

Junos[®] OS

Layer 2 Network Access Protocols User Guide for Routing Devices

Published
2020-03-16

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, 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 Layer 2 Network Access Protocols User Guide for Routing Devices
Copyright © 2020 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 <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation | ix

Documentation and Release Notes | ix

Using the Examples in This Manual | ix

 Merging a Full Example | x

 Merging a Snippet | x

Documentation Conventions | xi

Documentation Feedback | xiv

Requesting Technical Support | xiv

 Self-Help Online Tools and Resources | xv

 Creating a Service Request with JTAC | xv

1

Overview

Network Access Configuration Overview | 2

Network Access Configuration Overview | 2

2

Configuration

Configuring PPP and L2TP | 4

Configuring the PPP Authentication Protocol | 5

Configuring the PPP Password Authentication Protocol On a Logical Interface | 6

 Understanding PPP Password Authentication Protocol | 6

 Configuring the PPP Password Authentication Protocol On a Logical Interface | 7

Configuring the PPP Protocol Field Compression | 9

Monitoring a PPP Session | 11

Tracing Operations of the pppd Process | 11

Example: Configuring PPP CHAP | 13

Example: Configuring CHAP Authentication with RADIUS | 14

Configuring the Maximum Number of LCP Configure-Requests to be Sent | 17

Configuring the Maximum Number of NCP Configure-Requests to be Sent | 18

Configuring L2TP for Enabling PPP Tunneling Within a Network | 18

Defining the Minimum L2TP Configuration | 20

Configuring the Address Pool for L2TP Network Server IP Address Allocation | 21

Example: Configuring an Address-Assignment Pool	23
Configuring the Group Profile for Defining L2TP Attributes	24
Configuring L2TP for a Group Profile	25
Configuring the PPP Attributes for a Group Profile	25
Example: PPP MP for L2TP	27
Configuring Access Profiles for L2TP or PPP Parameters	28
Configuring the Access Profile	28
Configuring the L2TP Properties for a Profile	29
Configuring the PPP Properties for a Profile	29
Configuring the Authentication Order	30
Configuring the Accounting Order	31
Example: Access Profile Configuration	31
Configuring an IKE Access Profile	33
Configuring the L2TP Client	35
Example: Defining the Default Tunnel Client	35
Example: Defining the User Group Profile	36
Configuring the CHAP Secret for an L2TP Profile	37
Example: Configuring L2TP PPP CHAP	38
Referencing the Group Profile from the L2TP Profile	38
Configuring L2TP Properties for a Client-Specific Profile	39
Example: PPP MP for L2TP	40
Example: L2TP Multilink PPP Support on Shared Interfaces	41
Configuring the PAP Password for an L2TP Profile	43
Example: Configuring PAP for an L2TP Profile	43
Configuring PPP Properties for a Client-Specific Profile	44
Applying a Configured PPP Group Profile to a Tunnel	46
Example: Applying a User Group Profile on the M7i or M10i Router	47
Example: Configuring L2TP	48
Supported PPP Interface Standards on ACX Series	51
Configuring PPP Address and Control Field Compression	52
Configuring the PPP Restart Timers	54
Configuring PPP CHAP Authentication	55

Configuring the PPP Clear Loop Detected Timer | 55

Configuring Dynamic Profiles for PPP | 56

Configuring the PPP Challenge Handshake Authentication Protocol | 57

PPP Challenge Handshake Authentication Protocol | 57

Configuring the PPP Challenge Handshake Authentication Protocol | 57

Displaying the Configured PPP Challenge Handshake Authentication Protocol | 59

Configuring the PPP Password Authentication Protocol On a Physical Interface | 60

Understanding PPP Password Authentication Protocol | 60

Configuring the PPP Password Authentication Protocol On a Physical Interface | 61

Configuring the PPP Password Authentication Protocol On a Logical Interface | 62

PPP Encapsulation on ACX Series Routers | 64

Configuring Interface Encapsulation on Physical Interfaces in ACX Series | 66

Configuring the Encapsulation on a Physical Interface | 66

Encapsulation Capabilities | 68

Example: Configuring the Encapsulation on a Physical Interface | 69

Configuring RADIUS Authentication for L2TP | 71

Configuring RADIUS Authentication for L2TP | 71

RADIUS Attributes for L2TP | 73

RADIUS Local Loopback Interface Attribute for L2TP Overview | 77

Example: Configuring RADIUS Authentication for L2TP | 78

Configuring the RADIUS Disconnect Server for L2TP | 80

Configuring RADIUS Authentication for an L2TP Client and Profile | 81

Example: Configuring RADIUS Authentication for an L2TP Profile | 82

Example: Configuring RADIUS-Based Subscriber Authentication and Accounting | 83

Configuring Subscriber Session Timeout Options | 86

Configuring MLPPP | 88

Understanding MLPPP Bundles on ACX Series Routers | 88

Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers | 90

Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 96

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces | 97

Configuring MRRU on Multilink and Link Services Logical Interfaces | 98

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces | 100

Configuring Multiclass MLPPP on LSQ Interfaces | 100

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on ACX Series | 102

Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP | 106

Example: Configuring an MLPPP Bundle on ACX Series | 109

Configuration Statements | 112

Access Configuration Statements | 114

accounting (Access Profile) | 119

accounting-order | 120

accounting-stop-on-access-deny | 121

accounting-stop-on-failure | 122

address-assignment (Address-Assignment Pools) | 123

address-pool | 126

attributes (RADIUS Attributes) | 128

authentication-order | 131

cell-overhead | 133

circuit-type (DHCP Local Server) | 134

client | 136

compression (PPP Properties) | 140

dead-peer-detection | 141

default-chap-secret | 143

default-pap-password | 144

dhcp-attributes (Address-Assignment Pools) | 145

encapsulation-overhead | 152

exclude (RADIUS Attributes) | 153

framed-pool | 161

group-profile (Group Profile) | 162

host (Address-Assignment Pools) | 164

idle-timeout (Access) | 165

ike (Access Profile) | 167

immediate-update | 169

interface-description-format | 170

interface-id | 172

keepalive | 173

keepalive-retries | 174

l2tp (Group Profile) | 175

l2tp (Profile) | 176

lcp-renegotiation | 181

local-chap | 182

maximum-sessions-per-tunnel | 183

multilink | 184

nas-port-extended-format | 186

network | 188

option-82 (Address-Assignment Pools) | 189

option-match | 190

options (Access Profile) | 191

order | 201

pool (Address-Assignment Pools) | 202

ppp (Group Profile) | 204

ppp (Profile) | 207

primary-dns | 208

primary-wins | 209

profile (Access) | 210

radius (Access Profile) | 217

radius-disconnect | 221

radius-disconnect-port | 222

radius-server | 223

range (Address-Assignment Pools) | 229

revert-interval (Access) | 230

secondary-dns | 231

secondary-wins | 232

secret (RADIUS) | 233

session-options | 234

statistics (Access Profile) | 237

update-interval | 238

Administration

Administrative Commands | 241

clear network-access aaa statistics | 242

clear network-access aaa subscriber | 247

clear services l2tp session | 249

clear services l2tp tunnel statistics | 252

show services l2tp radius | 254

Monitoring Commands | 259

show services l2tp session | 260

show services l2tp radius | 271

show services l2tp summary | 276

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | ix
- Using the Examples in This Manual | ix
- Documentation Conventions | xi
- Documentation Feedback | xiv
- Requesting Technical Support | xiv

Use this guide to configure common Layer 2 protocols.

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 <https://www.juniper.net/documentation/>.

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 <https://www.juniper.net/books>.

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.xml; }
```

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

```
[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 [CLI Explorer](#).

Documentation Conventions

[Table 1 on page xii](#) 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 xii 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
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

Convention	Description	Examples
<i>Italic text like this</i>	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 <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">• 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 <i>metric</i>>;
(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 (<i>string1</i> <i>string2</i> <i>string3</i>)
# (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 [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		


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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (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 so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide 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 Juniper Care or Partner Support Services 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- 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 <https://support.juniper.net/support/requesting-support/>.

1

PART

Overview

Network Access Configuration Overview | 2

Network Access Configuration Overview

IN THIS CHAPTER

- [Network Access Configuration Overview | 2](#)

Network Access Configuration Overview

The Junos operating system (Junos OS) enables you to configure network access features for the device at the **[edit access]** hierarchy level. This includes Layer 2 Tunneling Protocol (L2TP), Point-to-Point Protocol (PPP), and Subscriber Access configuration.

The PPP is an encapsulation protocol for transporting IP traffic across point-to-point links. For M7i, M10i, and M120 routers, you can configure L2TP tunneling security services on an Adaptive Services or a MultiServices Physical Interface Card (PIC).

The L2TP protocol allows PPP) to be tunneled within a network.

For a complete hierarchy of access configuration statements, see [“Access Configuration Statements” on page 114](#).

