule, which is used for inline IP reassembly on the inline services interface. The IP reassembly rule can be attached to a service set to indicate that the service set is of type IP reassembly. For inline IP reassembly, each rule must include the **match-direction** statement, which specifies the direction in which the match is applied.

NOTE: **If you configure an IP reassembly rule, then you must configure the match-direction statement.**

Options
rule-name—Name of the IP reassembly rule.
Range: Up to 63 characters

The remaining statement is explained separately.

NOTE: **To create more than one IP reassembly rule, include the rule statement multiple times.**

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring IP Inline Reassembly for L2TP on page 204
- IP Packet Fragment Reassembly for L2TP Overview on page 203

rx-connect-speed-when-equal (L2TP LAC)

Syntax rx-connect-speed-when-equal

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 13.1.

Description Enable sending the receive connect speed, which is represented by AVP 38, even when its value is equal to that of the transmit connect speed, which is represented by AVP 24. By default, AVP 38 is sent from the LAC to the LNS during the establishment of an L2TP tunnel session, only when its value is different from AVP 24. You can override the default setting with this configuration statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 172
- Transmission of the Receive Connect Speed AVP When Transmit and Receive Connect Speeds Are Equal on page 171

rx-window-size (L2TP)

Syntax rx-window-size *packets*;

Hierarchy Level [edit services l2tp tunnel]

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure the L2TP receive window size for an L2TP tunnel.

Options *packets*—Number of packets that a peer can transmit without receiving an acknowledgment from the router. By default, this value is set to 4 packets. If the receive window size is configured to its default value, the router does not send the Receive Window Size AVP (AVP 10) in the first tunnel negotiation packet that is sent to its peer.

Range: 4 through 128

Default: 4

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the L2TP Receive Window Size on page 137
- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

 Copyright © 2016, Juniper Networks, Inc.

secret (Tunnel Profile)

Syntax secret *password*;

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the tunnel password sent to the LNS for authentication.

Options ***password***—Cleartext password.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

service-device-pool (L2TP)

Syntax service-device-pool *pool-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Assign a pool of service interfaces to the tunnel group to balance traffic across.

NOTE: The service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options ***pool-name***—Name of the service device pool.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

service-device-pools (L2TP Service Interfaces)

Syntax

```
service-device-pools {
  pool pool-name {
    interface service-interface-name;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure one or more pools of service interfaces that can be assigned to an L2TP tunnel group for traffic load-balancing. The service device pool is required for dynamic LNS sessions.

Options *pool-name*—Name of the service interface pool.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Pool of Inline Services Interfaces for Dynamic LNS Sessions on page 199

 Copyright © 2016, Juniper Networks, Inc.

service-interface (L2TP Processing)

Syntax service-interface *interface-name*;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Option **si-*fpc/pic/port*** introduced in Junos OS Release 11.4.

Description Specify the service interface responsible for handling L2TP processing.

NOTE: On MX Series routers, the service interface configuration is required for static LNS sessions. Either the service interface configuration or the service device pool configuration can be used for dynamic LNS sessions.

Options ***interface-name***—Name of the service interface. The interface type depends on the line card as follows:

- **sp-*fpc/pic/port***—On AS or Multiservices PICs on M7i, M10i, and M120 routers.
- **si-*fpc/pic/port***—On MPCs on MX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Local Gateway Address and PIC*
- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

session-mode

Syntax session-mode (automatic | multihop | singlehop);

Hierarchy Level [edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the session mode.

Options **automatic**—Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface.

multihop—Configure multihop BFD sessions.

single-hop—Configure single hop BFD sessions.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

 Copyright © 2016, Juniper Networks, Inc.

shared-secret

Syntax shared-secret *shared-secret*;

Hierarchy Level [edit access profile *profile-name* client *client-name* l2tp]

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure the shared secret.

Options ***shared-secret***—Shared secret key for authenticating the peer.

Required Privilege Level admin—To view this statement in the configuration.
admin–control—To add this statement to the configuration.

Related Documentation
- *Configuring L2TP Properties for a Client-Specific Profile*
- Configuring an L2TP Access Profile on the LNS on page 180

source-gateway (Tunnel Profile)

Syntax

```
source-gateway {
   address client-ip-address;
   gateway-name client-name;
}
```

Hierarchy Level [edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify the IP address and hostname of the source gateway at the local L2TP tunnel endpoint, the LAC.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring a Tunnel Profile for Subscriber Access on page 162

spi

Syntax

```
spi hexadecimal-value {
   algorithm (hmac-md5 | md5);
   entity-type (host | mobility-agent);
   key (hex | ascii) string;
   replay-method (none | timestamp seconds);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address*],
[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*],
[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*],
[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*],
[edit services mobile-ip peer ip-address *address*],
[edit services mobile-ip peer nai *user@domain*]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* services mobile-ip peer ip-address *address*]**, **[edit logical-systems *logical-system-name* services mobile-ip peer nai *user@domain*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*]**, **[edit routing-instances *routing-instances-name* services mobile-ip peer ip-address *address*]**, and **[edit routing-instances *routing-instances-name* services mobile-ip peer nai *user@domain*]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the security parameter index for identifying a security context between a pair of nodes among the contexts available in the Mobility Security Association. The index selects the authentication algorithm and key.

Options

hexadecimal-value—Security parameter index identifier.

Range: 100 to FFFFFFFF

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

Copyright © 2016, Juniper Networks, Inc.

statistics (Access Profile)

Syntax statistics (time | volume-time);

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Option **volume-time** introduced in Junos OS Release 9.4.

Description Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options **time**—Collect uptime statistics only.

volume-time—Collect both volume and uptime statistics. This option is not available for Mobile IP.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Mobile IP Home Agent Elements and Behavior on page 217
- *Configuring Authentication and Accounting Parameters for Subscriber Access*

tag (Access)

Syntax tag *route-tag*;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*],
[edit dynamic-profiles *profile-name* routing-options access route *prefix*]

Release Information Statement introduced in Junos OS Release 10.2.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options access route *prefix*]** and **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance routing-options rib *routing-table-name* access route *prefix*]** hierarchy levels introduced in Junos OS Release 10.1.

Description Dynamically configure the tag for an access route.

Options ***route-tag***—Either the specific tag you want to assign to the access route or either of the following tag variables:

- **$junos-framed-route-tag**—Tag assigned to an IPv4 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-Route attribute [22].
- **$junos-framed-route-ipv6-tag**—Tag assigned to an IPv6 access route; the variable is dynamically replaced with the preference value (Subattribute 6) from the RADIUS Framed-IPv6-Route attribute [99].

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Dynamic Access Routes for Subscriber Management on page 45

 Copyright © 2016, Juniper Networks, Inc.

threshold (detection-time)

Syntax threshold *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd detection-time],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd detection-time],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd detection-time]

Release Information
Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

NOTE: **The threshold time must be greater than or equal to the minimum-interval or the minimum-receive-interval.**

Options ***milliseconds***— Value for the detection time adaptation threshold.

Range: 1 through 255,000

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

threshold (transmit-interval)

Syntax threshold *milliseconds*;

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd transmit-interval],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd transmit-interval],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd transmit-interval]

Release Information Statement introduced in Junos OS Release 12.1.

Description Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.

Options ***milliseconds*** — Threshold value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

NOTE: **The threshold value specified in the threshold statement must be greater than the value specified in the minimum-interval statement for the transmit-interval statement.**

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

 Copyright © 2016, Juniper Networks, Inc.

timestamp-tolerance

Syntax timestamp-tolerance *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent virtual-network home-agent-address *ip-address*],
[edit services mobile-ip home-agent virtual-network home-agent-address *ip-address*]

Release Information Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description Configure the acceptable difference between a registration request timestamp and the local time of the home agent.

Options **timestamp-tolerance *seconds***—Acceptable difference in time between a registration request timestamp and the local time of the home agent.

Range: 1 through 255 seconds

Default: 7 seconds

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

tos-reflect (L2TP)

Syntax tos-reflect;

Hierarchy Level [edit services l2tp tunnel-group *name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the LNS to reflect the IP ToS value in the inner IP header to the outer IP header. When CoS rewrite rules are also configured (**[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]**), the rewrite is performed only on the inner IP ToS; the outer IP TOS value is not affected.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dynamic CoS for an L2TP LNS Inline Service*

 Copyright © 2016, Juniper Networks, Inc.

traceoptions (Services L2TP)

Syntax

```
traceoptions {
  debug-level level;
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    protocol name;
    user user@domain;
    user-name username;
  }
  flag flag;
  interfaces interface-name {
    debug-level level;
    flag flag;
  }
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit services l2tp]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define tracing operations for L2TP processes.

Options **debug-level** ***level***—Trace level for PPP, L2TP, RADIUS, and UDP; this option does not apply to L2TP on MX Series routers:

- **detail**—Trace detailed debug information.
- **error**—Trace error information.
- **packet-dump**—Trace packet decoding information.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

- **protocol** *name*—One of the following protocols; this option does not apply to L2TP on MX Series routers:
  - **l2tp**

  - **ppp**
  - **radius**
  - **udp**

- **user** *user@domain*—Username of a subscriber; this option does not apply to L2TP on M Series routers. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.
- **user-name** *username*—Username of a subscriber; this option does not apply to L2TP on MX Series routers.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **configuration**—Trace configuration events.
- **events**—Trace interface events.
- **general**—Trace general events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization.
- **ipc-rx**—Trace IPC receive events.
- **ipc-tx**—Trace IPC transmit events.
- **memory**—Trace memory management code.
- **message**—Trace message processing code.
- **packet-error**—Trace packet error events.
- **parse**—Trace parsing events.
- **protocol**—Trace L2TP events.
- **receive-packets**—Trace received L2TP packets.
- **routing-process**—Trace routing process interactions.
- **routing-socket**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **states**—Trace state machine events.
- **timer**—Trace timer events.
- **transmit-packets**—Trace transmitted L2TP packets.
- **tunnel**—Trace tunnel events.

Copyright © 2016, Juniper Networks, Inc.

interfaces ***interface-name***—Apply L2TP traceoptions to a specific services interface. This option does not apply to L2TP on MX Series routers.

- **debug-level** ***level***—Trace level for the interface; this option does not apply to L2TP on MX Series routers:
  - **detail**—Trace detailed debug information.
  - **error**—Trace error information.
  - **extensive**—Trace all PIC debug information.
- **flag** ***flag***—Tracing operation to perform for the interface. This option does not apply to L2TP on MX Series routers. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:
  - **all**—Trace everything.
  - **ipc**—Trace L2TP Inter-Process Communication (IPC) messages between the PIC and the Routing Engine.
  - **packet-dump**—Dump each packet content based on debug level.
  - **protocol**—Trace L2TP, PPP, and multilink handling.
  - **system**—Trace packet processing on the PIC.

level—Specify level of tracing to perform. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: **error**

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

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

Related Documentation

- *Tracing L2TP Operations*
- Tracing L2TP Operations for Subscriber Access on page 249

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Mobile IP)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip],
[edit routing-instances *routing-instances-name* services mobile-ip],
[edit services mobile-ip]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define tracing operations for Mobile IP processes.

Options

file *filename*— Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. Ensure that filenames are unique for each logical system or routing instance in which Mobile IP is configured.

NOTE: Global messages (common to all logical systems and routing instances) are always saved in **/var/log/mipd**. Messages that are specific to a logical system or routing instance are never saved in **/var/log/mipd**. If you do not configure a trace filename for a logical system or routing instance, then nothing is traced for that entity.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **authentication**—Trace authentication operations.
- **binding**—Trace bindings.

- **event**—Trace events.
- **ha-fsm**—Trace home agent state machine operations.
- **home-agent**—Trace home agent operations.
- **interface-database**—Trace interface database operations.
- **packet**—Trace packet decoding operations.
- **protocol**—Trace protocol operations.
- **rtsock**—Trace routing socket operations.
- **session-db**—Trace session database events.
- **signal**—Trace signal operations.
- **subscriber**—Trace subscriber events.
- **timer**—Trace timer events.
- **trace**—Trace changes in tracing.
- **tunnel**—Trace tunneling operations.
- **user-interface**—Trace user interface events.

level—Specify level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Default: error

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Copyright © 2016, Juniper Networks, Inc.

Required Privilege Level

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

Related Documentation

- Tracing Mobile IP Operations for Subscriber Access on page 257

traceoptions (Protocols PPP Service)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    aci regular-expression;
    ari regular-expresion;
    service-name regular-expresion;
    underlying-interface interface-name;
    user user@domain;
  }
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit protocols ppp-service]

Release Information Statement introduced in Junos OS Release 9.5.
Option **user** introduced in Junos OS Release 14.1.

Description Define tracing operations for PPP service processes.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

BEST PRACTICE: **Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.**

- **aci** ***regular-expression***—Regular expression to match the agent circuit identifier provided by PPP client.
- **ari** ***regular-expression***—Regular expression to match the agent remote identifier provided by PPP client.

 Copyright © 2016, Juniper Networks, Inc.

- **service** ***regular-expression***—Regular expression to match the name of PPPoE service.
- **underlying-interface** ***interface-name***—Name of a PPP underlying interface. You cannot use a regular expression for this filter option.
- **user** *user@domain*—Username of a subscriber. Optionally use an asterisk (*) as a wildcard to substitute for characters at the beginning or end of either term or both terms.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **accounting-statistics**—Trace accounting statistics events.
- **all**—Trace all operations.
- **authentication**—Trace authentication events.
- **chap**—Trace CHAP events.
- **events**—Trace interface events.
- **gres**—Trace GRES events.
- **init**—Trace daemon initialization events.
- **interface-db**—Trace interface database events.
- **lcp**—Trace LCP state machine events.
- **memory**—Trace memory processing events.
- **ncp**—Trace NCP state machine events.
- **packet-error**—Trace packet error events.
- **pap**—Trace PAP events.
- **parse**—Trace parsing events.
- **profile**—Trace libdynamic profile events.
- **receive-packets**—Trace received PPP packets.
- **routing-process**—Trace routing process interactions.
- **rtp**—Trace real-time priority events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace session database interactions.
- **smi-services-sentry**—Trace SMI services requests and retries.
- **states**—Trace state machine events.
- **transmit-packets**—Trace transmitted PPP packets.
- **tunnel**—Trace L2TP tunneling events.

Copyright © 2016, Juniper Networks, Inc.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Default: error

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

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

Related Documentation

- Tracing PPP Service Operations for Subscriber Access on page 241

Copyright © 2016, Juniper Networks, Inc.

traceoptions (Subscriber Management)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level [edit system services subscriber-management]

Release Information Statement introduced in Junos OS Release 11.1.

Description Define tracing operations for subscriber management interface processes.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the filename within quotation marks. All files are placed in the directory **/var/log**.

files ***number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **database**—Trace database events.
- **general**—Trace general events.
- **issu**—Trace unified ISSU events.
- **server**—Trace server events.
- **session-db**—Trace session database interactions.
- **ui**—Trace user interface events.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- *Tracing Subscriber Management Database Operations for Subscriber Access*

transmit-interval

Syntax

```
transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
}
```

Hierarchy Level
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd]

Release Information
Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description
Configure the Bidirectional Forwarding Detection (BFD) transmit interval.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78

Copyright © 2016, Juniper Networks, Inc.

tunnel (L2TP)

Syntax

```
tunnel {
  assignment-id-format (assignment-id | client-server-id);
  idle-timeout seconds;
  minimum-retransmission-timeout;
  nas-port-method;
  retransmission-count-established count;
  retransmission-count-not-established count;
  rx-window-size
  tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);
}
```

Hierarchy Level

[edit services l2tp]

Release Information

Statement introduced in Junos OS Release 11.4.
Option **rx-window-size** introduced in Junos OS Release 13.2.

Description

Configure L2TP tunnel characteristics.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring an L2TP LAC on page 149
- Configuring an L2TP LNS with Inline Service Interfaces on page 175

tunnel (Tunnel Profile)

Syntax

```
tunnel tunnel-id {
  identification name;
  logical-system logical-system-name;
  max-sessions number;
  medium type;
  nas-port-method cisco-avp;
  preference number;
  remote-gateway {
    address server-ip-address;
    gateway-name server-name;
  }
  routing-instance routing-instance-name;
  secret password;
  source-gateway {
    address client-ip-address;
    gateway-name client-name;
  }
  type tunnel-type;
}
```

Hierarchy Level

[edit access tunnel-profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Define the attributes of a tunnel for the tunnel profile. You can define up to 31 tunnels for each tunnel profile.

Options

tunnel-id—Unique integer that identifies a tunnel defined within a profile. For a subscriber, RADIUS attributes and VSAs can supply or override the attributes defined here for the tunnel.

Range: 1 through 31

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

Copyright © 2016, Juniper Networks, Inc.

tunnel-group

Syntax

```
tunnel-group group-name {
  aaa-access-profile profile-name;
  dynamic-profile profile-name;
  hello-interval seconds;
  hide-avps;
  l2tp-access-profile profile-name;
  local-gateway address {
    address address;
    gateway-name gateway-name;
  }
  maximum-send-window packets;
  ppp-access-profile profile-name;
  receive-window packets;
  retransmit-interval seconds;
  service-device-pool pool-name;
  service-interface interface-name;
  syslog {
    host hostname {
      services severity-level;
      facility-override facility-name;
      log-prefix prefix-value;
    }
  }
  tos-reflect;
  tunnel-switch-profile profile-name;
  tunnel-timeout seconds;
}
```

Hierarchy Level

[edit services l2tp]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for MX Series routers introduced in Junos OS Release 11.4.

Description

Specify the L2TP tunnel properties.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

group-name—Identifier for the tunnel group.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring L2TP Tunnel Groups*

- Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces on page 198

tunnel-profile (L2TP Tunnel Switching)

Syntax tunnel-profile *profile-name*;

Hierarchy Level [edit access tunnel-switch-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the name of an L2TP tunnel profile that defines the tunnel to which PPP subscriber traffic is switched.

Options ***profile-name***—Unique name that identifies the tunnel profile; configured with the **tunnel-profile** statement at the **[edit access]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring L2TP Tunnel Switching on page 135

 Copyright © 2016, Juniper Networks, Inc.

tunnel-profile (Tunnel Profile)

Syntax

```
tunnel-profile profile-name {
  tunnel tunnel-id {
    identification name;
    logical-system logical-system-name;
    max-sessions number;
    medium type;
    nas-port-method cisco-avp;
    preference number;
    remote-gateway {
      address server-ip-address;
      gateway-name server-name;
    }
    routing-instance routing-instance-name;
    secret password;
    source-gateway {
      address client-ip-address;
      gateway-name client-name;
    }
    type tunnel-type;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Description Define a tunnel profile for subscriber access.

Options *profile-name*—Unique name that identifies the tunnel profile. The profile can be referenced from within a domain map or by the RADIUS Tunnel-Group VSA [26-64].

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

tunnel-switch-profile (L2TP Tunnel Switching, Application)

Syntax tunnel-switch-profile *profile-name*;

Hierarchy Level [edit access domain map *domain-map-name*],
[edit services l2tp],
[edit services l2tp tunnel-group *group-name*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify a tunnel switch profile that determines whether packets in an L2TP session from a LAC are switched to another session that has a different destination LNS.

Options *profile-name*—Unique name that identifies the tunnel switch profile.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 135

tunnel-switch-profile (L2TP Tunnel Switching, Definition)

Syntax

```
tunnel-switch-profile profile-name {
  avp {
    bearer-type action;
    calling-number action;
    cisco-nas-port-info action;
  }
  tunnel-profile profile-name;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 13.2.

Description Define a tunnel switch profile for subscriber access; specify actions to take for L2TP AVPs in the switched packets and the profile for the tunnel to which the PPP traffic is switched.

The remaining statements are explained separately.

Options *profile-name*—Unique name that identifies the tunnel switch profile. The profile can be applied as a default or referenced from within a domain map, a tunnel group, or by the RADIUS Tunnel-Group VSA [26-64].

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring L2TP Tunnel Switching on page 135

 Copyright © 2016, Juniper Networks, Inc.

tx-address-change (L2TP LAC)

Syntax tx-address-change (accept | ignore | ignore-ip-address | ignore-udp-port);

Hierarchy Level [edit services l2tp tunnel]

Description Configure whether the LAC accepts or ignores requests from a peer to change the destination IP address, UDP port, or both.

Default The LAC accepts the IP address change from its peer.

Options

accept—Accept a change in IP address , UDP port, or both.

ignore—Ignore a change request for IP address or UDP port.

ignore-ip-address—Ignore a change request for IP address.

ignore-udp-port—Ignore a change request for UDP port.

Required Privilege Level

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

Related Documentation

- Configuring the LAC to Ignore Address and Port Changes Requested by the LNS on page 151

tx-connect-speed-method (L2TP LAC)

Syntax tx-connect-speed-method *method*;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 11.4.
Options **ancp**, **pppoe-ia-tag**, and **static** introduced in Junos OS Release 13.1.

Description Specify the method that determines the connection speed values sent from the LAC to the LNS in Incoming-Call-Connected (ICCN) messages. The transmit speed is conveyed in AVP 24 (Tx Connect Speed) and the receive speed is conveyed in AVP 38 (Rx Connect Speed). Both values are in bits per seconds (BPS).

Default static

Options *method*—Method used to derive the connection speed values.

- **ancp**—The speed is derived from the configured ANCP value for the underlying interface. You can change this rate after a subscriber has logged in, but those changes do not affect the actual rate used by the LNS. The **ancp** method gets the highest preference among the methods configured.
- **pppoe-ia-tags**—PPPoE IA tags sent from the DSLAM to the LAC. This speed value transmitted when a subscriber logs in and it cannot be subsequently changed. This value is used when the **ancp** value is not available. This speed does not apply to the subscribers that are already logged in; it is effective only for subscribers that log in after this setting has been saved. AVP 24 is the value of Actual-Data-Rate-Downstream (VSA 26-129). AVP 38 is the value of Actual-Data-Rate-Upstream (VSA 26-130), and is sent only when the VSA values differ.

  NOTE: **This rate does not affect the subscribers already logged in; however, new subscribers inherit the new rate.**

- **static**—The speed is derived from the recommended (advisory) shaping rate configured on the PPPoE logical interface underlying the subscriber interface. If the advisory speed is not configured on the underlying interface, then the actual port speed is used. The default method, when no other methods yield a value, is the **static** method or the advisory speed method. If the advisory speed is not configured, then the actual port speed is used. For ge and xe interfaces, the speed value is set to 10,000,000 and for ae interfaces, the speed value is set to 0 and sent in both AVP 24 and AVP 38.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2016, Juniper Networks, Inc.

Related Documentation

- Configuring the Method to Set the LAC Connection Speeds to the LNS on page 172

type (Tunnel Profile)

Syntax

type *tunnel-type*;

Hierarchy Level

[edit access tunnel-profile *profile-name* tunnel *tunnel-id*]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Specify the tunnel type (protocol).

Default

l2tp

Options

tunnel-type—Tunnel protocol type. The only value currently available is **l2tp**.

Required Privilege Level

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

Related Documentation

- Configuring a Tunnel Profile for Subscriber Access on page 162

unit (Dynamic PPPoE)

Syntax

```
unit logical-unit-number {
  keepalives interval seconds;
  no-keepalives;
  pppoe-options {
    underlying-interface interface-name;
    server;
  }
  ppp-options {
    authentication [ authentication-protocols ];
    chap {
      challenge-length minimum minimum-length maximum maximum-length;
    }
    pap;
  }
  family inet {
    unnumbered-address interface-name;
    address address;
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
      output {
        service-set service-set-name {
          service-filter filter-name;
        }
      }
    }
    filter {
      input filter-name {
        precedence precedence;
      }
      output filter-name {
        precedence precedence;
      }
    }
  }
  filter {
    input filter-name;
    output filter-name;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure a logical unit number for the dynamic PPPoE logical interface. You must configure a logical interface to be able to use the router.

 Copyright © 2016, Juniper Networks, Inc.

Options ***logical-unit-number***—Variable used to specify the unit number when the PPPoE logical interface is dynamically created. In the **unit *logical-unit-number*** statement for dynamic PPPoE logical interfaces, you must use the predefined variable **$junos-interface-unit** in place of ***logical-unit-number***. The **$junos-interface-unit** predefined variable is dynamically replaced with the unit number supplied by the router when the subscriber logs in.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring a Basic PPPoE Dynamic Profile*
- *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

user-group-profile

Syntax user-group-profile *profile-name*;

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply a configured PPP group profile to PPP users.

NOTE: If *<user-group-profile>* is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, will go down.

Options ***profile-name***—Name of a PPP group profile configured at the **[edit access group-profile *profile-name*]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- *Applying a Configured PPP Group Profile to a Tunnel*
- Configuring an L2TP Access Profile on the LNS on page 180

version (BFD)

Syntax

version (0 | 1 | automatic);

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp oam bfd-liveness-detection],
[edit logical-systems *logical-system-name* protocols ldp oam fec *address* bfd-liveness-detection],
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group *group-name* liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method bfd],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec *address* bfd-liveness-detection]

Release Information

Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the BFD protocol version to detect.

Options

0—Use BFD protocol version 0.

1—Use BFD protocol version 1.

automatic—Autodetect the BFD protocol version.

Default: automatic

Required Privilege Level

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

Related Documentation

- Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 83
- Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients on page 78
- *Configuring BFD for LDP LSPs*

Copyright © 2016, Juniper Networks, Inc.

weighted-load-balancing (L2TP LAC)

Syntax weighted-load-balancing;

Hierarchy Level [edit services l2tp]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify that the router chooses among multiple tunnels that share the same preference level by considering the maximum sessions configured per tunnel. The tunnel configured with the highest maximum number of sessions in the preference level has the highest weight. This tunnel is selected until the maximum number of sessions for the tunnel is reached. Then the router selects the tunnel with the next higher weight to establish connections until that tunnel's maximum session limit is reached, and so on.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Weighted Load Balancing for LAC Tunnel Sessions on page 166
- Configuring the L2TP LAC Tunnel Selection Parameters on page 165

wimax

Syntax wimax;

Hierarchy Level [edit services mobile-ip access-type],
[edit logical-systems *logical-system-name* services mobile-ip access-type],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* services mobile-ip access-type],
[edit routing-instances *routing-instance-name* services mobile-ip access-type]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable WiMAX features for Mobile IP home agent, including the ability to process, send, and receive WiMAX Vendor Specific Attributes (VSAs).

NOTE: Although this statement is available in the CLI for both default and nondefault router contexts, the commit operation is disallowed when you configure the statement in a nondefault router context.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Access Type for Mobile IP on page 232

virtual-network

Syntax

```
virtual-network {
  home-agent-address ip-address {
    registration-lifetime seconds;
    revocation-required;
    timestamp-tolerance seconds;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* services mobile-ip home-agent],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit routing-instances *routing-instances-name* services mobile-ip home-agent],
[edit services mobile-ip home-agent]

Release Information

Statement introduced in Junos OS Release 9.3.
Support at the **[edit logical-systems *logical-system-name* ...]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* ...]**, and **[edit routing-instances *routing-instances-name* ...]** hierarchy levels introduced in Junos OS Release 9.5.

Description

Define the virtual network for the Mobile IP home agent. Only one virtual network is supported.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Mobile IP on page 229
- Configuring the Mobile IP Home Agent on page 230

Copyright © 2016, Juniper Networks, Inc.

vlan-id (Dynamic Profiles)

Syntax vlan-id (*number* | none);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For VLAN demux, Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options *number*—A valid VLAN identifier. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN identifier.

none—Enable the use of untagged pseudo-wire frames on dynamic interfaces.

- For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.
- For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.
- VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Static Subscriber Interfaces Using VLAN Demux Interfaces*
- *Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles*

vlan-tags

Syntax vlan-tags outer [*tpid*].*vlan-id* [inner [*tpid*].*vlan-id*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

NOTE: The **inner-range *vid1–vid2*** option is supported on MX Series routers with IQE PICs only.

Options

inner [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN ID.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-stacked-vlan-id** predefined variable.

Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dual VLAN Tags*
- *stacked-vlan-tagging*

 Copyright © 2016, Juniper Networks, Inc.

CHAPTER 35

Operational Commands

Copyright © 2016, Juniper Networks, Inc.

- show system subscriber-management summary
- test services l2tp tunnel

Copyright © 2016, Juniper Networks, Inc.

clear mobile-ip binding

Syntax **clear mobile-ip binding** (all | ip-address *ip-address* | nai *nai-string*)
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 9.3.

Description Clear the Mobile IP binding.

Options **all**—Clear all Mobile IP bindings.

ip-address ***ip-address***—Clear the Mobile IP bindings for the specified IP home address (HoA).

nai ***nai-string***—Clear the Mobile IP bindings for the specified network access identifier.

logical-system ***logical-system-name***—(Optional) Clear the Mobile IP bindings for the specified logical system.

routing-instance ***routing-instance-name***—(Optional) Clear the Mobile IP bindings for the specified routing instance.

Required Privilege Level clear

Related Documentation
- show mobile-ip home-agent bindings on page 474

List of Sample Output clear mobile-ip binding on page 453

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

Sample Output

clear mobile-ip binding

```
user@host> clear mobile-ip binding all
```

clear services l2tp destination

Syntax

clear services l2tp destination
<all | local-gateway *gateway-address* | peer-gateway *gateway-address*>

Release Information

Command introduced in Junos OS Release 10.4.
Statistics option introduced in Junos OS Release 13.1

Description

Clear all Layer 2 Tunneling Protocol (L2TP) destinations and all tunnels and sessions that belong to the destinations. This command is available only for LAC on MX Series routers.

Options

all—Close all L2TP destinations.

local-gateway *gateway-address*—Clear only the L2TP destinations and all tunnels and sessions associated with the specified local gateway address.

peer-gateway *gateway-address*—Clear only the L2TP destinations and all tunnels and sessions associated with the peer gateway with the specified address.

Required Privilege Level

clear

Related Documentation

- show services l2tp destination on page 507

List of Sample Output

clear services l2tp destination all on page 454

Output Fields

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

Sample Output

clear services l2tp destination all

```
user@host> clear services l2tp destination all

  Destination 2 closed
```


Copyright © 2016, Juniper Networks, Inc.

clear services l2tp session

Syntax

clear services l2tp session (all | interface *interface-name* | local-gateway *gateway-address* | local-gateway-name *gateway-name* | local-session-id *session-id* | local-tunnel-id *tunnel-id* | peer-gateway *gateway-address* | peer-gateway-name *gateway-name* | tunnel-group *group-name* | user *username*)

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M10i and M7i routers only) Clear Layer 2 Tunneling Protocol (L2TP) sessions on LNS.

(MX Series routers only) Clear L2TP sessions on LAC and LNS.

Options

all—Close all L2TP sessions.

interface *interface-name*—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- **si-*fpc/pic/port***—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- **sp-*fpc/pic/port***—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*—Clear only the L2TP sessions associated with the specified local gateway address.

local-gateway-name *gateway-name*—Clear only the L2TP sessions associated with the specified local gateway name.

local-session-id *session-id* —Clear only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id *tunnel-id*—Clear only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear only the L2TP sessions associated with the peer gateway with the specified name.

tunnel-group *group-name*—Clear only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

user *username* —(M Series routers only) Clear only the L2TP sessions for the specified username.

Required Privilege Level

clear

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- clear services l2tp session statistics on page 457
- show services l2tp session on page 512

List of Sample Output

clear services l2tp session on page 456
clear services l2tp session interface on page 456

Output Fields

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

Sample Output

clear services l2tp session

```
user@host> clear services l2tp session 31694

  Session 31694 closed
```

Sample Output

clear services l2tp session interface

```
user@host> show services l2tp session Tunnel local ID: 17185
Local   Remote   State                Interface         Interface
  ID      ID                            unit              Name
  5117    1        Established          1073741828        si-2/0/0
  34915   2        Established          1073741829        si-2/1/0
  6454    3        Established          1073741830        si-2/0/0
  46142   4        Established          1073741831        si-2/1/0

user@host> clear services l2tp session interface si-2/0/0
  Session  5117 closed
  Session  6454 closed

user@host> show services l2tp session Tunnel local ID: 17185
  Local   Remote   State                Interface         Interface
  ID      ID                            unit              Name
  34915   2        Established          1073741829        si-2/1/0
  46142   4        Established          1073741831        si-2/1/0
```


Copyright © 2016, Juniper Networks, Inc.

clear services l2tp session statistics

Syntax

clear services l2tp session statistics (all | interface *interface-name* | local-gateway *gateway-address* | local-gateway-name *gateway-name* | local-session-id *session-id* | local-tunnel-id *tunnel-id* | peer-gateway *gateway-address* | peer-gateway-name *gateway-name* | tunnel-group *group-name* | user *username*)

Release Information

Command introduced before Junos OS Release 7.4.
Support for MX Series routers added in Junos OS Release 10.4.

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) sessions.

Options

all—Clear statistics for all L2TP sessions.

interface ***interface-name***—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- **si-*fpc/pic/port***—MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- **sp-*fpc/pic/port***—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway ***gateway-address***—Clear statistics for only the L2TP sessions associated with the local gateway with the specified address.

local-gateway-name ***gateway-name***—Clear statistics for only the L2TP sessions associated with the local gateway with the specified name.

local-session-id ***session-id*** —Clear statistics for only the L2TP sessions with this identifier for the local endpoint of the L2TP session.

local-tunnel-id ***tunnel-id***—Clear statistics for only the L2TP sessions associated with the specified local tunnel identifier.

peer-gateway ***gateway-address***—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified address.

peer-gateway-name ***gateway-name***—Clear statistics for only the L2TP sessions associated with the peer gateway with the specified name.

tunnel-group ***group-name***—Clear statistics for only the L2TP sessions associated with the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

user ***username*** **<statistics>**—Clear statistics for only the L2TP sessions for the specified username. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- clear services l2tp session on page 455
- show services l2tp session on page 512

List of Sample Output clear services l2tp session statistics all on page 458

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

Sample Output

clear services l2tp session statistics all

```
user@host> clear services l2tp session statistics all
  Session 26497 statistics cleared
```

 Copyright © 2016, Juniper Networks, Inc.

clear services l2tp tunnel

Syntax

clear services l2tp tunnel (all | interface sp-*fpc/pic/port* | local-gateway *gateway-address* | local-gateway-name *gateway-name* | local-tunnel-id *tunnel-id* | peer-gateway *gateway-address* | peer-gateway-name *gateway-name* | tunnel-group *group-name*)

Release Information

Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Clear Layer 2 Tunneling Protocol (L2TP) tunnels.

Options

all—Clear all L2TP tunnels.

sp-*fpc/pic/port*—(Optional) Clear only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*—Clear only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*—Clear only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*—Clear only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear only the L2TP tunnels associated with the peer gateway with the specified name.

tunnel-group *group-name*—Clear only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- clear services l2tp tunnel statistics on page 461
- show services l2tp tunnel on page 525

List of Sample Output

clear services l2tp tunnel on page 460

Output Fields

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

Sample Output

clear services l2tp tunnel