For information about configuring Subscriber Access, see *Broadband Subscriber Sessions User Guide*

RELATED DOCUMENTATION

[Access Configuration Statements | 114](#)

[Configuring the PPP Authentication Protocol | 5](#)

[Configuring L2TP for Enabling PPP Tunneling Within a Network | 18](#)

[Defining the Minimum L2TP Configuration | 20](#)

[Configuring the Address Pool for L2TP Network Server IP Address Allocation | 21](#)

[Configuring the Group Profile for Defining L2TP Attributes | 24](#)

[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

2

PART

Configuration

Configuring PPP and L2TP | **4**

Configuring RADIUS Authentication for L2TP | **71**

Configuring MLPPP | **88**

Configuration Statements | **112**

Configuring PPP and L2TP

IN THIS CHAPTER

- Configuring the PPP Authentication Protocol | 5
- Configuring the PPP Password Authentication Protocol On a Logical Interface | 6
- Configuring the PPP Protocol Field Compression | 9
- Monitoring a PPP Session | 11
- Tracing Operations of the pppd Process | 11
- Example: Configuring PPP CHAP | 13
- Example: Configuring CHAP Authentication with RADIUS | 14
- Configuring the Maximum Number of LCP Configure-Requests to be Sent | 17
- Configuring the Maximum Number of NCP Configure-Requests to be Sent | 18
- Configuring L2TP for Enabling PPP Tunneling Within a Network | 18
- Defining the Minimum L2TP Configuration | 20
- Configuring the Address Pool for L2TP Network Server IP Address Allocation | 21
- Example: Configuring an Address-Assignment Pool | 23
- Configuring the Group Profile for Defining L2TP Attributes | 24
- Example: PPP MP for L2TP | 27
- Configuring Access Profiles for L2TP or PPP Parameters | 28
- Configuring an IKE Access Profile | 33
- Configuring the L2TP Client | 35
- Example: Defining the Default Tunnel Client | 35
- Example: Defining the User Group Profile | 36
- Configuring the CHAP Secret for an L2TP Profile | 37
- Example: Configuring L2TP PPP CHAP | 38
- Referencing the Group Profile from the L2TP Profile | 38
- Configuring L2TP Properties for a Client-Specific Profile | 39
- Example: PPP MP for L2TP | 40
- Example: L2TP Multilink PPP Support on Shared Interfaces | 41
- Configuring the PAP Password for an L2TP Profile | 43
- Example: Configuring PAP for an L2TP Profile | 43

- [Configuring PPP Properties for a Client-Specific Profile | 44](#)
- [Applying a Configured PPP Group Profile to a Tunnel | 46](#)
- [Example: Applying a User Group Profile on the M7i or M10i Router | 47](#)
- [Example: Configuring L2TP | 48](#)
- [Supported PPP Interface Standards on ACX Series | 51](#)
- [Configuring PPP Address and Control Field Compression | 52](#)
- [Configuring the PPP Restart Timers | 54](#)
- [Configuring PPP CHAP Authentication | 55](#)
- [Configuring the PPP Clear Loop Detected Timer | 55](#)
- [Configuring Dynamic Profiles for PPP | 56](#)
- [Configuring the PPP Challenge Handshake Authentication Protocol | 57](#)
- [Configuring the PPP Password Authentication Protocol On a Physical Interface | 60](#)
- [PPP Encapsulation on ACX Series Routers | 64](#)
- [Configuring Interface Encapsulation on Physical Interfaces in ACX Series | 66](#)

Configuring the PPP Authentication Protocol

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. To configure PPP, you can configure the Challenge Handshake Authentication Protocol (CHAP). CHAP allows each end of a PPP link to authenticate its peer, as defined in RFC 1994. The authenticator sends its peer a randomly-generated challenge that the peer must encrypt using a one-way hash; the peer must then respond with that encrypted result. The key to the hash is a secret known only to the authenticator and authenticated. When the response is received, the authenticator compares its calculated result with the peer's response. If they match, the peer is authenticated.

Each end of the link identifies itself to its peer by including its name in the CHAP challenge and response packets it sends to the peer. This name defaults to the local hostname, or you can explicitly set it using the **local-name** option. When a host receives a CHAP challenge or CHAP response packet on a particular interface, it uses the peer identity to look up the CHAP secret key to use.

To configure CHAP, include the **profile** statement at the **[edit access]** hierarchy level:

```
[edit access]
  profile profile-name {
    client client-name chap-secret chap-secret;
  }
```

Then reference the CHAP profile name at the **[edit interfaces]** hierarchy level.

You can configure multiple CHAP profiles, and configure multiple clients for each profile.

Definitions:

- **profile** is the mapping between peer identifiers and CHAP secret keys. The identity of the peer contained in the CHAP challenge or response queries the profile for the secret key to use.
- **client** is the peer identity.
- **chap-secret** is the secret key associated with that peer.

RELATED DOCUMENTATION

[Example: Configuring PPP CHAP | 13](#)

[Example: Configuring CHAP Authentication with RADIUS | 14](#)

Configuring the PPP Password Authentication Protocol On a Logical Interface

IN THIS SECTION

- [Understanding PPP Password Authentication Protocol | 6](#)
- [Configuring the PPP Password Authentication Protocol On a Logical Interface | 7](#)

Understanding PPP Password Authentication Protocol

The Password Authentication Protocol (PAP) provides 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.

For interfaces with PPP encapsulation, you can configure interfaces to support the Password Authentication Protocol (PAP), as defined in RFC 1334, *PAP Authentication Protocols*. If authentication is configured, the PPP link negotiates using CHAP or PAP protocol for authentication during the Link Control Protocol (LCP)

negotiation phase. PAP is only performed after the link establishment phase (LCP up) portion of the authentication phase.

During authentication, the PPP link sends a PAP authentication-request packet to the peer with an ID and password. The authentication-request packet is sent every 2 seconds, similar to the CHAP challenge, until a response is received (acknowledgment packet, nonacknowledgment packet). If an acknowledgment packet is received, the PPP link transitions to the next state, the network phase. If a nonacknowledgment packet is received, an LCP terminate request is sent, and the PPP link goes back to the link establishment phase. If no response is received, and an optional retry counter is set to **true**, a new request acknowledgment packet is resent. If the retry counter expires, the PPP link transitions to the LCP negotiate phase.

You can configure the PPP link with PAP in passive mode. By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

Configuring the PPP Password Authentication Protocol On a Logical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password. If no matching access profile is found for the PAP authentication request that was received by the interface, the optionally configured default PAP password is used.

To configure the PPP password authentication protocol, on each logical interface with PPP encapsulation, perform the following steps.

1. The default PAP password is used when no matching PAP access profile exists, or if the PAP access profile name changes during PPP link negotiation. To configure the default PAP password, include the **pap-password** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]
user@host# set default-pap-password password
```

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
user@host# set local-name name
```

NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]  
user@host# set local-password password
```

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]  
user@host# set passive
```

NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

SEE ALSO

| [Configuring the PPP Authentication Protocol](#) | 5

Configuring the PPP Protocol Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure protocol field compression. By default, the protocol field is not compressed. This means PPP-encapsulated packets are transmitted with a two-byte protocol field. For example, IPv4 packets are transmitted with the protocol field set to 0x0021, and MPLS packets are transmitted with the protocol field set to 0x0281.

For all protocols with identifiers in the range 0x0000 through 0x00ff, you can configure the router to compress the protocol field to one byte, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. Protocol field compression (PFC) allows you to conserve bandwidth by transmitting less data.

NOTE: The protocol field cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure PFC:

1. In configuration mode, go to the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

```
[edit ]
user@host# edit interfaces interface-name ppp-options
```

2. Include the **compression** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level, and specify **pfc**.

```
[edit interfaces interface-name ppp-options]
compression pfc;
```

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **link flags** field for the physical interface. Successfully negotiated options are displayed in the **flags** field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.

```
user@router# run show interfaces so-0/1/1
```

```
Physical interface: so-0/1/1, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16,
  Payload scrambler: Enabled
```



```

Device flags      : Present Running
Interface flags: Point-To-Point SNMP-Traps 16384
Link flags       : No-Keepalives ACFC PFC
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Not-configured
CoS queues      : 4 supported
Last flapped    : 2004-12-29 10:49:32 PST (00:18:35 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
SONET alarms    : None
SONET defects   : None
Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
  Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
  Protocol inet, MTU: 4470
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255

```

This configuration causes the local router to try to negotiate PFC with its peer. If PFC is successfully negotiated, the local router sends packets with compressed protocol fields. When you include the **compression pfc** statement in the configuration, the PPP session restarts, and the local router sends the PFC option in the LCP Configure-Request packet. The PFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then PFC is negotiated. If PFC is successfully negotiated, the local router can receive packets with either 2-byte (uncompressed) or 1-byte (compressed) protocol fields.

RELATED DOCUMENTATION

ppp-options

[compression](#) | 140

pfc

Monitoring a PPP Session

You can monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged by default to `/var/log/pppd`, or to the file specified in the **traceoptions** statement.

To monitor a PPP session:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

```
[edit ]
user@host# edit protocols ppp
```

2. Include the **monitor-session** statement.

```
[edit protocols ppp]
user@host# monitor-session (interface-name | all);
```

When monitoring is configured, the operational mode commands **show ppp summary** and **show ppp interface** display a **Monitored** flag in the **Session flags** column or line.

RELATED DOCUMENTATION

| *monitor-session*

Tracing Operations of the pppd Process

You can trace the operations of the router's pppd process.

To trace the router's pppd process:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

```
[edit ]
user@host# edit protocols ppp
```

2. Include the **traceoptions** statement.

```
[edit protocols ppp]
```

```

traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
    flag flag;
    level severity-level;
    no-remote-trace;
}

```

- To specify more than one tracing operation, include multiple **flag** statements.