```
user@host> clear services l2tp tunnel 17185

  Tunnel 17185 closed
```


Copyright © 2016, Juniper Networks, Inc.

clear services l2tp tunnel statistics

Syntax

clear services l2tp tunnel statistics (all | interface sp-*fpc/pic/port* | local-gateway *gateway-address* | local-gateway-name *gateway-name* | local-tunnel-id *tunnel-id* | peer-gateway *gateway-address* | peer-gateway-name *gateway-name* | tunnel-group *group-name*)

Release Information

Command introduced before Junos OS Release 7.4.
Support for MX Series routers added in Junos OS Release 10.4.

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC only.) Clear statistics for Layer 2 Tunneling Protocol (L2TP) tunnels.

Options

all—Clear statistics for all L2TP tunnels.

interface sp-*fpc/pic/port*—Clear statistics for only the L2TP tunnels using the specified adaptive services interface. This option is not available for L2TP LAC on MX Series routers.

local-gateway *gateway-address*—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified address.

local-gateway-name *gateway-name*—Clear statistics for only the L2TP tunnels associated with the local gateway with the specified name.

local-tunnel-id *tunnel-id*—Clear statistics for only the L2TP tunnels that have the specified local tunnel identifier.

peer-gateway *gateway-address*—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified address.

peer-gateway-name *gateway-name*—Clear statistics for only the L2TP tunnels associated with the peer gateway with the specified name.

tunnel-group *group-name*—Clear statistics for only the L2TP tunnels in the specified tunnel group. This option is not available for L2TP LAC on MX Series routers.

Required Privilege Level

clear

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- clear services l2tp tunnel on page 459
- show services l2tp tunnel on page 525

List of Sample Output

clear services l2tp tunnel statistics all on page 462

Output Fields

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

Sample Output

clear services l2tp tunnel statistics all

```
user@host> clear services l2tp tunnel statistics all
  Tunnel  9933 statistics cleared
```


Copyright © 2016, Juniper Networks, Inc.

restart

List of Syntax

Syntax

```
restart
<adaptive-services |ancpd-service | application-identification |audit-process |
  auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control |
  class-of-service |clksyncd-service |database-replication|datapath-trace-service
  |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | iccp-service |idp-policy | immediately
  |interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service |link-management
  |local-policy-decision-function |mac-validation |mib-process | mobile-ip | mountd-service
  |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |
  packet-triggered-subscribers |peer-selection-service |pgcp-service | pgm |
  pic-services-logging | pki-service |ppp | ppp-service |pppoe |
  protected-system-domain-service | redundancy-interface-process | remote-operations |
  root-system-domain-service | routing <logical-system logical-system-name>  | sampling
  | sbc-configuration-process | sdk-service |service-deployment | services | services pgcp
  gateway gateway-name | snmp |soft |static-subscribers |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd |usb-control|
  vrrp |web-management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

```
restart
<adaptive-services |audit-process | auto-configuration | autoinstallation |chassis-control |
  class-of-service |clksyncd-service |database-replication| dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | immediately |interface-control |
  ipsec-key-management | l2-learning | lacp |link-management |mib-process | mobile-ip |
  mountd-service |mpls-traceroute |mspd | named-service | nfsd-service | pgm | pki-service
  |ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
  sdk-service |secure-neighbor-discovery | service-deployment | services | snmp |soft
  |statistics-service| subscriber-management | subscriber-management-helper | tunnel-oamd
  | vrrp>
```

Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
```

| lldpd-service | mib-process | mountd-service | multicast-snooping | pgm | redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-deployment | sflow-service | snmp | vrrp | web-management>

Syntax (Routing Matrix)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp>
<all | all-lcc | lcc *number*>
<gracefully | immediately | soft>

Syntax (J Series Routing Platform)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dialer-services | dlsw | event-processing | firewall | interface-control | ipsec-key-management | isdn-signaling | l2-learning | l2tp-service | mib-process | network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>

Syntax (TX Matrix Routers)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management | mib-process |pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | lcc *number* | scc>
<gracefully | immediately | soft>

Syntax (TX Matrix Plus Routers)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | all-sfc | lcc *number* | sfc *number*>
<gracefully | immediately | soft>

Syntax (MX Series Routers)

restart
<adaptive-services | ancpd-service |application-identification|audit-process | auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control | class-of-service |clksyncd-service |database-replication| datapath-trace-service |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | ethernet-connectivity-fault-management |ethernet-link-fault-management |event-processing | firewall | general-authentication-service | gracefully |iccp-service |idp-policy |immediately |interface-control | ipsec-key-management | kernel-replication | l2-learning |l2cpd-service |l2tp-service | l2tp-universal-edge | lacp | license-service |link-management |local-policy-decision-function |mac-validation |mib-process | mobile-ip |mountd-service |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |

Copyright © 2016, Juniper Networks, Inc.

packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service |ppp | ppp-service |pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system *logical-system-name*> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway *gateway-name* |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member *member-id*>

Syntax (J Series Routers)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsw | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

Syntax (QFX Series)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |*logical-system-name*> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsw** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.

- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description Restart a Junos OS process.

CAUTION: **Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.**

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

 Copyright © 2016, Juniper Networks, Inc.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dlsw—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

Copyright © 2016, Juniper Networks, Inc.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service— (M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

Copyright © 2016, Juniper Networks, Inc.

lacp—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc ***number***—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member ***member-id***—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace ***member-id*** with a value of **0** or **1**.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

 Copyright © 2016, Juniper Networks, Inc.

ppp-service—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level

reset

Related Documentation

- *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output

restart interfaces on page 473

Copyright © 2016, Juniper Networks, Inc.

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

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

show mobile-ip home-agent bindings

Syntax

show mobile-ip home-agent bindings
<ip-address *ip-address* | nai *nai-string* | summary>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.3.

Description

Display information about Mobile IP home agent bindings.

Options

ip-address *ip-address*—(Optional) Display information for the specified Mobile IP home address.

logical-system *logical-system-name*—(Optional) Display information for the specified logical system.

nai *nai-string*—(Optional) Display information for the specified Mobile IP network access identifier.

routing-instance *routing-instance-name*—(Optional) Display information for the specified routing instance.

summary—(Optional) Display only summary (total bindings) information.

Required Privilege Level

view

Related Documentation

- clear mobile-ip binding on page 453
- show mobile-ip home-agent overview on page 477
- show mobile-ip home-agent traffic on page 479
- show mobile-ip home-agent virtual-network on page 482
- show mobile-ip wimax release on page 484

List of Sample Output

show mobile-ip home-agent bindings on page 475
show mobile-ip home-agent bindings ip-address on page 475
show mobile-ip home-agent bindings nai on page 476
show mobile-ip home-agent bindings summary on page 476

Output Fields

Table 19 on page 474 lists the output fields for the **show mobile-ip home-agent bindings** command. Output fields are listed in the approximate order in which they appear.

Table 19: show mobile-ip home-agent bindings Output Fields

Field Name	Field Description
Home Address	Home address of the mobile node.
NAI	Network access identifier of the mobile node.

Copyright © 2016, Juniper Networks, Inc.

Table 19: show mobile-ip home-agent bindings Output Fields *(continued)*

Field Name	Field Description
Home agent	Home agent address of the mobile node.
Care-of-address	Care of address used by the mobile node.
Lifetime Granted	Lifetime granted for the mobile node.
Lifetime Remaining	Remaining lifetime for the mobile node.
Tunnel Type	Type of tunnel requested by the mobile node.
Tunnel ID	Tunnel ID the mobile node is using.
Tunnel Source	Tunnel source address the mobile node is using.
Tunnel Destination	Tunnel destination address the mobile node is using.
Identification	Identification value received from the mobile node.
Revocation Support	Whether registration revocation is supported for this binding.
Notify MN	Whether mobile node notification has been negotiated.
Total Bindings	Total number of Mobile IP home agent bindings.

Sample Output

show mobile-ip home-agent bindings

```
user@host> show mobile-ip home-agent bindings
Home address  NAI              Home agent   Care-of-address
10.1.1.3      abcde@def.com    10.1.1.1     50.50.50.1
30.1.1.3          -            55.55.55.1   50.50.50.1
20.1.1.3      def@def.com      20.1.1.1     60.50.50.1
```

show mobile-ip home-agent bindings ip-address

```
user@host> show mobile-ip home-agent bindings ip-address 10.1.1.3
Home address       : 10.1.1.3
NAI                : abcde@def.com
Home agent         : 10.1.1.1
Care-of-address    : 50.50.50.1
Lifetime Granted   : 180
Lifetime Remaining : 20
Tunnel Type        : IP-IP
Tunnel  ID         : 10
Tunnel Source      : 10.1.1.1
Tunnel Destination : 50.50.50.1
Identification     : ABCD1234.4321ABCD
Revocation Support : Enabled
Notify MN of Revocation : Enabled
```

show mobile-ip home-agent bindings nai

```
user@host> show mobile-ip home-agent bindings nai abcde@def.com
Home address       : 10.1.1.3
NAI                : abcde@def.com
Home agent         : 10.1.1.1
Care-of-address    : 50.50.50.1
Lifetime Granted   : 180
Lifetime Remaining : 20
Tunnel Type        : IP-IP
Tunnel  ID         : 10
Tunnel Source      : 10.1.1.1
Tunnel Destination : 50.50.50.1
Identification     : ABCD1234.4321ABCD
Revocation Support : Enabled
Notify MN          : Enabled
```

show mobile-ip home-agent bindings summary

```
user@host> show mobile-ip home-agent bindings summary
Total bindings : 3
```


Copyright © 2016, Juniper Networks, Inc.

show mobile-ip home-agent overview

Syntax

show mobile-ip home-agent overview
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.3.

Description

Display overview information for Mobile IP home agent.

Options

logical-system *logical-system-name*—(Optional) Display information for the specified logical system.

routing-instance *routing-instance-name*—(Optional) Display information for the specified routing instance.

Required Privilege Level

view

Related Documentation

- show mobile-ip home-agent bindings on page 474
- show mobile-ip home-agent traffic on page 479
- show mobile-ip home-agent virtual-network on page 482
- show mobile-ip wimax release on page 484

List of Sample Output

show mobile-ip home-agent overview on page 477

Output Fields

Table 20 on page 477 lists the output fields for the **show mobile-ip home-agent overview** command. Output fields are listed in the approximate order in which they appear.

Table 20: show mobile-ip home-agent overview Output Fields

Field Name	Field Description
Status	Total number of registration requests received.
Service Enabled on	Total number of registration requests forwarded.
Home Agents	Total number of registration requests denied.
Authentication	Total number of registration replies sent.

Sample Output

show mobile-ip home-agent overview

```
user@host> show mobile-ip home-agent overview
Status             : Active
Service Enabled on : ge-0/0/3.0, ge-0/0/2.0
Home agents        : 10.1.1.1, 20.1.1.1, 55.55.55.1
Authentication     : AAA
```

Copyright © 2016, Juniper Networks, Inc.

show mobile-ip home-agent traffic

Syntax

show mobile-ip home-agent traffic
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.3.

Description

Display information about Mobile IP home agent protocol statistics.

Options

logical-system *logical-system-name*—(Optional) Display information for the specified logical system.

routing-instance *routing-instance-name*—(Optional) Display information for the specified routing instance.

Required Privilege Level

view

Related Documentation

- show mobile-ip home-agent bindings on page 474
- show mobile-ip home-agent overview on page 477
- show mobile-ip home-agent virtual-network on page 482
- show mobile-ip wimax release on page 484

List of Sample Output

show mobile-ip home-agent traffic on page 480

Output Fields

Table 21 on page 479 lists the output fields for the **show mobile-ip home-agent traffic** command. Output fields are listed in the approximate order in which they appear.

Table 21: show mobile-ip home-agent traffic Output Fields

Field Name	Field Description
Registration request received	Total number of registration requests received.
Registration request forwarded	Total number of registration requests forwarded.
Registration request denied	Total number of registration requests denied.
Registration replies sent	Total number of registration replies sent.
Registration Errors unspecified	Total number of registration requests denied by the home agent for reasons unspecified.

Table 21: show mobile-ip home-agent traffic Output Fields *(continued)*

Field Name	Field Description
Registration Errors Administrative prohibited	Total number of registration requests denied by home agent as “administrative prohibited.”
Registration Errors Insufficient Resource	Total number of registration requests denied by the home agent for insufficient resources.
Registration Errors Bad request form	Total number of registration requests denied by the home agent due to a bad request form.
Registration Errors Too many Bindings	Total number of registration requests denied by the home agent for having too many bindings.
Registration Errors Unknown HA	Total number of registration requests denied by the home agent for having an unknown home agent.
Registration Errors ID mismatch	Total number of registration requests denied by the home agent for having a mismatched ID.
Registration Errors Authentication failed MN	Total number of registration requests denied by the home agent because the mobile node failed authentication.
Registration Errors Authentication failed FA	Total number of registration requests denied by the home agent because the foreign agent failed authentication.

Sample Output

show mobile-ip home-agent traffic

```
user@host> show mobile-ip home-agent traffic
Registration Request
    Received  : 10
    Forwarded : 5
    Denied    : 5
Registration Replies
    Sent      : 5
Registration Errors
    Unspecified    : 0
    Administrative prohibited  : 0
    Insufficient Resource      : 0
    Bad request form           : 0
    Too many Bindings          : 0
    Unknown HA                 : 0
    ID mismatch                : 0
    Unavailable Reverse tunnel : 0
    Unavailable Encapsulation  : 0
    Reverse Tunnel Mandatory   : 0
    Authentication failed MN   : 0
    Authentication failed FA   : 0
```


Copyright © 2016, Juniper Networks, Inc.

show mobile-ip home-agent virtual-network

Syntax

show mobile-ip home-agent virtual-network
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.3.

Description

Display information about Mobile IP home agent virtual networks.

Options

logical-system *logical-system-name*—(Optional) Display information for the specified logical system.

routing-instance *routing-instance-name*—(Optional) Display information for the specified routing instance.

Required Privilege Level

view

Related Documentation

- show mobile-ip home-agent bindings on page 474
- show mobile-ip home-agent overview on page 477
- show mobile-ip home-agent traffic on page 479
- show mobile-ip wimax release on page 484

List of Sample Output

show mobile-ip home-agent virtual-network on page 483

Output Fields

Table 22 on page 482 lists the output fields for the **show mobile-ip home-agent virtual-network** command. Output fields are listed in the approximate order in which they appear.

Table 22: show mobile-ip home-agent virtual-network Output Fields

Field Name	Field Description
Home agent address	Home agent address of the mobile node.
Registration Lifetime	Maximum registration lifetime that home agent allows.
Time Tolerance	Number of seconds the time stamp may differ.
Address Pool	Address pool configured.
Total MNs	Current number of mobile nodes that the home agent is serving.
Home address	Home address of the mobile node.
NAI	Network access identifier of the mobile node.

Copyright © 2016, Juniper Networks, Inc.

Table 22: show mobile-ip home-agent virtual-network Output Fields *(continued)*

Field Name	Field Description
Care-of-address	Care of address used by the mobile node.
RegLifetime Granted	Lifetime granted for the mobile node.
RegLifetime Remaining	Remaining lifetime for the mobile node.

Sample Output

show mobile-ip home-agent virtual-network

```
user@host> show mobile-ip home-agent virtual-network
Home Agent Address    : 55.55.55.55
Registration Lifetime : 1800
Time Tolerance        : 120
Address Pool          : 10.1.1.10 - 10.1.1.50
Total MN's            : 2

MN's :
 Home address         : 60.60.60.1
 NAI                  : abcde@def.com
 Care-of-address      : 50.50.50.1
 Reglifetime granted  : 120
 Reglifetime remaining: 100

 Home address         : 70.70.70.1
 NAI                  : def@def.com
 Care-of-address      : 80.80.80.1
 Reglifetime granted  : 120
 Reglifetime remaining: 100
```

Copyright © 2016, Juniper Networks, Inc.

show mobile-ip wimax release

Syntax

show mobile-ip wimax release
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.5.

Description

Display the WiMAX Forum Network Architecture release that is supported by the current Mobile IP implementation.

Options

logical-system *logical-system-name*—(Optional) Display information for the specified logical system.

routing-instance *routing-instance-name*—(Optional) Display information for the specified routing instance.

Required Privilege Level

view

Related Documentation

- show mobile-ip home-agent bindings on page 474
- show mobile-ip home-agent overview on page 477
- show mobile-ip home-agent traffic on page 479
- show mobile-ip home-agent virtual-network on page 482

List of Sample Output

show mobile-ip wimax release on page 484

Output Fields

Table 23 on page 484 lists the output fields for the **show mobile-ip wimax release** command. Output fields are listed in the approximate order in which they appear.

Table 23: show mobile-ip wimax release Output Fields

Field Name	Field Description
Release	WiMAX Forum Network Architecture release number.
Version	WiMAX Forum Network Architecture version number.

Sample Output

show mobile-ip wimax release

```
user@host> show mobile-ip wimax release
Release 1, Version 1.2
```


Copyright © 2016, Juniper Networks, Inc.

show ppp interface

Syntax show ppp interface *interface-name*
<extensive| terse>

Release Information Command introduced in Junos OS Release 7.5.

Description Display information about PPP interfaces.

Options ***interface-name***—Name of a logical interface.

extensive | terse—(Optional) Display the specified level of output.

Required Privilege Level view

List of Sample Output
show ppp interface on page 493
show ppp interface extensive on page 493
show ppp interface terse on page 493

Output Fields Table 24 on page 485 lists the output fields for the **show ppp interface** command. Output fields are listed in the approximate order in which they appear.

Table 24: show ppp interface Output Fields

Field Name	Field Description	Level of Output
Session	Name of the logical interface on which the session is running.	All levels
Type	Session type: PPP.	All levels
Phase	PPP process phase: **Authenticate**, **Pending**, **Establish**, **LCP**, **Network**, **Disabled**, and **Tunneled**.	All levels
Session flags	Special conditions present in the session: **Bundled**, **TCC**, **No-keepalives**, **Looped**, **Monitored**, and **NCP-only**.	All levels
***protocol* State**	Protocol state information. See specific protocol state fields for information.	None specified
AUTHENTICATION	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the **Authentication** field description for further information.	None specified

Copyright © 2016, Juniper Networks, Inc.

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Keepalive settings	Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS based PPP sessions are supported only on service interfaces (si). • **Interval**—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the **interval** and **Down-count** values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • **Up-count**—The number of keepalive packets a destination must receive to change a link's status from down to up. • **Down-count**—The number of keepalive packets a destination must fail to receive before the network takes down a link.	**extensive**
RE Keepalive statistics	Keepalive statistics for the packets handled by the Routing Engine. • **LCP echo req Tx**—LCP echo requests sent from the Routing Engine. • **LCP echo req Rx**—LCP echo requests received at the Routing Engine. • **LCP echo rep Tx**—LCP echo responses sent from the Routing Engine. • **LCP echo rep Rx**—LCP echo responses received at the Routing Engine. • **LCP echo req timeout**—Number of keepalive packets where the keepalive aging timer has expired. • **LCP Rx echo req Magic Num Failures**—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • **LCP Rx echo rep Magic Num Failures**—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

LCP	LCP information:
	• **State**—LCP protocol state (all platforms except M120 and M320 routers):
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—LCP protocol state (M120 and M320 routers):
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—LCP state start time.
	• **Last completed**—LCP state completion time.
		extensive

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	• **Negotiated options**:
	  • **ACFC**—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields.
	  • **Asynchronous map**—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link.
	  • **Authentication protocol**—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required.
	  • **Authentication algorithm**—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported.
	  • **Endpoint discriminator class**—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link.
	  • **Magic number**—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated.
	  • **MRU**—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets.
	  • **MRRU**—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets.
	  • **Multilink header suspendable classes**—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given.
	  • **Multilink header format classes**—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number.
	  • **PFC**—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field.
	  • **short sequence**—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers.	

Copyright © 2016, Juniper Networks, Inc.

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Authentication	CHAP or PAP authentication state information. For CHAP authentication: • **Chap-ans-rcvd**—Packet was sent from the peer, indicating that the peer received the **Chap-resp-sent** packet. • **Chap-ans-sent**—Packet was sent from the authenticator, indicating that the authenticator received the peer's **Chap-resp-rcvd** packet. • **Chap-chal-rcvd**—Challenge packet has been received by the peer. • **Chap-chal-sent**—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • **Chap-resp-rcvd**—CHAP response packet has been received by the authenticator. • **Chap-resp-sent**—CHAP response packet has been sent to the authenticator. • **Closed**—Link is not available for authentication. • **Failure**—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • **Success**—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. For PAP authentication: • **Pap-resp-sent**—PAP response sent to peer (ACK/NACK)t. • **Pap-req-rcvd**—PAP request packet received from peer. • **Pap-resp-rcvd**—PAP response received from the peer (ACK/NACK). • **Pap-req-sent**—PAP request packet sent to the peer. • **Closed**—Link is not available for authentication. • **Failure**—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • **Success**—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.	None specified

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

IPCP	Internet Protocol Control Protocol (IPCP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—IPCP state start time.
	• **Last completed**—IPCP state authentication completion time.
	• **Negotiated options:**
	  • **compression protocol**—Negotiate the use of a specific compression protocol. By default, compression is not enabled.
	  • **local address**—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.
	  • **primary DNS server**—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link.
	  • **primary WINS server**—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link.
	  • **remote address**—IP address of the remote end of the link in dotted quad notation.
	  • **secondary DNS server**—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link.
	  • **secondary WINS server**—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link.
	• **Negotiation mode**—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: **Active** or **Passive**
		extensive

Copyright © 2016, Juniper Networks, Inc.

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

IPV6CP	Internet Protocol version 6 Control Protocol (IPv6CP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—IPV6CP state start time.
	• **Last completed**—IPV6CP state authentication completion time.
	• **Negotiated options**:
	  • **local interface identifier**—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.
	  • **remote interface identifier**—IP address of the remote end of the link in dotted quad notation.
	• **Negotiation mode**—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: **Active** or **Passive**
		extensive

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

OSINLCP State	OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):
	• **State:**
	  • **Ack-rcvd**—Configure-Request has been sent and Configure-Ack has been received.
	  • **Ack-sent**—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—Attempt has been made to configure the connection.
	• **Last started**—OSINLCP state start time.
	• **Last completed**—OSINCLP state completion time.
		extensive

TAGCP	TAGCP information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—TAGCP state start time.
	• **Last completed**—TAGCP state authentication completion time.
		extensive
		none

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show ppp interface

```
user@host> show ppp interface si-1/3/0.0
Session si-1/3/0.0, Type: PPP, Phase: Authenticate
  Session flags: Monitored
    LCP State: Opened
    AUTHENTICATION: CHAP State: Chap-resp-sent, Chap-ans-sent
    IPCP State: Closed, OSINLCP State: Closed
```

show ppp interface extensive

```
user@host> show ppp interface si-0/0/3.0 extensive

    Session si-0/0/3.0, Type: PPP, Phase: Network
        Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
RE Keepalive statistics:
  LCP echo req Tx        :  657 (last sent 00:50:10 ago)
  LCP echo req Rx        :    0 (last seen: never)
  LCP echo rep Tx        :    0
  LCP echo rep Rx        :  657
  LCP echo req timout    :    0
  LCP Rx echo req Magic Num Failures :    0
  LCP Rx echo rep Magic Num Failures :    0
LCP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Authentication protocol: PAP, Magic number: 2341124815, MRU: 4470
Authentication: PAP
  State: Success
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
IPCP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Local address: 10.10.10.1, Remote address: 10.10.10.2
  Negotiation mode: Active
IPV6CP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Local interface identifier: 2a0:a522:64:d319, Remote interface identifier: 0:0:0:c
  Negotiation mode: Passive
```

show ppp interface terse

```
user@host> show ppp interface si-1/3/0 terse
Session name    Session type   Session phase      Session flags
si-1/3/0.0      PPP            Authenticate       Monitored
```

show ppp statistics

Syntax

show ppp statistics
<detail>
<memory>
<recovery>

Release Information

Command introduced in Junos OS Release 7.5.

Description

Display PPP interface statistics information.

Options

detail—(Optional) Display the detailed statistics.

memory—(Optional) Display PPP process memory statistics.

recovery—(Optional) Display recovery state of PPP after a GRES or restart. It is safe to force another GRES or restart only when the recovery state indicates the recovery is done.

NOTE: When you issue this command option during the recovery process, the command may time out or fail silently rather than display output. Recovery is not complete until the command displays **Recovery state: recovery done.**

Required Privilege Level

view

List of Sample Output

show ppp statistics on page 498
show ppp statistics detail on page 498
show ppp statistics recovery (Safe to Restart) on page 499
show ppp statistics recovery (Unsafe to Restart) on page 499

Output Fields

Table 25 on page 494 lists the output fields for the **show ppp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 25: show ppp statistics Output Fields

Field Name	Field Description	Level of Output
Total sessions	Number of PPP sessions on an interface.	none **detail**
Sessions in disabled phase	Number of PPP sessions disabled. Number of sessions where the link is either administratively or physically down. Once the PPP process learns from the kernel that Layer 2 is ready to send and receive traffic, it will do a phase transition from disabled to established. When LCP and NCP transitions through states, links transition to the establish phase when terminate packets are exchanged or some other failure, such as authentication or expiration of a timer occurs.	none **detail**

Copyright © 2016, Juniper Networks, Inc.

Table 25: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Sessions in establish phase	Number of PPP sessions in establish phase. In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure and test the data link.	none **detail**
Sessions in authenticate phase	Number of PPP sessions in authenticate phase. Each end of the PPP link must first send LCP packets to configure the data link during the link establishment phase. After the link has been established, PPP provides for an optional authentication phase before proceeding to the Network-Layer Protocol (NLP) phase.	none **detail**
Sessions in network phase	Number of PPP sessions in the network phase. After a link has been established and optional facilities have been negotiated as needed by the LCP, PPP must send Network Control Protocol (NCP) packets to choose and configure one or more network-layer protocols, such as IP, IPX, or AppleTalk. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.	none **detail**
Bundles in pending phase	Number of unique bundles to which PPP links are referring.	none **detail**
Type	Type of structure for which memory is allocated. • **Queued rtsock msgs**—Queued route socket messages. When a PPP process is unable to send a route socket message to the kernel (typically because of congestion of the route socket interface), the message is queued for deferred processing. • **PPP session**—Active PPP session. Stores all the information for a PPP session, such as authentication, sequence number, LCP session, and NCP session information. • **Interface address**—Interface address associated with a PPP connection. Stores the information about the interface address that PPP obtains from the kernel. • **Destination profile**—Stores the destination profile information associated with an interface address. • **ML link settings**—Stores information about an MLPPP link, such as the bundle name and compressed real-time transport protocol (CRTP) settings. • **IPCP blocked address**—When addresses are blocked in an address pool (for example, when the interface address is within the range of an address pool, it will be implicitly blocked), this structure is used to store the address in the pool. • **PPP session trace**—A PPP session trace is allocated for record keeping for each session listed at the [**set protocols ppp monitor-session**] hierarchy level. • **IFL redundancy state**—Stores redundancy state information needed for high availability (HA) operation. • **Protocol family**—Stores the information about the protocol family that PPP obtains from the kernel.	**detail**

Table 25: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type (continued)	• **ML bundle settings**—Multilink bundle settings. Stores the context information for a MLPPP bundle. • **PPP LCP session**—PPP Link Control Protocol session, used for establishing, configuring, and testing the data-link connection. Stores the information for an LCP session, such as negotiated options, current state, and statistics. • **PPP NCP session**—PPP Network Control Protocol (NCP) phase in the PPP link connection process. Stores the information for an NCP session, such as negotiated options, current state, address family, and statistics. • **Physical interface**—Stores the information about the physical interface that PPP obtains from the kernel. • **Access profile**—Stores the information found at the [**edit access profile**] hierarchy level for each profile. • **ML wait entry**—Created when there are MLPPP links joining a bundle. before its addition to the PPP process. Links are saved here, and when the bundle is added, are properly assigned to the bundle. • **Group profile**—Stores information set in the PPP stanza of a group profile, such as the primary and secondary Domain Name System (DNS), primary and secondary NDNS, and address pool name. • **Profile client**—Stores the per-client information of the access profile (information obtained from the [**set access profile** ***name*** **client** ***client-name***] hierarchy level. • **PPP Auth session**—PPP authentication session. Stores all the session-specific authentication protocol parameters. • **Logical interface**—Stores the information about the logical interface that PPP obtains from the kernel. • **Non-tagged**—Generic catch-all for allocations not of a particular structure type.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 25: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type	If you specify the **memory** keyword, the following memory statistics are displayed for Ethernet interfaces on M120 and M320 routers. • **authenticate**—Stores information common to all PPP authentication protocols. • **linkInterface**—Stores information about PPP link interfaces. • **pap**—Stores information about PPP PAP authentication protocol. Includes authenticator and authenticate state machines. • **lcp**—PPP Link Control Protocol session. Used for establishing, configuring and testing the data-link connection. Stores information for LCP session, such as negotiated options, state, and statistics. • **chap**—Stores information about PPP CHAP authentication protocol. Includes authenticator and authenticate state machines. • **eapBuffer**—Stores runtime authentication information for EAP. • **eap**—Stores information about PPP EAP authentication protocol. Includes authenticator and authenticate state machines. • **authNone**—Stores information about no PPP authentication. Includes the authenticator state machine. • **networkInterface**—Stores information about NCP portions of PPP protocol. • **ipNcp**—PPP IPCP session information. Used for configuring, negotiating, and establishing IPCP protocol. Stores the current state, and configured and negotiated options. • **ipv6Ncp**—PPP IPv6CP session information. Used for configuring, negotiating, and establishing IPv6CP protocol. Stores the current state, and configured and negotiated options. • **osiNcp**—PPP OSICP session information. Used for configuring, negotiating, and establishing OSICP protocol. Stores the current state, and configured and negotiated options. • **mplsNcp**—PPP MPLSCP session information. Used for configuring, negotiating, and establishing MPLSCP protocol. Stores the current state. • **trace**—Stores information for PPP debugging.	**memory**
Total	Total memory allocations.	**detail**
Size	Size of the structure.	**detail**
Active	Number of instances of the structure that are used.	**detail**
Free	Number of instances of the structure that are on the free list. Types with a number in the **Free** column are pooled structures, and are typically types that are often used.	**detail**
Limit	Maximum number of instances that can be on the free list. Types with a number in the **Limit** column are pooled structures, and are typically types that are often used.	**detail**
Total size	Total amount of memory being used by a type of structure (includes active and free instances).	**detail**
Requests	Number of allocation requests made by a type of structure.	**detail**

Copyright © 2016, Juniper Networks, Inc.

Table 25: show ppp statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Failures	Number of failed allocations.	**detail**
Recovery state	State of PPP recovery after a GRES or restart: • recovery done—All sessions have recovered; it is safe to force another GRES or restart. • recovery cleanup pending—Not all PPP sessions have recovered; it is not safe to force another GRES or restart.	none
Subscriber sessions pending retention	Number of PPP subscriber sessions that are in the process of being recovered.	none
Subscriber sessions recovered OK	Number of PPP subscriber sessions that have recovered after a GRES or restart.	none
Subscriber sessions recovery failed	Number of PPP subscriber sessions that have failed to recover after a GRES or restart.	none

Sample Output

show ppp statistics

```
user@host> show ppp statistics
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0
```

show ppp statistics detail

```
user@host> show ppp statistics detail
Session statistics from PPP process
  Total sessions: 0
    Sessions in disabled phase    : 0
    Sessions in establish phase   : 0
    Sessions in authenticate phase: 0
    Sessions in network phase     : 0
    Bundles in pending phase      : 0
Type                   Size  Active  Free  Limit  Total size  Requests Failures
Queued rtsock msgs       28       0     0  65535           0         0
PPP session              60       0                         0         0
Interface address        64       0     0  65535           0         0
Destination profile      65       0                         0         0
ML link settings         68       0                         0         0
```


Copyright © 2016, Juniper Networks, Inc.

```
IPCP blocked address     68      0                          0         0
PPP session trace        76      0                          0         0
IFL redundancy state     76      0                          0         0
Protocol family          84      0      0  65535           0         0
ML bundle settings      108      0                          0         0
PPP LCP session         120      0                          0         0
PPP NCP session         124      0                          0         0
Physical interface      124    170      0  65535       21080       170
Access profile          132      0                          0         0
ML wait entry           144      0      0     20           0         0
Group profile           164      0                          0         0
Profile client          272      0                          0         0
PPP Auth session        356      0                          0         0
Logical interface       524      0      0  65535           0         0
Non-tagged                                                 8         2
Total                                                  21088       172       0

Session statistics from PPP universal edge process
  Total subscriber sessions: 32
    Subscriber sessions in disabled phase    : 32
    Subscriber sessions in establish phase   : 0
    Subscriber sessions in authenticate phase: 0
    Subscriber sessions in network phase     : 0
Type                   Size  Active  Free  Limit  Total size Requests Failures
authenticate            224       1    99  16384         224        0        0
linkInterface           152       1    99  16384         152        0        0
pap                     256       1    99  16384         256        0        0
lcp                     272       1    99  16384         272        0        0
chap                    284       0     0  16384           0        0        0
eapBuffer              1464       0     0  16384           0        0        0
eap                     276       0     0  16384           0        0        0
authNone
networkInterface        220       1    99  16384         220        0        0
ipNcp                   256       1    99  16384         256        0        0
ipv6Ncp                 204       0     0  16384           0        0        0
osiNcp                  192       0     0  16384           0        0        0
mplsNcp                 188       0     0  16384           0        0        0
trace                  2052       0    16     16           0        0        0
Total                                                   1380        0        0
```

show ppp statistics recovery (Safe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery done
    Subscriber sessions recovered OK      : 32001
    Subscriber sessions recovery failed   : 0
```

show ppp statistics recovery (Unsafe to Restart)

```
user@host> show ppp statistics recovery
Recovery statistics from PPP universal edge process
  Recovery state: recovery cleanup pending
    Subscriber sessions pending retention : 32001
    Subscriber sessions recovered OK      : 0
    Subscriber sessions recovery failed   : 0
```

show ppp summary

Syntax show ppp summary

Release Information Command introduced in Junos OS Release 7.5.

Description Display PPP session summary information.

Options This command has no options.

Required Privilege Level view

List of Sample Output show ppp summary on page 500

Output Fields Table 26 on page 500 lists the output fields for the **show ppp summary** command. Output fields are listed in the approximate order in which they appear.

Table 26: show ppp summary Output Fields

Field Name	Field Description
Interface	Interface on which the PPP session is running. An interface type of pp0 indicates an Ethernet interface type on a M120 or M320 router.
Session type	Type of session: **PPP** or **Cisco-HDLC**.
Session phase	PPP process phases: **Authenticate**, **Pending**, **Establish**, **Network**, **Disabled**.
Session flags	Special conditions present in the session, such as **Bundled**, **TCC**, **No-keepalives**, **Looped**, **Monitored**, and **NCP-only**.

Sample Output

show ppp summary

```
user@host> show ppp summary
Interface       Session type  Session phase      Session flags
at-4/0/0.456    PPP           Network
lsq-0/3/0.0     PPP           Disabled
lsq-1/0/0.0     PPP           Disabled
rlsq0.0         PPP           Network            NCP-only
so-1/0/0.0      PPP           Authenticate
so-1/0/1.0      PPP           Disabled           Looped
so-2/0/0.0      Cisco-HDLC    Establish
so-4/0/0.0      PPP           Establish          Monitored
t1-1/3/0:1.0    PPP           Network            Bundled
t1-1/3/0:2.0    PPP           Network            Bundled
pp0.12          PPP           Network
```

 Copyright © 2016, Juniper Networks, Inc.

show services inline ip-reassembly statistics

Syntax

show services inline ip-reassembly statistics
<fpc *fpc-slot*>
<pfe *pfe-slot*>

Release Information

Statement introduced in Junos OS Release 12.2X49.

Description

Display the inline IP reassembly statistics for the Packet Forwarding Engines on one or more MPCs. Inline IP reassembly statistics are collected at the Packet Forwarding Engine level.

NOTE: For more information on MPCs that support inline IP reassembly, refer to *Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPCEs*.

Options

none—Displays standard inline IP reassembly statistics for all MPCs.

fpc *fpc*—(Optional) Displays inline IP reassembly statistics for the specified MPC.

pfe *pfe*—(Optional) Displays inline IP reassembly for the specified Packet Forwarding Engine slot. You must specify an FPC slot number before specifying a Packet Forwarding Engine slot.