You can specify the following flags in the **traceoptions** statement:

- **access**—Trace access code
- **address-pool**—Trace address pool code
- **all**—Trace all areas of code
- **auth**—Trace authentication code
- **chap**—Trace challenge handshake authentication protocol code
- **ci**—Trace CI code
- **config**—Trace configuration code
- **ifdb**—Trace interface database code
- **lcp**—Trace LCP state machine code
- **memory**—Trace memory management code
- **message**—Trace message processing code
- **mlppp**—Trace multilink point-to-point protocol code
- **ncp**—Trace NCP state machine code
- **pap**—Trace password authentication protocol code
- **ppp**—Trace PPP protocol processing code
- **radius**—Trace RADIUS processing code
- **redundancy**—Trace redundancy code
- **rtsock**—Trace routing socket code
- **session**—Trace session management code
- **signal**—Trace signal handling code

- **timer**—Trace timer code
- **ui**—Trace user interface code

RELATED DOCUMENTATION

traceoptions

Example: Configuring PPP CHAP

Configuration

The following example shows how to configure the profile **pe-A-ppp-clients** at the **[edit access]** hierarchy level; then reference it at the **[edit interfaces]** hierarchy level:

CLI Quick Configuration

```
[edit]
access {
  profile pe-A-ppp-clients {
    client cpe-1 chap-secret "$ABC123";
    # SECRET-DATA
    client cpe-2 chap-secret "$ABC123";
    # SECRET-DATA
  }
}
interfaces {
  so-1/1/1 {
    encapsulation ppp;
    ppp-options {
      chap {
        access-profile pe-A-ppp-clients;
        local-name "pe-A-so-1/1/1";
      }
    }
  }
  so-1/1/2 {
    encapsulation ppp;
    ppp-options {
      chap {
        passive;
      }
    }
  }
}
```

```

        access-profile pe-A-ppp-clients;
        local-name "pe-A-so-1/1/2";
    }
}
}
}

```

RELATED DOCUMENTATION

[Configuring the PPP Authentication Protocol](#) | 5

Example: Configuring CHAP Authentication with RADIUS

Configuration

You can send RADIUS messages through a routing instance to customer RADIUS servers in a private network. To configure the routing instance to send packets to a RADIUS server, include the **routing-instance** statement at the **[edit access profile profile-name radius-server]** hierarchy level and apply the profile to an interface with the **access-profile** statement at the **[edit interfaces interface-name unit logical-unit-number ppp-options chap]** hierarchy level.

In this example, PPP peers of interfaces **at-0/0/0.0** and **at-0/0/0.1** are authenticated by a RADIUS server reachable via routing instance **A**. PPP peers of interfaces **at-0/0/0.2** and **at-0/0/0.3** are authenticated by a RADIUS server reachable via routing instance **B**.

For more information about RADIUS authentication, see *Configuring RADIUS Server Authentication*.

CLI Quick Configuration

```

system {
    radius-server {
        1.1.1.1 secret $ABC123;
        2.2.2.2 secret $ABC123;
    }
}
routing-instances {
    A {
        instance-type vrf;
        ...
    }
}

```

```

B {
    instance-type vrf;
    ...
}
}
access {
    profile A-PPP-clients {
        authentication-order radius;
        radius-server {
            3.3.3.3 {
                port 3333;
                secret "$ABC123"; # # SECRET-DATA
                timeout 3;
                retry 3;
                source-address 99.99.99.99;
                routing-instance A;
            }
            4.4.4.4 {
                routing-instance A;
                secret $ABC123;
            }
        }
    }
    profile B-PPP-clients {
        authentication-order radius;
        radius-server {
            5.5.5.5 {
                routing-instance B;
                secret $ABC123;
            }
            6.6.6.6 {
                routing-instance B;
                secret $ABC123;
            }
        }
    }
}
interfaces {
    at-0/0/0 {
        atm-options {
            vpi 0;
        }
        unit 0 {
            encapsulation atm-ppp-llc;

```

```

ppp-options {
    chap {
        access-profile A-PPP-clients;
    }
}
keepalives {
    interval 20;
    up-count 5;
    down-count 5;
}
vci 0.128;
family inet {
    address 21.21.21.21/32 {
        destination 21.21.21.22;
    }
}
}
unit 1 {
    encapsulation atm-ppp-llc;
    ...
    ppp-options {
        chap {
            access-profile A-PPP-clients;
        }
    }
    ...
}
unit 2 {
    encapsulation atm-ppp-llc;
    ...
    ppp-options {
        chap {
            access-profile B-PPP-clients;
        }
    }
    ...
}
unit 3 {
    encapsulation atm-ppp-llc;
    ...
    ppp-options {
        chap {
            access-profile B-PPP-clients;
        }
    }
}

```

```

    }
    ...
  }
  ...
}
...
}

```

Users who log in to the router with telnet or SSH connections are authenticated by the RADIUS server **1.1.1.1**. The backup RADIUS server for these users is **2.2.2.2**.

Each profile may contain one or more backup RADIUS servers. In this example, PPP peers are CHAP authenticated by the RADIUS server **3.3.3.3** (with **4.4.4.4** as the backup server) or RADIUS server **5.5.5.5** (with **6.6.6.6** as the backup server).

RELATED DOCUMENTATION

[Configuring the Authentication Order | 30](#)

[Example: Configuring PPP CHAP | 13](#)

[Configuring the PPP Authentication Protocol | 5](#)

Configuring the Maximum Number of LCP Configure-Requests to be Sent

Link Control Protocol (LCP) Configure-Request is used to establish a link. You can configure the maximum number of LCP Configure-Requests to send. The router stops sending LCP Configure-Requests after the specified maximum number is sent. To configure the LCP Configure-Request maximum, use the **lcp-max-conf-req** statement at the **[edit interfaces *interface-name* unit *number* ppp-options]** hierarchy level. The **number** range is from 0 to 65,535; where 0 specifies no limit and the LCP Configure-Request is sent indefinitely. The default is 254.

RELATED DOCUMENTATION

[lcp-max-conf-req](#)

Configuring the Maximum Number of NCP Configure-Requests to be Sent

Network Control Protocol (NCP) Configure-Request is used to establish a link. You can configure the maximum number of NCP Configure-Requests to send. The router stops sending NCP Configure-Requests after the specified maximum number is sent. To configure the NCP Configure-Request maximum, use the **ncp-max-conf-req** statement at the **[edit interfaces *interface-name* unit *number* ppp-options]** hierarchy level. The ***number*** range is from 0 to 65,535; where 0 specifies no limit and NCP Configure-Request is sent indefinitely. The default is 254.

RELATED DOCUMENTATION

| [ppp-options](#)

Configuring L2TP for Enabling PPP Tunneling Within a Network

For M7i and M10i routers, you can configure Layer 2 Tunneling Protocol (L2TP) tunneling security services on an Adaptive Services Physical Interface Card (PIC) or a MultiServices PIC. The L2TP protocol allows Point-to-Point Protocol (PPP) to be tunneled within a network.

NOTE: For information about how to configure L2TP service, see the *Junos OS Services Interfaces Library for Routing Devices* and the *Junos OS Network Interfaces Library for Routing Devices*.

To configure L2TP, include the following statements at the **[edit access]** hierarchy level:

```
[edit access]
address-pool pool-name {
    address address-or-prefix;
    address-range low <lower-limit> high <upper-limit>;
}
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;
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
}
}
profile profile-name {
    authentication-order [ authentication-methods ];
    accounting-order radius;
    client client-name {
        chap-secret chap-secret;
        group-profile profile-name;
        l2tp {
            interface-id interface-id;
            lcp-renegotiation;
            local-chap;
            maximum-sessions-per-tunnel number;
            ppp-authentication (chap | pap);
            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;
    }
}
radius-disconnect-port port-number {
    radius-disconnect {

```

```

    client-address {
        secret password;
    }
}
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
}

```

RELATED DOCUMENTATION

[Defining the Minimum L2TP Configuration | 20](#)

[Configuring RADIUS Authentication for L2TP | 71](#)

Defining the Minimum L2TP Configuration

To define the minimum configuration for the Layer 2 Tunneling Protocol (L2TP), include at least the following statements at the **[edit access]** hierarchy level:

```

[edit access]
address-pool pool-name {
    address address-or-prefix;
    address-range low <lower-limit> high <upper-limit>;
}
profile profile-name {
    authentication-order [ authentication-methods ];
    client client-name {
        chap-secret chap-secret;
        l2tp {
            interface-id interface-id;
            maximum-sessions-per-tunnel number;
            ppp-authentication (chap | pap);
        }
    }
}

```

```

        shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
        framed-ip-address ip-address;
        framed-pool framed-pool;
        interface-id interface-id;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
    }
}
}
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    secret password;
}

```

NOTE: When the L2TP network server (LNS) is configured with RADIUS authentication, the default behavior is to accept the preferred RADIUS-assigned IP address. Previously, the default behavior was to accept and install the nonzero peer IP address received in the Internet Protocol Control Protocol (IPCP) configuration request packet.

RELATED DOCUMENTATION

| [Configuring the Address Pool for L2TP Network Server IP Address Allocation](#) | 21

Configuring the Address Pool for L2TP Network Server IP Address Allocation

With an address pool, you configure an address or address range. When you define an address pool for a client, the L2TP network server (LNS) allocates IP addresses for clients from an address pool. If you do not want to use an address pool, you can specify an IP address by means of the **framed-ip-address** statement at the [edit access profile *profile-name* client *client-name* ppp] hierarchy level. For information about specifying an IP address, see [“Configuring PPP Properties for a Client-Specific Profile” on page 44](#).

NOTE: When an address pool is modified or deleted, all the sessions using that pool are deleted.

To define an address or a range of addresses, include the **address-pool** statement at the **[edit access]** hierarchy level:

```
[edit access]
address-pool pool-name;
```

pool-name is the name assigned to the address pool.

To configure an address, include the **address** statement at the **[edit access address-pool pool-name]** hierarchy level:

```
[edit access address-pool pool-name]
address address-or-prefix;
```

address-or-prefix is one address or a prefix value.

When you specify an address range, it cannot exceed 65,535 IP addresses.

To configure the address range, include the **address-range** statement at the **[edit access address-pool pool-name]** hierarchy level:

```
[edit access address-pool pool-name]
address-range <low lower-limit> <high upper-limit>;
```

- **low lower-limit**—The lower limit of an address range.
- **high upper-limit**—The upper limit of an address range.

NOTE: The address pools for user access and Network Address Translation (NAT) can overlap. When you configure an address pool at the **[edit access address-pool pool-name]** hierarchy level, you can also configure an address pool at the **[edit services nat pool pool-name]** hierarchy level.

RELATED DOCUMENTATION

Example: Configuring an Address-Assignment Pool

Configuration

This example shows an address-assignment pool configuration that creates two pools, one for IPv4 DHCP clients (**isp_1**), and a second pool (**chi-fiber-ra**) that is used for router advertisement.

CLI Quick Configuration

```
[edit access]
address-assignment {
  network-discovery-router-advertisement chi-fiber-ra;
  pool isp_1 {
    family inet {
      network 192.168.0.0/16;
      range southeast {
        low 192.168.102.2 high 192.168.102.254;
      }
      range northeast {
        low 192.168.119.2 high 192.168.119.250;
      }
    }
    host host.example.net {
      hardware-address 00:00:5E:00:53:90;
      ip-address 192.168.44.12;
    }
    dhcp-attributes {
      option-match {
        option-82 {
          circuit-id fiber range northeast;
        }
        option-82 {
          circuit-id cable_net range southeast;
        }
      }
    }
    boot-file boot.client;
    boot-server 192.168.200.100;
    grace-period 3600;
    maximum-lease-time 18000;
    netbios-node-type p-node;
    router 192.168.44.44 192.168.44.45;
```

```

    }
  }
}
pool chi-fiber-ra {
  family inet6 {
    prefix 2001:db8:2008:2009:2010::/48;
    range fiber3 {
      low 2001:db8:2008:2009:2010::1/64;
      high 2001:db8:2008:2009:2010::5/64;
    }
  }
}
}

```

This example creates an IPv4 address-assignment pool named **isp-1**, which contains two named address ranges, **southeast** and **northeast**. The address-assignment pool also contains a static binding for client **host host.example.net**. The ISP_1 pool configuration also includes the **dhcp-attributes** statement, indicating that the pool is used for DHCP clients. If the option 82 **circuit-id** entry matches the string **fiber**, then DHCP assigns the client an address from the **northeast** range. If the option 82 **circuit-id** matches the string **cable_net**, DHCP assigns an address from the **southeast** range.

The second address-assignment pool created in this example is **chi-fiber-ra**. The **neighbor-discovery-router-advertisement** statement at the beginning of the syntax specifies that this named address-assignment pool is used for router advertisement. The syntax at the end of the example configures the address-assignment pool named **chi-fiber-ra**.

Configuring the Group Profile for Defining L2TP Attributes

Optionally, you can configure the group profile to define the Point-to-Point Protocol (PPP) or Layer 2 Tunneling Protocol (L2TP) attributes. Any client referencing the configured group profile inherits all the group profile attributes.

NOTE: The **group-profile** statement overrides the **user-group-profile** statement, which is configured at the **[edit access profile *profile-name*]** hierarchy level. The **profile** statement overrides the attributes configured at the **[edit access group-profile *profile-name*]** hierarchy level. For information about the **user-group-profile** statement, see [“Applying a Configured PPP Group Profile to a Tunnel” on page 46](#).

Tasks for configuring the group profile are:

1. [Configuring L2TP for a Group Profile | 25](#)
2. [Configuring the PPP Attributes for a Group Profile | 25](#)

Configuring L2TP for a Group Profile

To configure the Layer 2 Tunneling Protocol (L2TP) for the group profile, include the following statements at the `[edit access group-profile profile-name l2tp]` hierarchy level:

```
[edit access group-profile profile-name l2tp]
interface-id interface-id;
lcp-renegotiation;
local-chap;
maximum-sessions-per-tunnel number;
```

interface-id is the identifier for the interface representing an L2TP session configured at the `[edit interfaces interface-name unit local-unit-number dial-options]` hierarchy level.

You can configure the LNS so that it renegotiates the link control protocol (LCP) with the PPP client (in the **renegotiation** statement). By default, the PPP client negotiates the LCP with the L2TP access concentrator (LAC). When you do this, the LNS discards the last sent and the last received LCP configuration request attribute value pairs (AVPs) from the LAC; for example, the LCP negotiated between the PPP client and the LAC.

You can configure the Junos OS so that the LNS ignores proxy authentication AVPs from the LAC and reauthenticates the PPP client using a CHAP challenge (in the **local-chap** statement). When you do this, the LNS directly authenticates the PPP client. By default, the PPP client is not reauthenticated by the LNS.

number is the maximum number of sessions per L2TP tunnel.

Configuring the PPP Attributes for a Group Profile

To configure the Point-to-Point Protocol (PPP) attributes for a group profile, include the following statements at the `[edit access group-profile profile-name ppp]` hierarchy level:

```
[edit access group-profile profile-name ppp]
cell-overhead;
encapsulation-overhead bytes;
framed-pool pool-id;
idle-timeout seconds;
interface-id interface-id;
keepalive seconds;
```



```

ppp-options {
  aaa-options aaa-options-name;
  chap;
  ignore-magic-number-mismatch;
  initiate-ncp (ip | ipv6 | dual-stack-passive)
  ipcp-suggest-dns-option;
  mru;
  mtu;
  pap;
  peer-ip-address-optional;
}
primary-dns primary-dns;
primary-wins primary-wins;
secondary-dns secondary-dns;
secondary-wins secondary-wins;

```

The **cell-overhead** statement configures the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

bytes (in the **encapsulation-overhead** statement) configures the number of bytes used as overhead for class-of-service calculations.

pool-id (in the **framed-pool** statement) is the name assigned to the address pool.

seconds (in the **idle-timeout** statement) is the number of seconds a user can remain idle before the session is terminated. By default, idle timeout is set to 0. You can configure this to be a value in the range from 0 through 4,294,967,295.

interface-id (in the **interface-id** statement) is the identifier for the interface representing an L2TP session configured at the **[edit interfaces interface-name unit local-unit-number dial-options]** hierarchy level.

seconds (in the **keepalive** statement) is the time period that must elapse before the Junos OS checks the status of the PPP session by sending an echo request to the peer. For each session, Junos OS sends out three keepalives at 10-second intervals and the session is close if there is no response. By default, the time to send a keepalive message is set to 10 seconds. You configure this to be a value in the range from 0 through 32,767.

primary-dns (in the **primary-dns** statement) is an IP version 4 (IPv4) address.

secondary-dns (in the **secondary-dns** statement) is an IPv4 address.

primary-wins (in the **primary-wins** statement) is an IPv4 address.

secondary-wins (in the **secondary-wins** statement) is an IPv4 address.

SEE ALSO

[Example: PPP MP for L2TP | 27](#)[Defining the Minimum L2TP Configuration | 20](#)[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

Example: PPP MP for L2TP

Configuration

The following example shows how to configure an L2TP and PPP group profile:

CLI Quick Configuration

```
[edit access]
group-profile westcoast_users {
  ppp {
    framed-pool customer_a;
    keepalive 15;
    primary-dns 192.120.65.1;
    secondary-dns 192.120.65.2;
    primary-wins 192.120.65.3;
    secondary-wins 192.120.65.4;
    interface-id west
  }
}
group-profile eastcoast_users {
  ppp {
    framed-pool customer_b;
    keepalive 15;
    primary-dns 192.120.65.5;
    secondary-dns 192.120.65.6;
    primary-wins 192.120.65.7;
    secondary-wins 192.120.65.8;
    interface-id east;
  }
}
group-profile westcoast_tunnel {
  l2tp {
    maximum-sessions-per-tunnel 100;
  }
}
```

```
group-profile east_tunnel {
  l2tp {
    maximum-sessions-per-tunnel 125;
  }
}
```

RELATED DOCUMENTATION

[Configuring the Group Profile for Defining L2TP Attributes | 24](#)

[Defining the Minimum L2TP Configuration | 20](#)

[Referencing the Group Profile from the L2TP Profile | 38](#)

Configuring Access Profiles for L2TP or PPP Parameters

To validate Layer 2 Tunneling Protocol (L2TP) connections and session requests, you set up access profiles by configuring the profile statement at the **[edit access]** hierarchy level. You can configure multiple profiles. You can also configure multiple clients for each profile.

Tasks for configuring the access profile are:

1. [Configuring the Access Profile | 28](#)
2. [Configuring the L2TP Properties for a Profile | 29](#)
3. [Configuring the PPP Properties for a Profile | 29](#)
4. [Configuring the Authentication Order | 30](#)
5. [Configuring the Accounting Order | 31](#)
6. [Example: Access Profile Configuration | 31](#)

Configuring the Access Profile

To configure the profile, include the **profile** statement at the **[edit access]** hierarchy level:

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

profile-name is the name assigned to the profile.

NOTE: The **group-profile** statement overrides the **user-group-profile** statement, which is configured at the **[edit access profile *profile-name*]** hierarchy level. The **profile** statement overrides the attributes configured at the **[edit access group-profile *profile-name*]** hierarchy level. For information about the **user-group-profile** statement, see [“Applying a Configured PPP Group Profile to a Tunnel” on page 46](#).

When you configure a profile, you can only configure either L2TP or PPP parameters. You cannot configure both at the same time.