Required Privilege Level

view

Related Documentation

- ip-reassembly on page 345

List of Sample Output

show services inline ip-reassembly statistics fpc on page 505

Output Fields

Table 27 on page 501 lists the output fields for the **show services inline ip-reassembly statistics** command. Output fields are listed in the approximate order in which they appear.

Table 27: show services inline ip-reassembly statistics Output Fields

Field Name	Field Description
FPC	MPC slot number for which the statistics are displayed.
PFE	Packet Forwarding Engine on the MPC for which the statistics are displayed.

Table 27: show services inline ip-reassembly statistics Output Fields *(continued)*

Field Name	Field Description
NOTE: The output fields displayed (per Packet Forwarding Engine) are arranged in a logical sequence from top to bottom to enable users to understand how the inline IP reassembly statistics are gathered. The information about total number of fragments received is displayed first, and then the information about the reassembled packets and those pending reassembly are displayed. Then, the reasons why the fragments were dropped or not reassembled are displayed. Finally, the information about the fragments reassembled, fragments dropped, and fragments sent to the backup user plane PIC (services PIC) are displayed.	
Total Fragments Received	Total number of fragments received and the current rate of fragments received for inline IP reassembly. The following information is also displayed: • **First Fragments**—Number of first fragments received and current rate of first fragments processed. • **Intermediate Fragments**—Number of intermediate fragments received and current rate of intermediate fragments processed. • **Last Fragments**—Number and rate of last fragments received. NOTE: Current rate refers to the current number of fragments processed per second in the instant preceding the command's execution.
Total Packets Reassembled	Total number of packets reassembled and current rate, in the instant preceding the command's execution, at which the packets are reassembled.
Approximate Packets Pending Reassembly	Approximate number of packets pending reassembly.

Copyright © 2016, Juniper Networks, Inc.

Table 27: show services inline ip-reassembly statistics Output Fields *(continued)*

Field Name	Field Description
Fragments Dropped Reasons	Total number of fragments dropped reasons and the current rate of total fragment dropped reasons. The number of dropped reasons and rate corresponding to each of the following reasons are also displayed: • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment NOTE: • These fields indicate *why* a fragment was dropped. When a fragment is dropped, the corresponding reason field is incremented by 1. For example, when a fragment is dropped because the memory runs out, the **Buffers not available** field increases by 1. • The maximum number of fragments allowed for reassembly is 16. If the interface encounters a 17th fragment, it drops the entire packet and increments the **Fragment per packet exceeded** field by 17. • Current rate refers to the current number of fragment dropped reasons per second in the instant preceding the command's execution.
Reassembly Errors Reasons	Number of errors during reassembly and the current rate of reassembly errors. The number of errors and the rate for each of the following types of errors are also displayed: • Fragment not found • Fragment not in sequence • ASIC errors NOTE: Current rate refers to the current number of reassembly errors processed per second in the instant preceding the command's execution.
Aged out packets	Number of aged out packets and the current number of packets aged out per second in the instant preceding the command's execution. NOTE: In some cases, aged out packets can refer to aged out fragments. If previous fragments of the packet have already been discarded then linking of the dropped fragments to the aged out fragments cannot occur.
Total Fragments Successfully Reassembled	Number of fragments successfully reassembled and the current number of fragments reassembled per second in the instant preceding the command's execution.

Copyright © 2016, Juniper Networks, Inc.

Table 27: show services inline ip-reassembly statistics Output Fields *(continued)*

Field Name	Field Description
Total Fragments Dropped	Total number of fragments dropped and the current rate of total number of fragments dropped. The number of fragments dropped and rate corresponding to each of the following reasons are also displayed: • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment • Fragment not found • Fragment not in sequence • ASIC errors • Aged out fragments
Total fragments punted to UPIC	Number of fragments sent to the backup user plane PIC (services PIC) and current rate of fragments sent per second in the instant preceding the command's execution

The following information applies to the **Total Fragments Dropped** field.

- These fields indicate *how many* of the packet fragments received were then dropped due to a particular reason.

  For example, consider a packet that has 10 fragments, 9 of which have been received and stored in memory. When the tenth fragment arrives, if the memory runs out (Buffers not available), then this fragment is dropped. Because the tenth fragment has been dropped, the other 9 fragments must also be dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 10.

  For the next packet arriving, which also has 10 fragments, the first four fragments are stored but the memory runs out for the fifth fragment. Then the first 5 fragments (fifth and the first four) are dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 5.

  For the remaining fragments of the packet, if memory becomes available, the next 5 fragments (6 through 10) that arrive are stored in memory. The fragments are stored until the timeout period elapses, and are eventually dropped. In this case, the **Aged out packets** field is incremented by 1 and the **Aged out fragments** field (under the **Total Fragments Dropped** field) is incremented by 5.

  The fragment counters (after both packets have been processed) are as follows:

Copyright © 2016, Juniper Networks, Inc.

- **Fragments Dropped Reasons**
  - **Buffers not available** 2
  - **Aged out packets** 1
- **Total Fragment Dropped**
  - **Buffers not available** 15
  - **Aged out packets** 5

- Current rate refers to the current total number fragments dropped per second in the instant preceding the command's execution.

Sample Output

show services inline ip-reassembly statistics fpc

```
user@host> show services inline ip-reassembly statistics fpc 0
FPC: 0 PFE: 0
=============
                                                   Total   Current Rate
Total Fragments Received                       728177644          83529
   First Fragments                             260759430          29924
   Intermediate Fragments                      206658784          23681
   Last Fragments                              260759430          29924

Total Packets Successfully Reassembled         260746982          29924

Approximate Packets Pending Reassembly                 4

Fragments Dropped Reasons                          34558              3
   Buffers not available                               0              0
   Fragments per packet exceeded                       0              0
   Packet length exceeded                              0              0
   Record insert error                                 0              0
   Record in use error                             34558              3
   Duplicate first fragments                           0              0
   Duplicate last fragments                            0              0
   Missing first fragment                              0              0

Reassembly Errors Reasons                              0              0
   Fragment not found                                  0              0
   Fragment not in sequence                            0              0
   ASIC errors                                         0              0

Aged out packets                                      63              0

Total Fragments Successfully Reassembled       728142977          83528

Total Fragments Dropped                            34673              3
   Buffers not available                               0              0
   Fragments per packet exceeded                       0              0
   Packet length exceeded                              0              0
   Record insert error                                 0              0
   Record in use error                             34558              3
   Duplicate first fragments                           0              0
   Duplicate last fragments                            0              0
   Missing first fragment                              0              0
```

```
    Fragment not found                                0              0
    Fragment not in sequence                          0              0
    ASIC errors                                       0              0
    Aged out fragments                              115              0

Total fragments punted to UPIC                        0              0
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp destination

Syntax

show services l2tp destination
<brief | detail | extensive>
<local-gateway *gateway-address*>
<peer-gateway *gateway-address*>
<statistics>

Release Information

Command introduced in Junos OS Release 10.4.

Description

Display information about L2TP tunnel destinations.

Options

brief | detail | extensive—(Optional) Display the specified level of information.

local-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified local gateway address.

peer-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified peer gateway address.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with any of the level options, **brief**, **detail**, or **extensive**.

Required Privilege Level

view

Related Documentation

- clear services l2tp destination on page 454
- show services l2tp destination lockout on page 511
- show services l2tp session on page 512
- show services l2tp summary on page 520
- show services l2tp tunnel on page 525

List of Sample Output

show services l2tp destination on page 509
show services l2tp destination detail on page 509
show services l2tp destination extensive (LAC) on page 509
show services l2tp destination extensive (LNS) on page 510
show services l2tp destination statistics (LAC only on MX Series Routers) on page 510

Output Fields

Table 28 on page 507 lists the output fields for the **show services l2tp destination** command. Output fields are listed in the approximate order in which they appear.

Table 28: show services l2tp destination Output Fields

Field Name	Field Description	Level of Output
Local Name	Name of this destination.	All levels
Remote IP	IP address of the remote peer (LNS).	All levels

Table 28: show services l2tp destination Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tunnels	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed
Sessions	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed
State	Administrative state of the L2TP destination: • **Enabled**—No restrictions exist on creation or operation of sessions and tunnels for this destination. • **Disabled**—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the **Disabled** state.	All levels
Local IP	IP address of the local gateway (LAC).	**detail extensive**
Transport	Medium used for tunneling. Only **ipUdp** is supported.	**detail extensive**
Logical System	Logical system in which the tunnel is configured.	**detail extensive**
Router Instance	Routing instance in which the tunnel is configured.	**detail extensive**
Lockout State	Reachability state of the destination: • **not locked**—Destination is considered reachable. • **waiting for lockout timeout**—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	**detail extensive**
Access Line Information	State of tunnel support for L2TP AVPs that report attributes and traffic rates for subscriber access lines: • **enabled** • **disabled**	**detail extensive**
Speed Updates	State of tunnel support for L2TP AVP 97, which reports connection speed changes for subscriber access lines: • **enabled** • **disabled**	**detail extensive**
Connections	Number of total, active, and failed tunnel and session connections for the destination.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 28: show services l2tp destination Output Fields *(continued)*

Field Name	Field Description	Level of Output
Control Tx	Amount of control information transmitted, in packets and bytes.	**statistics**
Control Rx	Amount of control information received, in packets and bytes.	**statistics**
Data Tx	Amount of data transmitted, in packets and bytes.	**statistics**
Data Rx	Amount of data received, in packets and bytes.	**statistics**
Error Tx	Number of errors transmitted, in packets.	**statistics**
Error Rx	Number of errors received, in packets.	**statistics**

Sample Output

show services l2tp destination

```
user@host> show services l2tp destination
  Local Name    Remote IP        Tunnels       Sessions  State
  1             10.10.1.1        1             1         Enabled
```

show services l2tp destination detail

```
user@host> show services l2tp destination detail
 Local name: 1
    Remote IP: 10.1.1.1
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 10.1.1.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: enabled, Speed Updates: enabled
 Local name: 1
    Remote IP: 10.1.1.8
    Tunnels: 1, Sessions: 1
    State: Enabled
    Local IP: 10.1.1.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: waiting for lockout timeout
    Access Line Information: enabled, Speed Updates: enabled
```

show services l2tp destination extensive (LAC)

```
user@host> show services l2tp destination extensive
  Local name: 1
    Remote IP: 10.1.1.1
    State: Enabled
    Local IP: 10.1.1.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
    Access Line Information: enabled, Speed Updates: enabled
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination extensive (LNS)

```
user@host> show services l2tp destination extensive
  Local name: 3
    Remote IP: 11.1.1.3
    State: Enabled
    Local IP: 11.1.1.2
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
```

show services l2tp destination statistics (LAC only on MX Series Routers)

```
user@host> show services l2tp destination statistics
 Local name: 2, Tunnels: 1, Sessions: 210
                    Packets        Bytes
 Control Tx             680        63.3k
 Control Rx             283        10.6k
 Data Tx               1129        14.3k
 Data Rx                877        10.9k
 Errors Tx                0
 Errors Rx                0
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp destination lockout

Syntax **show services l2tp destination lockout**

Release Information Command introduced in Junos OS Release 13.2.

Description Display a list of destinations that are currently locked out and the time remaining for each to remain in the lockout state.

Options This command has no options.

Required Privilege Level view

Related Documentation

- clear services l2tp destination on page 454
- request services l2tp destination unlock on page 398
- show services l2tp destination on page 507
- show services l2tp session on page 512
- show services l2tp summary on page 520
- show services l2tp tunnel on page 525

List of Sample Output show services l2tp destination lockout on page 511

Output Fields Table 29 on page 511 lists the output fields for the **show services l2tp destination lockout** command. Output fields are listed in the approximate order in which they appear.

Table 29: show services l2tp destination lockout Output Fields

Field Name	Field Description
Destination	Name of the destination.
Time Remaining	Time remaining for the destination to be locked out.
L2TP lockout destinations found	Total count of lockout destinations.

Sample Output

show services l2tp destination lockout

```
user@host> show services l2tp destination lockout
  Destination   Time Remaining
  4             45
  5             43
  6             8
3 L2TP lockout destinations found
```

show services l2tp session

Syntax

show services l2tp session
<brief | detail | extensive>
<interface *interface-name*>
<local-gateway *gateway-address*>
<local-gateway-name *gateway-name*>
<local-session-id *session-id*>
<local-tunnel-id *tunnel-id*>
<peer-gateway *gateway-address*>
<peer-gateway-name *gateway-name*>
<statistics>
<tunnel-group *group-name*>
<user *username*>

Release Information

Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.

Description

(M10i and M7i routers only) Display information about active L2TP sessions for LNS.

(MX Series routers only) Display information about active L2TP sessions for LAC and LNS.

Options

none—Display standard information about all active L2TP sessions.

brief | detail | extensive—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display L2TP session information for only the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- **si-*fpc/pic/port***— MPCs on MX Series routers only. This option is not available for L2TP on M Series routers.
- **sp-*fpc/pic/port***—AS or Multiservices PICs on M7i, M10i, and M120 routers only. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified local gateway address.

local-gateway-name *gateway-name*—(Optional) Display L2TP session information for only the specified local gateway name.

local-session-id *session-id*—(Optional) Display L2TP session information for only the specified local session identifier.

local-tunnel-id *tunnel-id*—(Optional) Display L2TP session information for only the specified local tunnel identifier.

peer-gateway *gateway-address*—(Optional) Display L2TP session information for only the specified peer gateway address.

Copyright © 2016, Juniper Networks, Inc.

peer-gateway-name ***gateway-name***—(Optional) Display L2TP session information for only the specified peer gateway name.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with any of the level options, **brief**, **detail**, or **extensive**.

tunnel-group ***group-name***—(Optional) Display L2TP session information for only the specified tunnel group. To display information about L2TP CPU and memory usage, you can include the tunnel group name in the **show services service-sets memory-usage** ***group-name*** and **show services service-sets cpu-usage** ***group-name*** commands. This option is not available for L2TP LAC on MX Series routers.

user ***username***—(M Series routers only) (Optional) Display L2TP session information for only the specified username.

Required Privilege Level

view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*
- clear services l2tp session on page 455

List of Sample Output

show services l2tp session (LNS on M Series Routers) on page 516
show services l2tp session (LNS on MX Series Routers) on page 516
show services l2tp session (LAC) on page 516
show services l2tp session detail (LAC) on page 516
show services l2tp session extensive (LAC) on page 517
show services l2tp session extensive (LAC on MX Series Routers) on page 517
show services l2tp session extensive (LNS on M Series Routers) on page 517
show services l2tp session extensive (LNS on MX Series Routers) on page 518
show services l2tp session statistics (MX Series Routers) on page 518

Output Fields

Table 30 on page 513 lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

Table 30: show services l2tp session Output Fields

Field Name	Field Description	Level of Output
Interface	(LNS only) Name of an adaptive services interface.	All levels
Tunnel group	(LNS only) Name of a tunnel group.	All levels
Tunnel local ID	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 30: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
State	State of the L2TP session: • **Established**—Session is operating. This is the only state supported for the LAC. • **closed**—Session is being closed. • **destroyed**—Session is being destroyed. • **clean-up**—Session is being cleaned up. • **lns-ic-accept-new**—New session is being accepted. • **lns-ic-idle**—Session has been created and is idle. • **lns-ic-reject-new**—New session is being rejected. • **lns-ic-wait-connect**—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
Bundle ID	(LNS only) Bundle identifier. Indicates the session is part of a multilink bundle. Sessions that have a blank **Bundle** field are not participating in the Multilink Protocol. Sessions in a multilink bundle might belong to different L2TP tunnels. For L2TP output organized by bundle ID, issue the **show services l2tp multilink extensive** command.	All levels
Mode	(LNS) Mode of the interface representing the session: **shared** or **exclusive**. (LAC) Mode of the interface representing the session: **shared** or **dedicated**. Only **dedicated** is currently supported for the LAC.	**extensive**
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	**extensive**
Remote IP	IP address of remote endpoint of the PPP session.	**extensive**
Username	(LNS only) Name of the user logged in to the session.	All levels
Assigned IP address	(LNS only) IP address assigned to remote client.	**extensive**
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	**extensive**
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	**extensive**
Local MRU	(LNS only) Maximum receive unit (MRU) setting of the local device, in bytes.	**extensive**
Remote MRU	(LNS only) MRU setting of the remote device, in bytes.	**extensive**
Tx speed	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps). Both the initial (**initial**) and current (**update**) line speeds can be displayed on MX Series routers.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 30: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Rx speed	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps). Both the initial (**initial**) and current (**update**) line speeds can be displayed on MX Series routers.	**extensive**
Bearer type	Type of bearer enabled: • **0**—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • **1**—Digital access requested. • **2**—Analog access requested. • **4**—Asynchronous Transfer Mode (ATM) bearer support.	**extensive**
Framing type	Type of framing enabled: • **1**—Synchronous framing • **2**—Asynchronous framing	**extensive**
LCP renegotiation	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: **On** or **Off**.	**extensive**
Authentication	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	**extensive**
Interface ID	(LNS only) Identifier used to look up the logical interface for this session.	**extensive**
Interface unit	Logical interface for this session.	All levels
Call serial number	Unique serial number assigned to the call.	**extensive**
Policer bandwidth	Maximum policer bandwidth configured for this session.	**extensive**
Policer burst size	Maximum policer burst size configured for this session.	**extensive**
Firewall filter	Configured firewall filter name.	**extensive**
Session encapsulation overhead	Overhead allowance configured for this session, in bytes.	**extensive**
Session cell overhead	Cell overhead activation (**On** or **Off**).	**extensive**
Create time	Date and time when the call was created.	**extensive**
Up time	Length of time elapsed since the call became active, in hours, minutes, and seconds.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 30: show services l2tp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	**extensive**
Statistics since	Date and time when collection of the following statistics began: • **Control Tx**—Amount of control information transmitted, in packets and bytes. • **Control Rx**—Amount of control information received, in packets and bytes. • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes. • **Errors Tx**—Number of errors transmitted, in packets. • **Errors Rx**—Number of errors received, in packets. • **LCP echo req Tx**—Number of LCP echo requests transmitted, in packets. • **LCP echo req Rx**—Number of LCP echo requests received, in packets. • **LCP echo rep Tx**—Number of LCP echo responses transmitted, in packets. • **LCP echo rep Rx**—Number of LCP echo responses received, in packets. • **LCP echo Req timout**—Number of LCP echo requests that timed out. • **LCP echo Req error**—Number of errors received for LCP echo packets. • **LCP echo Rep error** —Number of errors transmitted for LCP echo packets.	**extensive**

Sample Output

show services l2tp session (LNS on M Series Routers)

```
user@host> show services l2tp session
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID:  8802
  Local Remote Interface State               Bundle Username
  ID    ID     unit
  37966      5         2 Established
```

show services l2tp session (LNS on MX Series Routers)

```
user@host> show services l2tp session
Tunnel local ID: 40553
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  17967  1        Established          1073749824        si-5/2/0
```

show services l2tp session (LAC)

```
user@host> show services l2tp session
Tunnel local ID: 31889
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
   31694      1   Established                311         pp0
```

show services l2tp session detail (LAC)

```
user@host> show services l2tp session detail
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1, Interface unit: 311
    State: Established, Interface: pp0, Mode: Dedicated
```


Copyright © 2016, Juniper Networks, Inc.