Configuring the L2TP Properties for a Profile

To configure the Layer 2 Tunneling Protocol (L2TP) properties for a profile, include the following statements at the **[edit access profile *profile-name*]** hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
accounting-order radius;
client client-name {
  group-profile profile-name;
  l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
    ppp-authentication (chap | pap);
    shared-secret shared-secret;
  }
}
user-group-profile profile-name;
```

Configuring the PPP Properties for a Profile

To configure the PPP properties for a profile, include the following statements at the **[edit access profile *profile-name*]** hierarchy level:

```
[edit access profile profile-name]
authentication-order [ authentication-methods ];
client client-name {
  chap-secret chap-secret;
```

```

group-profile profile-name;
pap-password pap-password;
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-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;
}

```

NOTE: When you configure PPP properties for a profile, you typically configure the **chap-secret** statement or **pap-password** statement.

Configuring the Authentication Order

You can configure the order in which the Junos OS tries different authentication methods when authenticating peers. For each access attempt, the software tries the authentication methods in order, from first to last.

To configure the authentication order, include the **authentication-order** statement at the **[edit access profile *profile-name*]** hierarchy level:

```

[edit access profile profile-name]
authentication-order [ authentication-methods ];

```

In ***authentication-methods***, specify one or more of the following in the preferred order, from first tried to last tried:

- **radius**—Verify the client using RADIUS authentication services.
- **password**—Verify the client using the information configured at the **[edit access profile *profile-name* client *client-name*]** hierarchy level.

NOTE: When you configure the authentication methods for L2TP, only the first configured authentication method is used.

For L2TP, RADIUS authentication servers are configured at the **[edit access radius-server]** hierarchy level. For more information about configuring RADIUS authentication servers, see [“Configuring RADIUS Authentication for L2TP” on page 71](#).

If you do not include the **authentication-order** statement, clients are verified by means of **password** authentication.

Configuring the Accounting Order

You can configure RADIUS accounting for an L2TP profile.

With RADIUS accounting enabled, Juniper Networks routers or switches, acting as RADIUS clients, can notify the RADIUS server about user activities such as software logins, configuration changes, and interactive commands. The framework for RADIUS accounting is described in RFC 2866.

To configure RADIUS accounting, include the **accounting-order** statement at the **[edit access profile profile-name]** hierarchy level:

```
[edit access profile profile-name]
  accounting-order radius;
```

When you enable RADIUS accounting for an L2TP profile, it applies to all the clients within that profile. You must enable RADIUS accounting on at least one L2TP profile for the RADIUS authentication server to send accounting stop and start messages.

NOTE: When you enable RADIUS accounting for an L2TP profile, you do not need to configure the **accounting-port** statement at the **[edit access radius-server server-address]** hierarchy level. When you enable RADIUS accounting for an L2TP profile, accounting is triggered on the default port of 1813.

For L2TP, RADIUS authentication servers are configured at the **[edit access radius-server]** hierarchy level.

Example: Access Profile Configuration

The following example shows a configuration of an access profile:

```
[edit access]
profile westcoast_bldg_1 {
  client white {
    chap-secret "$9$3s2690leK8X7VKM7VwgaJn/Ctu1hclv87Ct87";
    # SECRET-DATA
    ppp {
      idle-timeout 22;
      primary-dns 192.120.65.10;
      framed-ip-address 12.12.12.12/32;
    }
    group-profile westcoast_users;
  }
  client blue {
    chap-secret "$9$eq1KWxbwgZUHNdjqmTF3uO1Rhr-dsoJDNd";
    # SECRET-DATA
    group-profile sunnyvale_users;
  }
  authentication-order password;
}
profile westcoast_bldg_1_tunnel {
  client test {
    l2tp {
      shared-secret "$9$r3HKvLg4ZUDkX7JGjif5p0BIRS8LN";
      # SECRET-DATA
      maximum-sessions-per-tunnel 75;
      ppp-authentication chap;
    }
    group-profile westcoast_tunnel;
  }
  client production {
    l2tp {
      shared-secret "$9$R2QErv8X-goGylVwg4jiTz36/t0BEleWFnRh
      rIXbs2aJDHqf3nCP5";
      # SECRET-DATA
      ppp-authentication chap;
    }
    group-profile westcoast_tunnel;
  }
}
```

SEE ALSO

[Defining the Minimum L2TP Configuration](#) | 20

Configuring an IKE Access Profile

An Internet Key Exchange (IKE) access profile is used to negotiate IKE and IPsec security associations with dynamic peers. You can configure only one tunnel profile per service set for all dynamic peers. The configured preshared key in the profile is used for IKE authentication of all dynamic peers terminating in that service set. You can also use the digital certificate method for IKE authentication with dynamic peers. Include the **ike-policy *policy-name*** statement at the **[edit access profile *profile-name* client * ike]** hierarchy level. ***policy-name*** is the name of the IKE policy you define at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level.

The IKE tunnel profile specifies all the information you need to complete the IKE negotiation. Each protocol has its own statement hierarchy within the client statement to configure protocol-specific attribute value pairs, but only one client configuration is allowed for each profile. The following is the configuration hierarchy.

```
[edit access]
profile profile-name {
  client * {
    ike {
      allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
      }
      dead-peer-detection {
        interval seconds
        threshold number
      }
      ike-policy policy-name;
      initiate-dead-peer-detection;
      interface-id string-value;
      ipsec-policy ipsec-policy;
      pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
      reverse-route
    }
  }
}
```

For dynamic peers, the Junos OS supports only IKE main mode with both the preshared key and digital certificate methods. In this mode, an IPv6 or IPv4 address is used to identify a tunnel peer to obtain the

preshared key or digital certificate information. The client value * (wildcard) means that configuration within this profile is valid for all dynamic peers terminating within the service set accessing this profile.

The following statement makes up the IKE profile:

- **allowed-proxy-pair**—During phase 2 IKE negotiation, the remote peer supplies its network address (**remote**) and its peer's network address (**local**). Since multiple dynamic tunnels are authenticated through the same mechanism, this statement must include the list of possible combinations. If the dynamic peer does not present a valid combination, the phase 2 IKE negotiation fails.

By default, **remote 0.0.0.0/0 local 0.0.0.0/0** is used if no values are configured.

- **dead-peer-detection**—Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peer devices. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE) to peers and waiting for DPD acknowledgements (R-U-THERE-ACK). Use the option **interval** to specify the seconds between which messages should be sent. Use the **threshold** option to specify the maximum number of messages (1-10) to be sent.
- **ike-policy**—Name of the IKE policy that defines either the local digital certificate or the preshared key used to authenticate the dynamic peer during IKE negotiation. You must include this statement to use the digital certificate method for IKE authentication with a dynamic peer. You define the IKE policy at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level.
- **initiate-dead-peer-detection**—Detects dead peers on dynamic IPsec tunnels.
- **interface-id**—Interface identifier, a mandatory attribute used to derive the logical service interface information for the session.
- **ipsec-policy**—Name of the IPsec policy that defines the IPsec policy information for the session. You define the IPsec policy at the **[edit services ipsec-vpn ipsec policy *policy-name*]** hierarchy level. If no policy is set, any policy proposed by the dynamic peer is accepted.
- **pre-shared-key**—Key used to authenticate the dynamic peer during IKE phase 1 negotiation. This key is known to both ends through an out-of-band secure mechanism. You can configure the value either in **hexadecimal** or **ascii-text** format. It is a mandatory value.
- **reverse-route** —(M Series and MX Series routers with an AS or MultiServices PIC only) Configure a reverse route for dynamic endpoint IPsec tunnels.

RELATED DOCUMENTATION

Configuring the L2TP Client

To configure the client, include the **client** statement at the **[edit access profile *profile-name*]** hierarchy level:

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

client-name is the peer identity.

For L2TP, you can optionally use the wildcard (*) to define a default tunnel client to authenticate multiple LACs with the same secret and L2TP attributes. If an LAC with a specific name is not defined in the configuration, the wildcard tunnel client authenticates it.

NOTE: The * for the default client configuration applies only to M Series routers. On MX Series routers, use **default** instead. See *Configuring an L2TP Access Profile on the LNS* for more about MX Series routers.

RELATED DOCUMENTATION

[Example: Defining the Default Tunnel Client](#) | 35

Example: Defining the Default Tunnel Client

Configuration

Use the wildcard (*) to define a default tunnel client to authenticate multiple LACs with the same secret:

CLI Quick Configuration

```
[edit access profile profile-name]  
client * {  
  l2tp {  
    interface-id interface1;  
    lcp-renegotiation;  
    local-chap;  
    maximum-sessions-per-tunnel 500;
```

```

    ppp-authentication chap;
    shared-secret "$ABC123";
  }
}

```

For any tunnel client, you can optionally use the user group profile to define default PPP attributes for all users coming in through a tunnel. The user group profile must define PPP attributes. If the user group profile is specified, all users (PPP sessions) use the PPP attributes specified in the user group profile. The PPP attributes specified in the local or RADIUS server take precedence over those specified in the user group profile.

Optionally, you can use a wildcard client to define a user group profile. When you do this, any client entering this tunnel uses the PPP attributes (defined user group profile attributes) as its default PPP attributes.

RELATED DOCUMENTATION

[Configuring the L2TP Client | 35](#)

[Example: Defining the User Group Profile | 36](#)

Example: Defining the User Group Profile

Configuration

Use a wildcard client to define a user group profile:

CLI Quick Configuration

```

[edit access profile profile]
client * {
    user-group-profile user-group-profile1;
}

```

RELATED DOCUMENTATION

[Applying a Configured PPP Group Profile to a Tunnel | 46.](#)

Configuring the CHAP Secret for an L2TP Profile

CHAP allows each end of a PPP link to authenticate its peer, as defined in RFC 1994. The authenticator sends its peer a randomly generated challenge that the peer must encrypt using a one-way hash; the peer must then respond with that encrypted result. The key to the hash is a secret known only to the authenticator and authenticated. When the response is received, the authenticator compares its calculated result with the peer's response. If they match, the peer is authenticated.

Each end of the link identifies itself to its peer by including its name in the CHAP challenge and response packets it sends to the peer. This name defaults to the local hostname, or you can explicitly set it using the **local-name** option. When a host receives a CHAP challenge or CHAP response packet on a particular interface, it uses the peer identity to look up the CHAP secret key to use.

NOTE: When you configure PPP properties for a Layer 2 Tunneling Protocol (L2TP) profile, you typically configure the **chap-secret** statement or **pap-password** statement.

To configure CHAP, include the **profile** statement and specify a profile name at the **[edit access]** hierarchy level:

```
[edit access]
profile profile-name {
  client client-name chap-secret data;
}
```

Then reference the CHAP profile name at the **[edit interfaces interface-name ppp-options chap]** hierarchy level.

You can configure multiple profiles. You can also configure multiple clients for each profile.

profile is the mapping between peer identifiers and CHAP secret keys. The identity of the peer contained in the CHAP challenge or response queries the profile for the secret key to use.

client is the peer identity.

chap-secret secret is the secret key associated with that peer.

RELATED DOCUMENTATION

[Example: Configuring L2TP PPP CHAP](#) | 38

Example: Configuring L2TP PPP CHAP

Configuration

Configure the profile **westcoast_bldg1** at the **[edit access]** hierarchy level, then reference it at the **[edit interfaces]** hierarchy level:

CLI Quick Configuration

```
[edit]
access {
  profile westcoast_bldg1 {
    client cpe-1 chap-secret "$ABC123";
    # SECRET-DATA
    client cpe-2 chap-secret "$ABC123";
    # SECRET-DATA
  }
}
```

RELATED DOCUMENTATION

[Configuring the CHAP Secret for an L2TP Profile](#) | 37

Referencing the Group Profile from the L2TP Profile

You can reference a configured group profile from the L2TP tunnel profile.

To reference the group profile configured at the **[edit access [group-profile](#) *profile-name*]** hierarchy level, include the **group-profile** statement at the **[edit access profile *profile-name* client *client-name*]** hierarchy level:

```
[edit access profile profile-name client client-name]
group-profile profile-name;
```

profile-name references a configured group profile from a PPP user profile.

RELATED DOCUMENTATION

[Example: Defining the User Group Profile | 36](#)

[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Configuring L2TP Properties for a Client-Specific Profile

To define L2TP properties for a client-specific profile, include one or more of the following statements at the **[edit access profile *profile-name* client *client-name* l2tp]** hierarchy level:

NOTE: When you configure the profile, you can configure either L2TP or PPP parameters, but not both at the same time.

```
[edit access profile profile-name client client-name l2tp]
interface-id interface-id;
lcp-renegotiation;
local-chap;
maximum-sessions-per-tunnel number;
multilink {
    drop-timeout milliseconds;
    fragment-threshold bytes;
}
ppp-authentication (chap | pap);
shared-secret shared-secret;
```

interface-id (in the **interface-id** statement) is the identifier for the interface representing an L2TP session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level.

number (in the **maximum-sessions-per-tunnel** statement) is the maximum number of sessions for an L2TP tunnel.

shared-secret (in the **shared-secret** statement) is the shared secret for authenticating the peer.

You can specify PPP authentication (in the **ppp-authentication** statement). By default, the PPP authentication uses CHAP. You can configure this to use Password Authentication Protocol (PAP).

You can configure LNS so it renegotiates LCP with the PPP client (in the **lcp-negotiation** statement). By default, the PPP client negotiates the LCP with the LAC. When you do this, the LNS discards the last sent LCP configuration request and last received LCP configuration request AVPs from the LAC; for example, the LCP negotiated between the PPP client and LAC.

You can configure the Junos OS so that the LNS ignores proxy authentication AVPs from the LAC and reauthenticates the PPP client using a CHAP challenge (in the **local-chap** statement). By default, the PPP client is not reauthenticated by the LNS. When you do this, the LNS directly authenticates the PPP client.

You can configure the PPP MP for L2TP if the PPP sessions that are coming into the LNS from the LAC have multilink PPP negotiated. When you do this, you join multilink bundles based on the endpoint discriminator (in the **multilink** statement).

- **milliseconds** (in the **drop-timeout** statement) specifies the number of milliseconds for the timeout that associated with the first fragment on the reassembly queue. If the timeout expires before all the fragments have been collected, the fragments at the beginning of the reassembly queue are dropped. If the drop timeout is not specified, the Junos OS holds on to the fragments (fragments may still be dropped if the multilink reassembly algorithm determines that another fragment belonging to the packet on a reassembly queue has been lost).

NOTE: The drop timeout and fragmentation threshold for a bundled multilink might belong to different tunnels. The different tunnels might have different drop timeout and fragmentation thresholds. We recommend configuring group profiles instead of profiles when you have L2TP tunnels.

- **bytes** specifies the maximum size of a packet, in bytes (in the **fragment-threshold** statement). If a packet exceeds the fragmentation threshold, the Junos OS fragments it into two or more multilink fragments.

RELATED DOCUMENTATION

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

[Example: PPP MP for L2TP | 40](#)

[Example: L2TP Multilink PPP Support on Shared Interfaces | 41](#)

Example: PPP MP for L2TP

Configuration

Join multilink bundles based on the endpoint discriminator:

CLI Quick Configuration

```
[edit access]
profile tunnel-profile {
  client remote-host {
    l2tp {
      multilink {
        drop-timeout 600;
        fragmentation-threshold 100;
      }
    }
  }
}
```

RELATED DOCUMENTATION

[Referencing the Group Profile from the L2TP Profile | 38](#)

[Example: L2TP Multilink PPP Support on Shared Interfaces | 41](#)

Example: L2TP Multilink PPP Support on Shared Interfaces

Configuration

On M7i and M10i routers, L2TP multilink PPP sessions are supported on both dedicated and shared interfaces. This example shows how to configure many multilink bundles on a single ASP shared interface.

CLI Quick Configuration

```
[edit]
interfaces {
  sp-1/3/0 {
    traceoptions {
      flag all;
    }
    unit 0 {
      family inet;
    }
    unit 20 {
      dial-options {
        l2tp-interface-id test;
        shared;
      }
    }
  }
}
```



```

    }
    family inet;
  }
}
}
access {
  profile t {
    client cholera {
      l2tp {
        interface-id test;
        multilink;
        shared-secret "$9$n8HX6A01RhLvL1R"; # SECRET-DATA
      }
    }
  }
  profile u {
    authentication-order radius;
  }
  radius-server {
    192.168.65.63 {
      port 1812;
      secret "$9$Vyb4ZHkPQ39mf9pORlexNdbgoZUjqP5"; # SECRET-DATA
    }
  }
}
services {
  l2tp {
    tunnel-group 1 {
      tunnel-access-profile t;
      user-access-profile u;
      local-gateway {
        address 10.70.1.1;
      }
      service-interface sp-1/3/0;
    }
    traceoptions {
      flag all;
      debug-level packet-dump;
      filter {
        protocol l2tp;
        protocol ppp;
        protocol radius;
      }
    }
  }
}

```

```
}
}
```

RELATED DOCUMENTATION

[Referencing the Group Profile from the L2TP Profile | 38](#)

Configuring the PAP Password for an L2TP Profile

When you configure PPP properties for an L2TP profile, you typically configure the **chap-secret** statement or **pap-password** statement. For information about how to configure the CHAP secret, see [“Configuring the CHAP Secret for an L2TP Profile” on page 37](#).

To configure the Password Authentication Protocol (PAP) password, include the **pap-password** statement at the **[edit access profile *profile-name* client *client-name*]** hierarchy level:

```
[edit access profile profile-name client client-name]
pap-password pap-password;
```

pap-password is the password for PAP.

RELATED DOCUMENTATION

[Example: Configuring PAP for an L2TP Profile | 43](#)

Example: Configuring PAP for an L2TP Profile

Configuration

The following examples shows you how to configure the password authentication protocol for an L2TP profile:

CLI Quick Configuration

```
[edit access]
```

```

profile sunnyvale_bldg_2 {
  client green {
    pap-password "$ABC123";
    ppp {
      interface-id west;
    }
    group-profile sunnyvale_users;
  }
  client red {
    chap-secret "$ABC123";
    group-profile sunnyvale_users;
  }
  authentication-order radius;
}
profile Sunnyvale_bldg_1_tunnel {
  client test {
    l2tp {
      shared-secret "$ABC123";
      ppp-authentication pap;
    }
  }
}

```

RELATED DOCUMENTATION

[Configuring the PAP Password for an L2TP Profile](#) | 43

Configuring PPP Properties for a Client-Specific Profile

To define PPP properties for a profile, include one or more of the following statements at the **[edit access profile *profile-name* client *client-name* ppp]** hierarchy level.

NOTE: The properties defined in the profile take precedence over the values defined in the group profile.

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