```
    Local IP: 10.1.1.2:1701, Remote IP: 10.1.1.1:1701
    Local name: ce-lac, Remote name: ce-lns
```

show services l2tp session extensive (LAC)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 10.10.1.2:1701, Remote IP: 10.10.1.1:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 0, Rx speed: 0
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LAC on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:     1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 10.10.1.2:1701, Remote IP: 10.10.1.1:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: initial 64000, Update 256000
    Rx speed: initial 64000, Update 256000
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LNS on M Series Routers)

```
user@host> show services l2tp session extensive
Interface: sp-1/2/0, Tunnel group: group1, Tunnel local ID: 62746
  Session local ID: 56793, Session remote ID: 53304
    State: Established, Bundle ID: 5, Mode: shared
    Local IP: 10.128.1.1:1701, Remote IP: 10.128.1.2:1701
    Username: usr1@juniper_1.net, Assigned IP address: 10.50.2.1/32
    Local MRU: 4000, Remote MRU: 1500, Tx speed: 64000, Rx speed: 64000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_20
    Interface unit: 20, Call serial number: 4137941434
    Policer bandwidth: 64000, Policer burst size: 51200
    Firewall filter: f1
    Session encapsulation overhead: 16, Session cell overhead: On
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:16:41
    Idle time: 00:00:00
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx              4           88
```

```
      Control Rx              2           28
      Data Tx                 0            0
      Data Rx               461        29.0k
      Errors Tx               0
      Errors Rx               0

Interface: sp-1/2/0, Tunnel group: group_company_dns, Tunnel local ID: 37266
  Session local ID: 39962, Session remote ID: 53303
    State: Established, Bundle ID: 5, Mode: shared
    Local IP: 10.128.11.1:1701, Remote IP: 10.128.11.2:1701
    Username: usr1@company.com, Assigned IP address: 10.46.2.3/24
    Local name: router-1, Remote name: router-2
    Local MRU: 4470, Remote MRU: 4470, Tx speed: 155000000, Rx speed: 155000000
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: CHAP, Interface ID: unit_31
    Interface unit: 31, Call serial number: 4137941433
    Policer bandwidth: 64000, Policer burst size: 51200
    Firewall filter: f1
    Create time: Tue Mar 23 14:13:17 2004, Up time: 01:16:39
    Idle time: 01:16:36
    Statistics since: Tue Mar 23 14:13:15 2004
                        Packets        Bytes
      Control Tx              6          196
      Control Rx              4          150
      Data Tx                 0            0
      Data Rx                 1           80
      Errors Tx               0
      Errors Rx               0
```

show services l2tp session extensive (LNS on MX Series Routers)

```
user@host> show services l2tp session extensive
Tunnel local ID: 40553
  Session local ID: 17967, Session remote ID: 1
    Interface unit: 1073749824
    State: Established
    Interface: si-5/2/0
    Mode: Dedicated
    Local IP: 11.1.1.2:1701, Remote IP: 11.1.1.3:1701
    Local name: lns-mx960, Remote name: testlac
    Tx speed: 56000, Rx speed: 0
    Bearer type: 2, Framing type: 1
    LCP renegotiation: Off, Authentication: None
    Call serial number: 1
    Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:48
    Idle time: N/A
    Statistics since: Mon Apr 25 20:27:50 2011
                        Packets        Bytes
      Control Tx              4          219
      Control Rx              4          221
      Data Tx                 0            0
      Data Rx                10          228
      Errors Tx               0
      Errors Rx               0
```

show services l2tp session statistics (MX Series Routers)

```
user@host>show services l2tp session statistics local session-id 1
Tunnel local ID: 17185
Session local ID: 1, Session remote ID: 14444, Interface unit: 1073788352
State: Established
```


Copyright © 2016, Juniper Networks, Inc.

```
Statistics since: Mon Aug 1 13:27:47 2011
    Packets   Bytes
Data Tx   4    51
Data Rx   3    36
```

show services l2tp summary

Syntax

show services l2tp summary
<interface sp-*fpc/pic/port*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Support for LAC on MX Series routers introduced in Junos OS Release 10.4.
Support for LNS on MX Series routers introduced in Junos OS Release 11.4.
Support for **statistics** option introduced in Junos OS Release 13.1.

Description

(M10i and M7i routers: LNS only. MX Series routers: LAC and LNS.) Display Layer 2 Tunneling Protocol (L2TP) summary information.

Options

none—Display complete L2TP summary information. For LNS on M Series routers, display L2TP summary information for all adaptive services interfaces. For LNS on MX Series routers, display L2TP summary information for all inline services interfaces.

interface sp-*fpc/pic/port*—(Optional) Display L2TP summary information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.

statistics—(Optional) Display a summary of control packets and bytes transmitted and received.

Required Privilege Level

view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*

List of Sample Output

show services l2tp summary (LAC on M Series routers) on page 523
show services l2tp summary (LAC on MX Series routers) on page 523
show services l2tp summary (LNS on MX Series routers) on page 523
show services l2tp summary (LNS on M Series routers) on page 523
show services l2tp summary statistics (MX Series routers) on page 524

Output Fields

Table 31 on page 520 lists the output fields for the **show services l2tp summary** command. Output fields are listed in the approximate order in which they appear.

Table 31: show services l2tp summary Output Fields

Field Name	Field Description
Failover within a preference level	State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level. Not displayed for LNS on M Series routers.

Copyright © 2016, Juniper Networks, Inc.

Table 31: show services l2tp summary Output Fields *(continued)*

Field Name	Field Description
Weighted load balancing	State of this tunnel selection method on the LAC. When enabled, the maximum session limit of a tunnel determines its weight within a preference level. Tunnel selection proceeds from greatest to least weight. When disabled, selection defaults to a round robin method. Not displayed for LNS on M Series routers.
Tunnel authentication challenge	State of tunnel authentication, indicating whether the LAC and LNS exchange an authentication challenge and response during the establishment of the tunnel. The state is **Enabled** when a secret is configured in the tunnel profile or on the RADIUS server in the Tunnel-Password attribute [69]. The state is **Disabled** when the secret is not present. Not displayed for LNS on M Series routers.
Calling number avp	When the state is **Enabled**, the LAC includes the value of the Calling Number AVP 22 in ICRQ packets sent to the LNS. When the state is **Disabled**, the attribute is not sent to the LNS. Not displayed for LNS on M Series routers.
Failover Protocol	When the state is enabled, the LAC operates in the default *failover-protocol-fall-back-to-silent-failover* manner. When the state is disabled, the **disable-failover-protocol** statement has been issued and the LAC operates only in silent failover mode. Not displayed for LNS on M Series routers.
Tx connect speed method	The connection speed method configured to send the speed values in the L2TP Tx Connect Speed (AVP 24) and L2TP Rx Connect Speed (AVP 38). Possible values are: • **ancp** • **pppoe-ia-tag** • **static**
Rx speed avp when equal	Indicates if the Rx connect speed when equal configuration is **enabled** or **disabled**.
Tunnel assignment id	Format of the tunnel name. Format of the tunnel name, based on RADIUS attributes returned from the AAA server: • **authentication-id**—Name consists of only Tunnel Assignment-Id [82]. This is the default value. • **client-server-id**—Name is a combination of Tunnel-Client-Auth-Id [90], Tunnel-Server-Endpoint [67], and Tunnel-Assignment-Id [82]. This format is available only on MX Series routers.

Copyright © 2016, Juniper Networks, Inc.

Table 31: show services l2tp summary Output Fields *(continued)*

Field Name	Field Description
Tunnel Tx Address Change	Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both: • **accept**—Accepts change requests for the IP address or UDP port. This is the default action. • **ignore**—Ignores all change requests. • **ignore-ip-address**—Ignores change requests for the IP address but accepts them for the UDP port. • **ignore-udp-port**—Ignores change requests for the UDP port but accepts them for the IP address.
Min Retransmission Timeout for control packets	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
Max Retransmissions for Established Tunnel	Maximum number of times control messages are retransmitted for established tunnels.
Max Retransmissions for Not Established Tunnel	Maximum number of times control messages are retransmitted for tunnels that are not established.
Tunnel Idle Timeout	Period that a tunnel can be inactive—that is, carrying no traffic—before it times out and is torn down.
Destruct Timeout	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
Destination Lockout Timeout	Timeout period for which all future destinations are locked out, meaning that they are not considered for selection when a new tunnel is created.
Reassembly Service Set	Indicates active IP reassembly configured for the interface.
Destinations	Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers.
Tunnels	Number of L2TP tunnels established on the router.
Sessions	Number of L2TP sessions established on the router.
Switched sessions	Number of L2TP tunnel-switched sessions established on the router.
Control	Count of L2TP control packets and bytes sent and received.
Data	Count of L2TP data packets and bytes sent and received.
Errors	Count of L2TP error packets and bytes sent and received.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show services l2tp summary (LAC on M Series routers)

```
user@host> show services l2tp summary
Failover within a preference level is Disabled
Weighted load balancing is Enabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Disabled
Tunnel assignment id format is authentication-id
Destinations: 1 Tunnels: 1, Sessions: 1
   Tx packets     Rx packets  Memory (bytes)
Control    260             144              11513856
Data       7.5k           16.9k             8.3k
Errors        0               0
```

show services l2tp summary (LAC on MX Series routers)

```
user@host> show services l2tp summary
Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is enabled
  Tunnel Tx Address Change is Accept
  Min Retransmissions Timeout for control packets is 2 seconds
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
Reassembly Service Set is ssnr3
Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0
```

show services l2tp summary (LNS on MX Series routers)

```
user@host show services l2tp summary
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Enabled
  Tx Connect speed method is static

reassembly Service Set is ssnr3  Destinations: 4, Tunnels: 19, Sessions: 65,
Switched sessions: 2
```

show services l2tp summary (LNS on M Series routers)

```
user@host> show services l2tp summary
Tunnels: 2, Sessions: 2, Errors: 0
    Tx packets   Rx packets    Memory (bytes)
Control     6k           9k                 688k
Data       70k          70k                 3054
```

show services l2tp summary statistics (MX Series routers)

```
user@host>show services l2tp summary statistics
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is advisory
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 4 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 secondsDestinations: 1, Tunnels: 1, Sessions:
 31815, Switched sessions: 0
     Tx packets    Rx packets  Memory (bytes)
  Control          90.4k           32.0k       245678080
  Data             127.3k          100.8kk         0
  Errors              0               0
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp tunnel

Syntax

show services l2tp tunnel
<brief | detail | extensive>
<interface sp-*fpc/pic/port*>
<local-gateway *gateway-address*>
<local-gateway-name *gateway-name*>
<local-tunnel-id *tunnel-id*>
<peer-gateway *gateway-address*>
<peer-gateway-name *gateway-name*>
<statistics>
<tunnel-group *group-name*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M10i and M7i routers only) Display information about active Layer 2 Tunneling Protocol (L2TP) tunnels for LNS.

(MX Series routers only) Display information about L2TP tunnels for LAC and LNS.

Options

none—Display standard information about all active L2TP tunnels.

brief | detail | extensive—(Default) Display the specified level of output.

interface sp-*fpc/pic/port*—(Optional) Display L2TP tunnel information for only the specified adaptive services interface. This option is not available for L2TP on MX Series routers.

local-gateway *gateway-address*—(Optional) Display L2TP tunnel information for only the specified local gateway address.

local-gateway-name *gateway-name*—(Optional) Display L2TP tunnel information for only the specified local gateway name.

local-tunnel-id *tunnel-id*—(Optional) Display L2TP tunnel information for only the specified local tunnel identifier.

peer-gateway *gateway-address*—(Optional) Display L2TP tunnel information for only the specified peer gateway address.

peer-gateway-name *gateway-name*—(Optional) Display L2TP tunnel information for only the specified peer gateway name.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with any of the level options, **brief**, **detail**, or **extensive**.

tunnel-group *group-name*—(Optional) Display L2TP tunnel information for only the specified tunnel group.

Required Privilege Level

view

Related Documentation

- *L2TP Services Configuration Overview*
- *L2TP Minimum Configuration*

List of Sample Output

show services l2tp tunnel (LAC) on page 528
show services l2tp tunnel detail (LAC) on page 528
show services l2tp tunnel detail (LAC on MX Series Routers) on page 528
show services l2tp tunnel detail (LNS on MX Series Routers) on page 528
show services l2tp tunnel extensive (LAC) on page 529
show services l2tp tunnel extensive (LNS on M Series Routers) on page 529
show services l2tp tunnel extensive (LNS on MX Series Routers) on page 530
show services l2tp tunnel statistics (MX Series Routers) on page 530

Output Fields

Table 32 on page 526 lists the output fields for the **show services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 32: show services l2tp tunnel Output Fields

Field Name	Field Description
Interface	(LNS only) Name of an adaptive services interface.
Tunnel group	(LNS only) Name of a tunnel group.
Local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.
Remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.
Remote IP	IP address of the peer endpoint of the tunnel.
Sessions	Number of L2TP sessions established through the tunnel.

Copyright © 2016, Juniper Networks, Inc.

Table 32: show services l2tp tunnel Output Fields *(continued)*

Field Name	Field Description
State	State of the L2TP tunnel: • **cc_responder_accept_new**—The tunnel has received and accepted the start control connection request (SCCRQ). • **cc_responder_reject_new**—The tunnel has received and rejected the SCCRQ. • **cc_responder_idle**—The tunnel has just been created. • **cc_responder_wait_ctl_conn**—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • **clean-up**—The tunnel is being cleaned up. • **closed**—The tunnel is being closed. • **destroyed**—The tunnel is being destroyed. • **Established**—The tunnel is operating. This is the only state supported for the LAC. • **Terminate**—The tunnel is terminating. • **Unknown**—The tunnel is not connected to the router.
Tunnel Name	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].
Local IP	IP address of the local endpoint of the tunnel.
Local name	Name used for local tunnel endpoint during tunnel negotiation.
Remote name	Name used for remote tunnel endpoint during tunnel negotiation.
Effective Peer Resync Mechanism	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.
Nas Port Method	NAS port method (type), which indicates whether the LAC sends Cisco NAS Port Info AVP (100) in ICRQs to the LNS: • **cisco-avp**—sends the AVP. • **none**—does not send the AVP.
Tunnel Logical System	Logical system in which the L2TP tunnel is brought up.
Tunnel Routing Instance	Routing instance in which the L2TP tunnel is brought up.
Max sessions	Maximum number of sessions that can be established on this tunnel.
Window size	Number of control messages that can be sent without receipt of an acknowledgment.
Hello interval	Interval between the transmission of hello messages, in seconds.

Copyright © 2016, Juniper Networks, Inc.

Table 32: show services l2tp tunnel Output Fields *(continued)*

Field Name	Field Description
Create time	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to **Unknown** and the **Create time** value resets.
Up time	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.
Idle time	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.
Statistics since	Date and time when collection of the following statistics began: • **Control Tx**—Amount of control information transmitted, in packets and bytes. • **Control Rx**—Amount of control information received, in packets and bytes. • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes. • **Errors Tx**—Number of errors transmitted, in packets. • **Errors Rx**—Number of errors received, in packets.