```

cell-overhead;
encapsulation-overhead bytes;
framed-ip-address ip-address;
framed-pool pool-id;
idle-timeout seconds;
interface-id interface-id;
keepalive seconds;
keepalive-retries number-of-retries;
primary-dns primary-dns;
primary-wins primary-wins;
secondary-dns secondary-dns;
secondary-wins secondary-wins;

```

NOTE: When you configure a profile, you can configure either L2TP or PPP parameters, but not both at the same time.

The **cell-overhead** statement configures the session to use ATM-aware egress shaping on the IQ2 PIC.

bytes (in the **encapsulation-overhead** statement) configures the number of bytes used as overhead for class-of-service calculations.

ip-address (in the **framed-ip-address** statement) is the IPv4 prefix.

pool-id (in the **framed-pool** statement) is a configured address pool.

seconds (in the **idle-timeout** statement) is the number of seconds a user can remain idle before the session is terminated. By default, idle timeout is set to 0. You can configure this to be a value in the range from 0 through 4,294,967,295.

interface-id (in the **interface-id** statement) is the identifier for the interface representing an L2TP session configured at the [edit interfaces *interface-name* unit *local-unit-number* dial-options] hierarchy level.

keepalive seconds is the time period that must elapse before the Junos OS checks the status of the PPP session by sending an echo request to the peer. For each session, Junos OS sends a maximum of ten keepalives at 10-second intervals and the session is closed if there is no response. By default, the time to send a keepalive messages is set to 10 seconds. You can configure this to be a value in the range from 0 through 32,767 seconds.

keepalive-retries number-of-retries is the number of retry attempts for checking the keepalive status of a Point-to-Point (PPP) protocol session. Configuring a lower number of retries helps reduce the detection time for PPP client session failures or timeouts if you have configured a **keepalive seconds** value. By default, the number of retries is set to 10 times. You can configure this to be a value in the range from 3 through 32,767 times.

primary-dns (in the **primary-dns** statement) is an IPv4 address.

secondary-dns (in the **secondary-dns** statement) is an IPv4 address.

primary-wins (in the **primary-wins** statement) is an IPv4 address.

secondary-wins (in the **secondary-wins** statement) is an IPv4 address.

RELATED DOCUMENTATION

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Applying a Configured PPP Group Profile to a Tunnel

On Mi7 and M10i routers, you can optionally apply a configured PPP group profile to a tunnel. For any tunnel client, you can use the **user-group-profile** statement to define default PPP attributes for all users coming in through a tunnel. The user group profile must define PPP attributes. If the user group profile is specified, all users (PPP sessions) use the PPP attributes specified in the user group profile.

When a PPP client enters a tunnel, the Junos OS first applies the PPP user group profile attributes and then any PPP attributes from the local or RADIUS server. The PPP attributes defined in the RADIUS or local server take precedence over the attributes defined in the user group profile.

To apply configured PPP attributes to a PPP client, include the **user-group-profile** statement at the **[edit access profile *profile-name* client *client-name*]** hierarchy level:

```
[edit access profile profile-name client client-name]  
  user-group-profile profile-name;
```

profile-name is a PPP group profile configured at the **[edit access group-profile *profile-name*]** hierarchy level. When a client enters this tunnel, it uses the **user-group-profile** attributes as the default attributes.

RELATED DOCUMENTATION

[Example: Applying a User Group Profile on the M7i or M10i Router | 47](#)

[Example: Defining the User Group Profile | 36](#)

Example: Applying a User Group Profile on the M7i or M10i Router

IN THIS SECTION

- [\[xref target has no title\]](#)
- [\[xref target has no title\]](#)
- [Configuration | 47](#)

Configuration

CLI Quick Configuration

The following example shows how to apply a configured PPP group profile to a tunnel:

```
[edit access]
group-profile example_users {
  ppp {
    idle-timeout 100;
  }
}
group-profile example_default_configuration {
  ppp {
    framed-pool customer_b;
    idle-timeout 20;
    interface-id west;
    primary-dns 192.168.65.5;
    secondary-dns 192.168.65.6;
    primary-wins 192.168.65.7;
    secondary-wins 192.168.65.8;
  }
}
profile example_bldg_1_tunnel {
  client test {
    l2tp {
      interface-id west;
      shared-secret "$ABC123";
      # SECRET-DATA
      maximum-sessions-per-tunnel 75;
      ppp-authentication chap;
    }
  }
}
```

```

        user-group-profile example_default_configuration; # Apply default PPP
    }
}
profile example_bldg_1 {
    client white {
        chap-secret "$ABC123";
        # SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.168.65.9;
            framed-ip-address 192.0.2.0/24;
        }
        group-profile example_users; # Reference the west_users group
    }
}

```

RELATED DOCUMENTATION

| [Applying a Configured PPP Group Profile to a Tunnel](#) | 46

Example: Configuring L2TP

Configuration

The following example shows how to configure L2TP:

CLI Quick Configuration

```

[edit]
access {
    address-pool customer_a {
        address 1.1.1.1/32;
    }
    address-pool customer_b {
        address-range low 2.2.2.2 high 2.2.3.2;
    }
    group-profile westcoast_users {
        ppp {
            framed-pool customer_a;
            idle-timeout 15;
        }
    }
}

```

```

        primary-dns 192.120.65.1;
        secondary-dns 192.120.65.2;
        primary-wins 192.120.65.3;
        secondary-wins 192.120.65.4;
        interface-id west;
    }
}
group-profile eastcoast_users {
    ppp {
        framed-pool customer_b;
        idle-timeout 20;
        primary-dns 192.120.65.5;
        secondary-dns 192.120.65.6;
        primary-wins 192.120.65.7;
        secondary-wins 192.120.65.8;
        interface-id east;
    }
}
group-profile westcoast_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 100;
    }
}
group-profile east_tunnel {
    l2tp {
        maximum-sessions-per-tunnel 125;
    }
}
profile westcoast_bldg_1 {
    client white {
        chap-secret "$ABC123";
        # SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.120.65.10;
            framed-ip-address 12.12.12.12/32;
        }
        group-profile westcoast_users;
    }
    client blue {
        chap-secret "$ABC123";
        # SECRET-DATA
        group-profile sunnyvale_users;
    }
}

```



```

    authentication-order password;
}
profile west-coast_bldg_2 {
    client red {
        pap-password "$ABC123";
        # SECRET-DATA
        ppp {
            idle-timeout 22;
            primary-dns 192.120.65.11;
            framed-ip-address 12.12.12.12/32;
        }
        group-profile westcoast_users;
    }
}
profile westcoast_bldg_1_tunnel {
    client test {
        l2tp {
            shared-secret "$ABC123";
            # SECRET-DATA
            maximum-sessions-per-tunnel 75;
            ppp-authentication chap;# The default for PPP authentication is CHAP.
        }
        group-profile westcoast_tunnel;
    }
    client production {
        l2tp {
            shared-secret "$ABC123
            ABC123"; # SECRET-DATA
            ppp-authentication chap;
        }
        group-profile westcoast_tunnel;
    }
}
profile westcoast_bldg_2_tunnel {
    client black {
        l2tp {
            shared-secret "$ABC123
            ABC123";
            # SECRET-DATA
            ppp-authentication pap;
        }
        group-profile westcoast_tunnel;
    }
}

```

```
}
```

RELATED DOCUMENTATION

| [Configuring L2TP for Enabling PPP Tunneling Within a Network](#) | 18

Supported PPP Interface Standards on ACX Series

Junos OS substantially supports the following RFCs, which define standards for Point-to-Point Protocol (PPP) interfaces.

- RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
- RFC 1334, *PPP Authentication Protocols*
- RFC 1661, *The Point-to-Point Protocol (PPP)*

RELATED DOCUMENTATION

| [Accessing Standards Documents on the Internet](#)

Configuring PPP Address and Control Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure compression of the Data Link Layer address and control fields, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. By default, the address and control fields are not compressed. This means PPP-encapsulated packets are transmitted with two 1-byte fields (0xff and 0x03). If you configure address and control field compression (ACFC) and ACFC is successfully negotiated with the local router's peer, the local router transmits packets without these 2 bytes. ACFC allows you to conserve bandwidth by transmitting less data.

On M320, M120, and T Series routers, ACFC is not supported for any ISO family protocols. Do not include the **acfc** statement at the **[edit interfaces *interface-name* ppp-options compression]** hierarchy level when you include the **family iso** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

NOTE: The address and control fields cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure ACFC:

1. In configuration mode, go to the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

```
[edit ]
user@host# edit interfaces interface-name ppp-options
```

2. Include the **compression** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level, and specify **acfc**.

```
[edit interfaces interface-name ppp-options]
compression acfc;
```

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **link flags** field for the physical interface. Successfully negotiated options are displayed in the **flags** field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.

```
user@router# run show interfaces so-0/1/1
```

```

Physical interface: so-0/1/1, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16
  Payload scrambler: Enabled
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags       : No-Keepalives ACFC PFC
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Not-configured
  CoS queues      : 4 supported
  Last flapped    : 2004-12-29 10:49:32 PST (00:18:35 ago)
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)
  SONET alarms     : None
  SONET defects    : None
Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
  Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
  Protocol inet, MTU: 4470
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255

```

This configuration causes the local router to try to negotiate ACFC with its peer. If ACFC is successfully negotiated, the local router sends packets with compressed address and control fields. When you include the **compression acfc** statement in the configuration, the PPP session restarts, and the local router sends the ACFC option in the LCP Configure-Request packet. The ACFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then ACFC is negotiated. If ACFC is successfully negotiated, the local router can receive packets with or without the address and control bytes included.

RELATED DOCUMENTATION

ppp-options

[compression](#) | 140

acfc

Configuring the PPP Restart Timers

You can configure a restart timer for the Link Control Protocol (LCP) and Network Control Protocol (NCP) components of a PPP session. You can configure the LCP restart timer on interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations. You can configure the NCP restart timer on interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces.

To configure the restart timer for the NCP component of a PPP session, include the **ncp-restart-timer** statement, and specify the number of milliseconds.

To configure the restart timer for the LCP component of a PPP session, include the **lcp-restart-timer** statement, and specify the number of milliseconds:

```
lcp-restart-timer milliseconds;
ncp-restart-timer milliseconds;
```

You can include these statements at the following hierarchy levels:

- [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]

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **PPP parameters** field for the physical interface.