Sample Output

show services l2tp tunnel (LAC)

```
user@host> show services l2tp tunnel
  Local ID  Remote ID  Remote IP                Sessions  State
     17185          1  10.10.1.1:1701                  1  Established
```

show services l2tp tunnel detail (LAC)

```
user@host> show services l2tp tunnel detail
  Tunnel local ID: 31889, Tunnel remote ID:     1
    Remote IP: 100.1.1.1:1701
    Sessions: 1, State: Established
    Tunnel Name: 1/tunnel-to-LNS-1
    Local IP: 100.1.1.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
```

show services l2tp tunnel detail (LAC on MX Series Routers)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 10.10.1.1:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 100.1.1.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel detail (LNS on MX Series Routers)

```
user@host> show services l2tp tunnel detail
Tunnel local ID: 17301, Tunnel remote ID: 1
    Remote IP: 12.1.1.15:1701
```


Copyright © 2016, Juniper Networks, Inc.

```
    Sessions: 1, State: Established
    Tunnel Name: 2/2
    Local IP: 12.1.1.5:1701
    Local name: ce-bras-mx240-e, Remote name: testlac2
    Effective Peer Resync Mechanism: silent failover
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
```

show services l2tp tunnel extensive (LAC)

```
user@host> show services l2tp tunnel extensive
  Tunnel local ID: 17185, Tunnel remote ID:     1
    Remote IP: 10.10.1.1:1701
    Sessions: 1, State: Established
    Tunnel Name: 2/tunnel-to-LNS-2
    Local IP: 100.1.1.2:1701
    Local name: ce-lac, Remote name: ce-lns
    Effective Peer Resync Mechanism: failover protocol
    Max sessions: 32000, Window size: 4, Hello interval: 60
    Create time: Tue Nov  9 15:23:29 2010, Up time: 00:00:26
    Idle time: 00:00:00
```

show services l2tp tunnel extensive (LNS on M Series Routers)

```
user@host> show services l2tp tunnel extensive
Interface: sp-1/2/0, Tunnel group: group1
  Tunnel local ID: 62746, Tunnel remote ID: 16930
    Remote IP: 10.128.1.2:1701
    Sessions: 1, State: Established
    Local IP: 10.128.1.1:1701
    Local name: router-1, Remote name: router-2
    Max sessions: 50, Window size: 32, Hello interval: 60
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:58
    Idle time: 00:00:07
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx             80         1152
      Control Rx              3          272
      Data Tx                 0            0
      Data Rx               450        28.0k
      Errors Tx               0
      Errors Rx               0

Interface: sp-1/2/0, Tunnel group: group_company_dns
  Tunnel local ID: 37266, Tunnel remote ID: 36217
    Remote IP: 10.128.11.2:1701
    Sessions: 1, State: Established
    Local IP: 10.128.11.1:1701
    Local name: router-1, Remote name: router-2
    Max sessions: unlimited, Window size: 32, Hello interval: 60
    Create time: Tue Mar 23 14:13:15 2004, Up time: 01:14:59
    Idle time: 01:14:55
    Statistics since: Tue Mar 23 14:13:13 2004
                        Packets        Bytes
      Control Tx             81         1164
      Control Rx              3          273
      Data Tx                 0            0
      Data Rx                 1           80
      Errors Tx               0
      Errors Rx               0
```

show services l2tp tunnel extensive (LNS on MX Series Routers)

```
user@host> show services l2tp tunnel extensive
Tunnel local ID: 40553, Tunnel remote ID: 1
    Remote IP: 192.168.1.3:1701
    Sessions: 1, State: Established
    Tunnel Name: 3/1838
    Local IP: 10.1.1.2:1701
    Local name: lns-mx960, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: vrf1
    Max sessions: 60000, Window size: 4, Hello interval: 60
    Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:11
    Idle time: 00:00:00, ToS Reflect: Enabled
    Tunnel Group Name: tg1
    Statistics since: Mon Apr 25 20:27:50 2011
                        Packets        Bytes
      Control Tx              4          219
      Control Rx              4          221
      Data Tx                 0            0
      Data Rx                 6           64
      Errors Tx               0
      Errors Rx
```

show services l2tp tunnel statistics (MX Series Routers)

```
user@host>show services l2tp tunnel statistics
Tunnel local ID: 17185, Tunnel remote ID: 1
Sessions: 31.8k, State: Established
Statistics since: Mon Aug 1 13:21:38 2011
     Packets   Bytes
Control Tx     90.3k    9.0M
Control Rx     32.0k    1296.9k
Data Tx       127.3k   1591.6k
Data Rx       100.8k   1273.4k
Errors Tx      0
Errors Rx      0
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp tunnel-switch destination

Syntax

show services l2tp tunnel-switch destination
< detail | extensive>
<statistics>

Release Information

Command introduced in Junos OS Release 13.2.

Description

Display information about L2TP switched tunnel destinations.

Options

none—Display standard information for all L2TP switched tunnel destinations.

detail | extensive—(Optional) Display the specified level of information.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the destination. You cannot include this option with either of the level options, **detail** or **extensive**.

Required Privilege Level

view

Related Documentation

- show services l2tp tunnel-switch session on page 535
- show services l2tp tunnel-switch summary on page 540
- show services l2tp tunnel-switch tunnel on page 542

List of Sample Output

show services l2tp tunnel-switch destination on page 533
show services l2tp tunnel-switch destination detail on page 533
show services l2tp destination extensive on page 533
show services l2tp destination statistics on page 534

Output Fields

Table 33 on page 531 lists the output fields for the **show services l2tp tunnel-switch destination** command. Output fields are listed in the approximate order in which they appear.

Table 33: show services l2tp tunnel-switch destination Output Fields

Field Name	Field Description	Level of Output
Local Name	Name of this destination.	All levels
Remote IP	IP address of the remote peer (LNS).	All levels
Tunnels	Number of tunnel connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed

Table 33: show services l2tp tunnel-switch destination Output Fields *(continued)*

Field Name	Field Description	Level of Output
Sessions	Number of session connections for the destination in the following categories: • total • active • failed	All levels for total **extensive** for active and failed
Switched-sessions	Number of L2TP sessions established by tunnel switching.	All levels
State	Administrative state of the L2TP destination: • **Enabled**—No restrictions exist on creation or operation of sessions and tunnels for this destination. • **Disabled**—Existing sessions and tunnels for this destination have been disabled and no new sessions or tunnels are created while in the **Disabled** state.	All levels
Local IP	IP address of the local gateway (LAC).	**detail extensive**
Transport	Medium used for tunneling. Only **ipUdp** is supported.	**detail extensive**
Logical System	Logical system in which the tunnel is configured.	**detail extensive**
Router Instance	Routing instance in which the tunnel is configured.	**detail extensive**
Lockout State	Reachability state of the destination: • **not locked**—Destination is considered reachable. • **waiting for lockout timeout**—Destination is locked out by L2TP because it is unreachable, so no attempts are made to reach the destination until the lockout timeout (300 seconds) expires, unless this is the only destination available for tunneling the subscriber.	**detail extensive**
Connections	Number of total, active, and failed tunnel and session connections for the destination.	**extensive**
Control Tx	Amount of control information transmitted, in packets and bytes.	**extensive statistics**
Control Rx	Amount of control information received, in packets and bytes.	**extensive statistics**
Data Tx	Amount of data transmitted, in packets and bytes.	**extensive statistics**
Data Rx	Amount of data received, in packets and bytes.	**extensive statistics**
Error Tx	Number of errors transmitted, in packets.	**extensive statistics**
Error Rx	Number of errors received, in packets.	**extensive statistics**

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show services l2tp tunnel-switch destination

```
user@host> show services l2tp tunnel-switch destination
  Local Name   Remote IP       Tunnels  Sessions Switched-sessions State
  1            192.168.20.3    1        1        1                 Enabled
  2            10.1.1.10       1        1        1                 Enabled
```

show services l2tp tunnel-switch destination detail

```
user@host> show services l2tp destination detail
  Local name: 1
    Remote IP: 192.168.20.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 10.50.1.1
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
  Local name: 2
    Remote IP: 172.20.1.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 10.30.1.1
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
```

show services l2tp destination extensive

```
user@host> show services l2tp destination extensive
Waiting for statistics...
  Local name: 1
    Remote IP: 192.168.20.3
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 10.50.1.1
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              6          239
      Control Rx              6          267
      Data Tx                67          815
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
  Local name: 2
    Remote IP: 172.20.1.10
    Tunnels: 1, Sessions: 1, Switched sessions: 1
    State: Enabled
    Local IP: 10.30.1.1
    Transport: ipUdp, Logical System: default, Router Instance: default
    Lockout State: not locked
      Connections     Totals        Active       Failed
      Tunnels              1             1            0
      Sessions             1             1            0
                        Packets        Bytes
      Control Tx              7          462
      Control Rx              6          171
```

```
    Data Tx                 0            0
    Data Rx                66          798
    Errors Tx               0
    Errors Rx               0
```

show services l2tp destination statistics

```
user@host> show services l2tp destination statistics
Waiting for statistics...
  Local name: 2, Tunnels: 1, Sessions: 1
                        Packets        Bytes
    Control Tx              5          452
    Control Rx              4          147
    Data Tx                 0            0
    Data Rx                 4           54
    Errors Tx               0
    Errors Rx               0
  Local name: 1, Tunnels: 1, Sessions: 1
                        Packets        Bytes
    Control Tx              4          184
    Control Rx              4          243
    Data Tx                 5           71
    Data Rx                 0            0
    Errors Tx               0
    Errors Rx               0
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp tunnel-switch session

Syntax

show services l2tp tunnel-switch session
<detail | extensive>
<statistics>

Release Information

Command introduced in Junos OS Release 13.2.

Description

Display information about L2TP switched tunnel sessions.

Options

none—Display standard information about all active L2TP switched tunnel sessions.

detail | extensive—(Optional) Display the specified level of output.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the session. You cannot include this option with either of the level options, **detail** or **extensive**.

Additional Information

Required Privilege Level

view

Related Documentation

- show services l2tp tunnel-switch destination on page 531
- show services l2tp tunnel-switch summary on page 540
- show services l2tp tunnel-switch tunnel on page 542

List of Sample Output

show services l2tp tunnel-switch session on page 537
show services l2tp tunnel-switch session detail on page 538
show services l2tp tunnel-switch session extensive on page 538

Output Fields

Table 34 on page 535 lists the output fields for the **show services l2tp tunnel-switch session** command. Output fields are listed in the approximate order in which they appear.

Table 34: show services l2tp tunnel-switch session Output Fields

Field Name	Field Description	Level of Output
Tunnel local ID	Identifier of the local endpoint of the tunnel, as assigned by the L2TP network server (LNS).	All levels
Local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	none
Remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	none

Copyright © 2016, Juniper Networks, Inc.

Table 34: show services l2tp tunnel-switch session Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the L2TP session: • **Established**—Session is operating. This is the only state supported for the LAC. • **closed**—Session is being closed. • **destroyed**—Session is being destroyed. • **clean-up**—Session is being cleaned up. • **lns-ic-accept-new**—New session is being accepted. • **lns-ic-idle**—Session has been created and is idle. • **lns-ic-reject-new**—New session is being rejected. • **lns-ic-wait-connect**—Session is waiting for the peer's incoming call connected (ICCN) message.	All levels
Interface unit	Logical interface for this session.	All levels
Interface Name	(LNS only) Name of an adaptive services interface.	none
Session local ID	Identifier of the local endpoint of the L2TP session, as assigned by the LNS.	**detail extensive**
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	**detail extensive**
Tunnel switch profile name	Name of a tunnel switch profile.	**detail extensive**
Mode	(LNS) Mode of the interface representing the session: **shared** or **exclusive**. (LAC) Mode of the interface representing the session: **shared** or **dedicated**. Only **dedicated** is currently supported for the LAC.	**detail extensive**
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	**detail extensive**
Remote IP	IP address of remote endpoint of the PPP session.	**detail extensive**
Local name	For LNS, name of the LNS instance in which the session was created. For LAC, name of the LAC.	**detail extensive**
Remote name	For LNS, name of the LAC from which the session was created. For LAC, name of the LAC instance.	**detail extensive**
Bearer type	Type of bearer enabled: • **0**—Might indicate that the call was not received over a physical link (for example, when the LAC and PPP are located in the same subsystem). • **1**—Digital access requested. • **2**—Analog access requested. • **4**—Asynchronous Transfer Mode (ATM) bearer support.	**extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 34: show services l2tp tunnel-switch session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Framing type	Type of framing enabled: • **1**—Synchronous framing • **2**—Asynchronous framing	**extensive**
LCP renegotiation	(LNS only) Whether Link Control Protocol (LCP) renegotiation is configured: **On** or **Off**.	**extensive**
Authentication	Type of authentication algorithm used: Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).	**extensive**
Interface ID	(LNS only) Identifier used to look up the logical interface for this session.	**extensive**
Call serial number	Unique serial number assigned to the call.	**extensive**
Tx speed	Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	**extensive**
Rx speed	Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps).	**extensive**
Create time	Day, date, and time when the call was created.	**extensive**
Up time	Length of time elapsed since the call became active, in hours, minutes, and seconds.	**extensive**
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	**extensive**
ToS Reflect	Status of IP ToS value reflection, **Disabled** or **Enabled**.	**extensive**
Statistics since	Date and time when collection of the following statistics began: • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes.	**extensive**

Sample Output

show services l2tp tunnel-switch session

```
user@host> show services l2tp tunnel-switch session
Tunnel local ID: 37602
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  13545  1        Established          1073741842        si-2/1/0

Tunnel local ID: 37060
  Local  Remote   State                Interface         Interface
  ID     ID                            unit              Name
  58296  1        Established          1073741843        si-2/1/0
```

show services l2tp tunnel-switch session detail

```
user@host> show services l2tp tunnel-switch session detail
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1, Interface unit: 1073741842
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 10.50.1.1:1701, Remote IP: 192.168.20.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1, Interface unit: 1073741843
    State: Established, Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 10.30.1.1:1701, Remote IP: 172.20.1.10:1701
    Local name: lns, Remote name: lns
```

show services l2tp tunnel-switch session extensive

```
user@host> show services l2tp tunnel-switch session extensive
Tunnel local ID: 37602
  Session local ID: 13545, Session remote ID: 1
    Interface unit: 1073741842
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 10.50.1.1:1701, Remote IP: 192.168.20.3:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Bearer type: 2, Framing type: 1
    LCP renegotiation: On, Authentication: None, Interface ID: si-2/1/0
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:06:50
    Idle time: N/A, ToS Reflect: Disabled
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets        Bytes
      Data Tx                85         1031
      Data Rx                 0            0

Tunnel local ID: 37060
  Session local ID: 58296, Session remote ID: 1
    Interface unit: 1073741843
    State: Established
    Interface: si-2/1/0
    Tunnel switch profile name: ce-lts-profile
    Mode: Dedicated
    Local IP: 10.30.1.1:1701, Remote IP: 172.20.1.10:1701
    Local name: lns, Remote name: lns
    Bearer type: 2, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Call serial number: 0
    Tx speed: 56000, Rx speed: 0
    Create time: Fri Jan 18 03:01:14 2013, Up time: 00:06:48
    Idle time: N/A
    Statistics since: Fri Jan 18 03:01:14 2013
                        Packets        Bytes
      Data Tx                 0            0
      Data Rx                84         1014
```


Copyright © 2016, Juniper Networks, Inc.

show services l2tp tunnel-switch summary

Syntax

show services l2tp tunnel-switch summary
<statistics>

Release Information

Command introduced in Junos OS Release 13.2.

Description

Display L2TP tunnel switch summary information.

Options

none—Display complete L2TP switched tunnel summary information.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for all switched tunnels and sessions.

Additional Information

Required Privilege Level

view

Related Documentation

- show services l2tp tunnel-switch destination on page 531
- show services l2tp tunnel-switch session on page 535
- show services l2tp tunnel-switch tunnel on page 542

List of Sample Output

show services l2tp tunnel-switch summary on page 541

Output Fields

Table 35 on page 540 lists the output fields for the **show services l2tp tunnel-switch summary** command. Output fields are listed in the approximate order in which they appear.

Table 35: show services l2tp tunnel-switch summary Output Fields

Field Name	Field Description
Tunnel switch profile name	Name of a tunnel switch profile.
LNS local session id	Identifier assigned by the LNS function on the LTS to the local endpoint of the L2TP session originating on a remote LAC (the first session)
LAC local session id	Identifier assigned by the LAC function on the LTS to the local endpoint of the L2TP session originating on the LTS (the second session).
LNS state	State of the L2TP session (the first session) between a remote LAC and the LNS function on the LTS.
LAC state	State of the L2TP session (the second session) between the LAC function on the LTS and a remote LNS.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

show services l2tp tunnel-switch summary

```
user@host> show services l2tp tunnel-switch summary
Tunnel switch profile name: ce-lts-profile
  LNS local   LAC local   LNS state      LAC state      Interface
  session ID  session ID                                name
  13545       58296       established    established    si-2/1/0
```

show services l2tp tunnel-switch tunnel

Syntax

show services l2tp tunnel-switch tunnel
<detail | extensive>
<statistics>

Release Information

Command introduced in Junos OS Release 13.2.

Description

Display information about L2TP switched tunnels.

Options

none—Display standard information about all active L2TP tunnels.

detail | extensive—(Default) Display the specified level of output.

statistics—(Optional) Display the number of control packets and bytes transmitted and received for the tunnel. You cannot include this option with either of the level options, **detail** or **extensive**.

Additional Information

Required Privilege Level

view

Related Documentation

- show services l2tp tunnel-switch destination on page 531
- show services l2tp tunnel-switch session on page 535
- show services l2tp tunnel-switch summary on page 540

List of Sample Output

show services l2tp tunnel-switch tunnel on page 544
show services l2tp tunnel-switch tunnel detail on page 545
show services l2tp tunnel-switch tunnel extensive on page 545

Output Fields

Table 36 on page 542 lists the output fields for the **show services l2tp tunnel-switch tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 36: show services l2tp tunnel-switch tunnel Output Fields

Field Name	Field Description	Level of Output
Local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	none
Remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	none
Remote IP	IP address of the peer endpoint of the tunnel.	All levels

Copyright © 2016, Juniper Networks, Inc.

Table 36: show services l2tp tunnel-switch tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
Sessions	Number of L2TP sessions established through the tunnel.	All levels
Switched-sessions	Number of L2TP sessions established by tunnel switching.	All levels
State	State of the L2TP tunnel: • **cc_responder_accept_new**—The tunnel has received and accepted the start control connection request (SCCRQ). • **cc_responder_reject_new**—The tunnel has received and rejected the SCCRQ. • **cc_responder_idle**—The tunnel has just been created. • **cc_responder_wait_ctl_conn**—The tunnel has sent the start control connection response (SCCRP) and is waiting for the start control connection connected (SCCCN) message. • **clean-up**—The tunnel is being cleaned up. • **closed**—The tunnel is being closed. • **destroyed**—The tunnel is being destroyed. • **Established**—The tunnel is operating. This is the only state supported for the LAC. • **Terminate**—The tunnel is terminating. • **Unknown**—The tunnel is not connected to the router.	All levels
Tunnel local ID	On the LNS, number assigned by the LNS that identifies the local endpoint of the tunnel relative to the LNS: the LNS. On the LAC, number assigned by the LAC that identifies the local endpoint of the tunnel relative to the LAC: the LAC.	**detail extensive**
Tunnel remote ID	On the LNS, number assigned by the LAC that identifies the remote endpoint of the tunnel relative to the LNS: the LAC. On the LAC, number assigned by the LNS that identifies the remote endpoint of the tunnel relative to the LAC: the LNS.	**detail extensive**
Tunnel Name	(LAC only) Name of the created tunnel. This value includes the destination name followed by the value of the RADIUS Tunnel-Assignment-ID VSA [82].	**detail extensive**
Local IP	IP address of the local endpoint of the tunnel.	**detail extensive**
Local name	Name used for local tunnel endpoint during tunnel negotiation.	**detail extensive**
Remote name	Name used for remote tunnel endpoint during tunnel negotiation.	**detail extensive**
Effective Peer Resync Mechanism	(LAC only) Peer resynchronization mechanism (PRM) in effect for the tunnel: • Failover protocol • Silent failover—Recovery takes place in the failed endpoint only using the proprietary silent failover protocol.	**detail extensive**

Copyright © 2016, Juniper Networks, Inc.

Table 36: show services l2tp tunnel-switch tunnel Output Fields *(continued)*

Field Name	Field Description	Level of Output
NAS Port Method	(LAC only) Status of interoperation with Cisco LNS devices: • none—NAS port method is not enabled for interoperation. • cisco-avp—NAS port method is enabled for interoperation.	**detail extensive**
Tunnel Logical System	Logical system in which the L2TP tunnel is brought up.	**detail extensive**
Tunnel Routing Instance	Routing instance in which the L2TP tunnel is brought up.	**detail extensive**
Max sessions	Maximum number of sessions that can be established on this tunnel.	**extensive**
Window size	Number of control messages that can be sent without receipt of an acknowledgment.	**extensive**
Hello interval	Interval between the transmission of hello messages, in seconds.	**extensive**
Create time	Date and time when the tunnel was created. While the LNS and LAC are connected, this value should correspond to the router's uptime. If connection to the LAC is severed, the State changes to **Unknown** and the **Create time** value resets.	**extensive**
Up time	Amount of time elapsed since the tunnel became active, in hours, minutes, and seconds.	**extensive**
Idle time	Amount of time elapsed since the tunnel became idle, in hours, minutes, and seconds.	**extensive**
ToS Reflect	Status of IP ToS value reflection, **Disabled** or **Enabled**.	**extensive**
Interface Name	(LNS only) Name of an adaptive services interface.	**extensive**
Tunnel Group Name	(LNS only) Name of a tunnel group.	**extensive**
Statistics since	Date and time when collection of the following statistics began: • **Control Tx**—Amount of control information transmitted, in packets and bytes. • **Control Rx**—Amount of control information received, in packets and bytes. • **Data Tx**—Amount of data transmitted, in packets and bytes. • **Data Rx**—Amount of data received, in packets and bytes. • **Errors Tx**—Number of errors transmitted, in packets. • **Errors Rx**—Number of errors received, in packets.	**extensive**

Sample Output

show services l2tp tunnel-switch tunnel

```
user@host> show services l2tp tunnel-switch tunnel
```


Copyright © 2016, Juniper Networks, Inc.

```
  Local ID Remote ID Remote IP              Sessions Switched-sessions State
  37602    1         192.168.20.3:1701            1                1 Established

  37060    1         172.20.1.10:1701             1                1 Established
```

show services l2tp tunnel-switch tunnel detail

```
user@host> show services l2tp tunnel-switch tunnel detail
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.168.20.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 1/1
    Local IP: 10.50.1.1:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 172.20.1.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 10.30.1.1:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
```

show services l2tp tunnel-switch tunnel extensive

```
user@host> show services l2tp tunnel-switch tunnel extensive
Waiting for statistics...
  Tunnel local ID: 37602, Tunnel remote ID: 1
    Remote IP: 192.168.20.3:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 1/1
    Local IP: 10.50.1.1:1701
    Local name: ce-bras-mx240-f, Remote name: testlac
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
    Tunnel Logical System: default, Tunnel Routing Instance: default
    Max sessions: 128100, Window size: 4, Hello interval: 60
    Create time: Fri Jan 18 03:01:11 2013, Up time: 00:07:49
    Idle time: 00:00:00, ToS Reflect: Disabled
    Interface Name: si-2/1/0, Tunnel Group Name: ce-l2tp-tunnel-group
    Statistics since: Fri Jan 18 03:01:11 2013
                        Packets        Bytes
      Control Tx              7          259
      Control Rx              7          279
      Data Tx                97         1175
      Data Rx                 0            0
      Errors Tx               0
      Errors Rx               0
  Tunnel local ID: 37060, Tunnel remote ID: 1
    Remote IP: 172.20.1.10:1701
    Sessions: 1, Switched sessions: 1, State: Established
    Tunnel Name: 2/1
    Local IP: 10.30.1.1:1701
    Local name: lns, Remote name: lns
    Effective Peer Resync Mechanism: silent failover
    Nas Port Method: none
```

```
Tunnel Logical System: default, Tunnel Routing Instance: default
Max sessions: 128100, Window size: 4, Hello interval: 60
Create time: Fri Jan 18 03:01:14 2013, Up time: 00:07:46
Idle time: 00:00:00
Statistics since: Fri Jan 18 03:01:14 2013
                    Packets        Bytes
  Control Tx              8          482
  Control Rx              7          183
  Data Tx                 0            0
  Data Rx                96         1158
  Errors Tx               0
  Errors Rx               0
```


Copyright © 2016, Juniper Networks, Inc.

show subscribers

Syntax

show subscribers
<detail | extensive | terse>
<aci-interface-set-name *aci-interface-set-name*>
<address *address*>
<agent-circuit-identifier *agent-circuit-identifier-substring*>
<client-type *client-type*>
<count>
<id>
<interface *interface*>
<logical-system *logical-system*>
<mac-address *mac-address*>
<physical-interface *physical-interface-name*>
<profile-name *profile-name*>
<routing-instance *routing-instance*>
<stacked-vlan-id *stacked-vlan-id*>
<subscriber-state *subscriber-state*>
<user-name *user-name*>
<vci *vci-identifier*>
<vpi *vpi-identifier*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 9.3.
Command introduced in Junos OS Release 9.3 for EX Series switches.
client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.
Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information for active subscribers.

Options

detail | extensive | terse—(Optional) Display the specified level of output.

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.168.17.1). If you specify the IP address as a prefix with a netmask (for example,

192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is **0** through **255**.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is **0** through **65535**.

 Copyright © 2016, Juniper Networks, Inc.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** ***stacked-vlan-id*** option to match the outer VLAN tag.

NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

Related Documentation

- show subscribers summary on page 565
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show subscribers stacked-vlan-id detail on page 563
show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 563
show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 563
show subscribers user-name detail on page 563
show subscribers vlan-id on page 564
show subscribers vlan-id detail on page 564
show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session) on page 564

Output Fields Table 37 on page 550 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 37: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is **Tunnel-switched**.
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
Primary DNS Address	IP address of primary DNS server.
Secondary DNS Address	IP address of secondary DNS server.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.

 Copyright © 2016, Juniper Networks, Inc.

Table 37: show subscribers Output Fields *(continued)*

Field Name	Field Description
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface Type	Whether the subscriber interface is **Static** or **Dynamic**.
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.
Interface Set Type	Interface type of the ACI interface set: **Dynamic**. This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (**Init**, **Configured**, **Active**, **Terminating**, **Tunneled**).
L2TP State	Current state of the L2TP session, **Tunneled** or **Tunnel-switched**. When the value is **Tunnel-switched**, two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.

Copyright © 2016, Juniper Networks, Inc.

Table 37: show subscribers Output Fields *(continued)*

Field Name	Field Description
Agent Remote ID	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	**len** = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.

Copyright © 2016, Juniper Networks, Inc.

Table 37: show subscribers Output Fields *(continued)*

Field Name	Field Description
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
Interface               IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824     100                                    default:default
demux0.1073741824       10.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825       11.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826       12.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers (IPv6)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name           LS:RI
ge-1/0/0.0         2001:db8::c0:0:0:0/74    WHOLESALER-CLIENT   default:default
     *             2001:db8::1/128          subscriber-25       default:default
```

show subscribers (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers
Interface           IP Address/VLAN ID         User Name
LS:RI
demux0.1073741834   0x8100.1002 0x8100.1
```

```
default:default
demux0.1073741835   0x8100.1001 0x8100.1
default:default
pp0.1073741836      61.1.1.1                dualstackuser1@EXAMPLE1.com
default:ASP-1
*                   2041:1:1::/48
*                   2061:1:1:1::/64
pp0.1073741837      23.1.1.3                dualstackuser2@EXAMPLE1.com
default:ASP-1
*                   2001:db8:1:2:5::/64
```

show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name         LS:RI
si-4/0/0.1         192.168.4.1          xyz@example.com   default:default
```

show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface           IP Address/VLAN ID    User Name              LS:RI
si-2/1/0.1073741842 Tunnel-switched       ap@example.com        default:default

si-2/1/0.1073741843 Tunnel-switched       ap@example.com        default:default
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 10.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 10.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT
```

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```


Copyright © 2016, Juniper Networks, Inc.

show subscribers address detail (IPv6)

```
user@host> show subscribers address 10.16.12.137 detail
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 10.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: ifl3720
Agent Remote ID: ifl3720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 10.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
```

```
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@example.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user1@example.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST
```


Copyright © 2016, Juniper Networks, Inc.

```
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 172.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@EXAMPLE1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
```

```
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2001:db8:db2:ffff:1::/64
Logical System: default
Routing Instance: default
```


Copyright © 2016, Juniper Networks, Inc.

```
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user1@example.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
```

```
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@EXAMPLE1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
```

 Copyright © 2016, Juniper Networks, Inc.

```
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
Effective shaping-rate: 31000000k
...
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
```

```
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
```


Copyright © 2016, Juniper Networks, Inc.

```
Interface           IP Address/VLAN ID   User Name         LS:RI
demux0.1073741825     11.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     12.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
```

```
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface           IP Address                              User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```


Copyright © 2016, Juniper Networks, Inc.

show subscribers summary

Syntax

show subscribers summary
<all>
< detail | extensive | terse>
<count>
<physical-interface *physical-interface-name*>
<logical-system *logical-system* pic | port | routing-instance *routing-instance* | slot>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display summary information for subscribers.

Options

all—(Optional) Display full subscriber summary.

detail | extensive | terse—(Optional) Display the specified level of output.

count—(Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type and LS:RI.

pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.

port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.

NOTE: **Due to display limitations, logical system and routing instance output values are truncated when necessary.**

Required Privilege Level

view

Related Documentation

- show subscribers on page 547

List of Sample Output

show subscribers summary on page 567

show subscribers summary all on page 567
show subscribers summary physical-interface on page 567
show subscribers summary physical-interface pic on page 568
show subscribers summary physical-interface port on page 568
show subscribers summary physical-interface slot on page 568
show subscribers summary pic on page 568
show subscribers summary pic (Aggregated Ethernet Interfaces) on page 569
show subscribers summary port on page 569
show subscribers summary slot on page 569
show subscribers summary terse on page 569

Output Fields Table 38 on page 566 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 38: show subscribers summary Output Fields

Field Name	Field Description
Subscribers by State	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.
Subscribers by Client Type	Number of subscribers summarized by client type. Client types can include DHCP, L2TP, PPP, PPPOE, STATIC-INTERFACE, and VLAN. Also displays the total number of subscribers for all client types (Total).
Subscribers by LS:RI	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session. For aggregated Ethernet interfaces, the output of the **summary (pic | port | slot)** options prefixes the interface name with **ae0:**.
Count	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the **summary** option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.
Total Subscribers	Total number of subscribers for all physical interfaces, all PICS, all ports, or all LS:RI slots.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.

 Copyright © 2016, Juniper Networks, Inc.

Sample Output

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
 Init      3
 Configured    2
 Active   183
 Terminating    2
 Terminated    1

 TOTAL        191

Subscribers by Client Type
 DHCP     107
 PPP       76
 VLAN         8

 TOTAL        191
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
 Init         3
 Configured       2
 Active      183
 Terminating       2
 Terminated       1

 TOTAL          191

Subscribers by Client Type
 DHCP      107
 PPP          76
 VLAN          8

 TOTAL         191

Subscribers by LS:RI
 default:default   1
 default:ri1     28
 default:ri2     16
 ls1:default    22
 ls1:riA     38
 ls1:riB     44
 logsysX:routinstY  42

 TOTAL     191
```

show subscribers summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
   Active: 3998
   Total: 3998

Subscribers by Client Type
   DHCP: 3998
```

```
   Total: 3998

Subscribers by LS:RI
   default:default: 3998
   Total: 3998
```

show subscribers summary physical-interface pic

```
user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
Interface           Count
ge-1/0              1000
ge-1/3              1000

Total Subscribers: 2000
```


Copyright © 2016, Juniper Networks, Inc.

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
Interface           Count
ae0: ge-1/0         801
ae0: ge-1/3         801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface           IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824   100                                   default:default
demux0.1073741824     100.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show system subscriber-management summary

Syntax show system subscriber-management summary

Release Information Command introduced in Junos OS Release 11.1.

Description Display complete subscriber management database summary information.

Options **none**—This command has no options.

Required Privilege Level view

Related Documentation

- *show database-replication statistics*
- *show database-replication summary*

List of Sample Output show system subscriber-management summary on page 571

Output Fields Table 39 on page 570 lists the output fields for the **show system subscriber-management summary** command. Output fields are listed in the approximate order in which they appear.

Table 39: show system subscriber-management summary Output Fields

Field Name	Field Description
Graceful Restart	State of graceful Routing Engine switchover (GRES): • **Enabled** • **Disabled**
Mastership	State of the Routing Engine: • **Master** • **Standby**
Database	State of the subscriber management database: • **Available** • **Init** • **Not-available**

 Copyright © 2016, Juniper Networks, Inc.

Table 39: show system subscriber-management summary Output Fields *(continued)*

Field Name	Field Description
Chassisd ISSU State	State of unified ISSU chassis daemon: • ABORT • DAEMON_ISSU_PREPARE • DAEMON_ISSU_PREPARE_DONE • DAEMON_SWITCHOVER_PREPARE • DAEMON_SWITCHOVER_PREPARE_DONE • FRU_ISSU • FRU_ISSU_DONE • IDLE • UNKNOWN
ISSU State	State of unified ISSU aggregate daemon: • ABORT • IDLE • PREPARE • READY • SWITCHOVER_PREPARE • SWITCHOVER_READY • UNKNOWN
ISSU Wait	Amount of time, in seconds, requested by a daemon to perform cleanup. If multiple daemons request time, the displayed value is the highest wait time requested by a daemon.

Sample Output

show system subscriber-management summary

```
user@host> show system subscriber-management summary
General:
    Graceful Restart     Enabled
    Mastership           Master
    Database             Available
    Chassisd ISSU State  DAEMON_ISSU_PREPARE
    ISSU State           PREPARE
    ISSU Wait            198
```

test services l2tp tunnel

Syntax

test services l2tp tunnel user ***user-name***
<password ***user-password*****>**
<tunnel-name ***name*****>**

Release Information

Command introduced in Junos OS Release 11.4.

Description

(MX Series routers only) Test and verify Layer 2 Tunneling Protocol (L2TP) tunnel configurations from the L2TP access concentrator (LAC). The test determines whether the user can be authenticated and tunneled according to the L2TP configuration. The establishment of all tunnels associated with the user is tested. You can optionally specify a particular tunnel to test for the user.

Options

user ***user-name***—Name of the user under test. You must use an existing configured username, although it can be created solely for testing a tunnel configuration.

password ***user-password***—(Optional) Authentication password for the specified user. If you omit this option, the test generates a dummy password—*testpass*—for the user.

tunnel-name ***name***—(Optional) Name of a tunnel to test.

Required Privilege Level

view

Related Documentation

- Testing L2TP Tunnel Configurations from the LAC on page 212

List of Sample Output

test services l2tp tunnel (User authentication fails) on page 573
test services l2tp tunnel (Multiple tunnels tested) on page 573
test services l2tp tunnel (Specific tunnel tested) on page 573

Output Fields

Table 40 on page 572 lists the output fields for the **test services l2tp tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 40: test services l2tp tunnel Output Fields

Field Name	Field Description
Tunnel-name	Name of the tunnel as configured in the local tunnel profile.
Tunnel-peer	IP address of the tunnel's remote peer (the L2TP network server [LNS]).
Logical-System	Logical system in which the tunnel is created.
Routing-Instance	Routing instance in which the tunnel is created.
Status	Status of the tunnel.

Copyright © 2016, Juniper Networks, Inc.

Sample Output

test services l2tp tunnel (User authentication fails)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication failed
```

test services l2tp tunnel (Multiple tunnels tested)

```
user@host> test services l2tp tunnel user testuser@example.com
Subscriber: testuser@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  test1tunnel    192.168.2.3      default         default         Up
  test2tunnel    172.24.3.3       default         default       Peer unresponsive

  test3tunnel    172.24.5.1       default           test          Up
```

test services l2tp tunnel (Specific tunnel tested)

```
user@host> test services l2tp tunnel user testuser@example.com tunnel-name test1tunnel
Subscriber: testuser@example.com, authentication success, l2tp tunneled
  Tunnel-name   Tunnel-peer   Logical-System  Routing-Instance Status
  test1tunnel    192.168.2.3      default         default         Up
```

Copyright © 2016, Juniper Networks, Inc.

PART 2

Index

Copyright © 2016, Juniper Networks, Inc.

Index

Symbols

A

B

C

D

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

Copyright © 2016, Juniper Networks, Inc.

K

L

Copyright © 2016, Juniper Networks, Inc.

M

Copyright © 2016, Juniper Networks, Inc.

N

O

P

Copyright © 2016, Juniper Networks, Inc.

Q

R

Copyright © 2016, Juniper Networks, Inc.

S

Copyright © 2016, Juniper Networks, Inc.

T

Copyright © 2016, Juniper Networks, Inc.

U

V

W

Copyright © 2016, Juniper Networks, Inc.