```
user@host> run show interfaces t1-0/0/0:1:1.0 detail
```

```
Logical interface t1-0/0/0:1:1.0 (Index 67) (SNMP ifIndex 40)
(Generation 156)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps 0x4000
Encapsulation: PPP
PPP parameters:
  LCP restart timer: 2000 msec
  NCP restart timer: 2000 msec
Protocol inet, MTU: 1500, Generation: 163, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255,
```

Configuring PPP CHAP Authentication

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*. When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer.

For information about configuring CHAP, see [“Configuring the PPP Challenge Handshake Authentication Protocol” on page 57](#).

Configuring the PPP Clear Loop Detected Timer

When a Point-to-Point Protocol (PPP) session detects a loop, the loop detected flag is set. If the flag is not cleared by the protocol after the loopback is cleared, the clear loop detected timer clears the flag after the specified time has elapsed.

To configure the clear loop detected timer for the LCP component of a PPP session, include the **loopback-clear-timer** statement, and specify the number of seconds.

```
loopback-clear-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [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]

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

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* in the *Junos Subscriber Access Configuration Guide*.

For information about using dynamic profiles to authenticate PPP subscribers, see *Configuring Dynamic Authentication for PPP Subscribers*.

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

Configuring the PPP Challenge Handshake Authentication Protocol

IN THIS SECTION

- [PPP Challenge Handshake Authentication Protocol | 57](#)
- [Configuring the PPP Challenge Handshake Authentication Protocol | 57](#)
- [Displaying the Configured PPP Challenge Handshake Authentication Protocol | 59](#)

PPP Challenge Handshake Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer. By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges. To enable CHAP, you must create an access profile, and you must configure the interfaces to use CHAP.

Configuring the PPP Challenge Handshake Authentication Protocol

When you configure an interface to use CHAP, you must assign an access profile to the interface. When an interface receives CHAP challenges and responses, the access profile in the packet is used to look up the shared secret, as defined in RFC 1994. If no matching access profile is found for the CHAP challenge that was received by the interface, the optionally configured default CHAP secret is used. The default CHAP secret is useful if the CHAP name of the peer is unknown, or if the CHAP name changes during PPP link negotiation.

To enable CHAP, you must create an access profile, and you must configure the interfaces to use PAP. For more information on how to configure access profile, see [“Configuring Access Profiles for L2TP or PPP Parameters”](#) on page 28.

To configure the PPP challenge handshake authentication protocol, on each physical interface with PPP encapsulation, perform the following steps.

1. To assign an access profile to an interface, include the **access-profile** statement at the **[edit interfaces interface-name ppp-options chap]** hierarchy level.

```
[edit interfaces interface-name ppp-options chap]
user@host# set access-profile name
```


NOTE: You must include the **access-profile** statement when you configure the CHAP authentication method. If an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped unless a default CHAP secret has been configured.

2. The default CHAP secret is used when no matching CHAP access profile exists, or if the CHAP name changes during PPP link negotiation. To configure a default CHAP secret for an interface, include the **default-chap-secret** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level.

```
[edit interfaces interface-name ppp-options chap]
user@host# set default-chap-secret name
```

3. To configure the name the interface uses in CHAP challenge and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
user@host# set local-name name
```

NOTE:

- The local name is any string from 1 through 32 characters in length, starting with an alphanumeric or underscore character, and including only the following characters:

```
a-z A-Z 0-9 % @ # / \ . _ -
```

- By default, when CHAP is enabled on an interface, the interface uses the router's system hostname as the name sent in CHAP challenge and response packets.

4. You can configure the interface not to challenge its peer, and only respond when challenged. To configure the interface not to challenge its peer, include the **passive** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
user@host# set passive;
```

NOTE: By default, when CHAP is enabled on an interface, the interface always challenges its peer and responds to challenges from its peer.

SEE ALSO

[Configuring the PPP Authentication Protocol | 5](#)

Displaying the Configured PPP Challenge Handshake Authentication Protocol

Purpose

To display the configured PPP CHAP at the **[edit access]** and **[edit interfaces]** hierarchy levels.

- Access profile—**pe-A-ppp-clients**
- default CHAP secret data—**"\$ABC123"**
- hostname for the CHAP challenge and response packets—**"pe-A-so-1/1/1"**
- Interface—**so-1/1/2**

Action

- Run the **show** command at the **[edit access]** hierarchy level.

```
profile pe-A-ppp-clients;
client cpe-1 chap-secret "$ABC123";
    # SECRET-DATA
[edit interfaces so-1/2/0]
encapsulation ppp;
ppp-options {
    chap {
        access-profile pe-A-ppp-clients;
        default-chap-secret "$ABC123";
        local-name "pe-A-so-1/1/1";
    }
}
```

- Run the **show** command at the **[edit interfaces s0-1/1/2]** hierarchy level.

```
ppp-options {
    chap {
        access-profile pe-A-ppp-clients;
```

```

    default-chap-secret "$ABC123";
    local-name "pe-A-so-1/1/2";
  }
}

```

Meaning

The configured CHAP and its associated set options are displayed as expected.

Configuring the PPP Password Authentication Protocol On a Physical Interface

IN THIS SECTION

- [Understanding PPP Password Authentication Protocol | 60](#)
- [Configuring the PPP Password Authentication Protocol On a Physical Interface | 61](#)
- [Configuring the PPP Password Authentication Protocol On a Logical Interface | 62](#)

Understanding PPP Password Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the Password Authentication Protocol (PAP), as defined in RFC 1334, *PAP Authentication Protocols*. If authentication is configured, the PPP link negotiates using CHAP or PAP protocol for authentication during the Link Control Protocol (LCP) negotiation phase. PAP is only performed after the link establishment phase (LCP up) portion of the authentication phase.

During authentication, the PPP link sends a PAP authentication-request packet to the peer with an ID and password. The authentication-request packet is sent every 2 seconds, similar to the CHAP challenge, until a response is received (acknowledgment packet, nonacknowledgment packet). If an acknowledgment packet is received, the PPP link transitions to the next state, the network phase. If a nonacknowledgment packet is received, an LCP terminate request is sent, and the PPP link goes back to the link establishment phase. If no response is received, and an optional retry counter is set to **true**, a new request acknowledgment packet is resent. If the retry counter expires, the PPP link transitions to the LCP negotiate phrase.

You can configure the PPP link with PAP in passive mode. By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP

passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

Configuring the PPP Password Authentication Protocol On a Physical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password.

To enable PAP, you must create an access profile, and you must configure the interfaces to use PAP. For more information on how to configure access profile, see [“Configuring Access Profiles for L2TP or PPP Parameters” on page 28](#).

To configure the PPP password authentication protocol, on each physical interface with PPP encapsulation, perform the following steps.

1. To assign an access profile to an interface, include the **access-profile** statement at the **[edit interfaces interface-name ppp-options pap]** hierarchy level.

```
[edit interfaces interface-name ppp-options pap]
user@host# set access-profile name
```

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces interface-name ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
user@host# set local-name name
```

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces interface-name ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
user@host# set local-password password
```

NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
user@host# set passive
```

NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

SEE ALSO

| [Configuring the PPP Authentication Protocol | 5](#)

Configuring the PPP Password Authentication Protocol On a Logical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password. If no matching access profile is found for the PAP authentication request that was received by the interface, the optionally configured default PAP password is used.

To configure the PPP password authentication protocol, on each logical interface with PPP encapsulation, perform the following steps.

1. To configure the default PAP password, include the **pap-password** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]
user@host# set default-pap-password password
```

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
```

```
user@host# set local-name name
```

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]  
user@host# set local-password password
```

NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]  
user@host# set passive
```

NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

SEE ALSO

| [Configuring the PPP Authentication Protocol](#) | 5

PPP Encapsulation on ACX Series Routers

You can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on ACX Series routers. PPP provides a standard method for transporting multiprotocol datagrams over a point-to-point link. PPP uses the High-Speed Data Link Control (HDLC) protocol for its physical interface and provides a packet-oriented interface for the network-layer protocols.

PPP is supported on the following MICs on ACX Series routers:

- On ACX1000 routers with 8-port built-in T1/E1 TDM MICs.
- On ACX2000, ACX2100, ACX2200, and ACX4000 routers with 16-port built-in T1/E1 TDM MICs.
- On ACX4000 routers with 16-Port Channelized E1/T1 Circuit Emulation MICs.

Starting with Release 12.3X54, you can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP on ACX4000 Series routers

On ACX Series routers, E1, T1, and NxDS0 interfaces support PPP encapsulation.

PPP is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

To configure the encapsulation on a physical interface, include the **encapsulation ppp** statement at the **[edit interfaces interface-name]** hierarchy level.

IP class of service (CoS) is not supported on PPP interfaces. All the traffic is sent to the best effort queue (queue 0) and CoS code points are not processed. Also, fixed classifiers are not supported. Circuit cross-connect (CCC) version of PPP (**ppp-ccc** option) and translational cross-connect (TCC) version of PPP (**ppp-tcc** option) are not supported for configuration with the **encapsulation** statement.

PPP is supported only for IPv4 networks. If you configure PPP encapsulation, you can configure an INET family by including the **family inet** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. MPLS family is not supported on logical interfaces if you configured PPP encapsulation. On interfaces with PPP encapsulation, configure PPP-specific interface properties by including the **ppp-options** statement at the **[edit interfaces interface-name]** hierarchy level. For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP).

For full T1/E1 interfaces on which PPP encapsulation needs to be enabled, create the T1/E1 interfaces out of channelized T1/E1 interfaces (CT1/CE1) by including the **framing** statement at the **[edit chassis fpc fpc-slot pic pic-slot]** hierarchy level:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@host# set framing (t1 | e1);
```

Configure a CT1 port down to a T1 channel. On the CT1 interface, set the **no-partition** option and then set the interface type as T1.

```
[edit interfaces ct1-mpc-slot/mic-slot/port-number]
user@host# set no-partition interface-type t1
```

Configure a CE1 port down to an E1 channel. On the CE1 interface, set the **no-partition** option and then set the interface type as E1.

```
[edit interfaces ce1-mpc-slot/mic-slot/port-number]
user@host# set no-partition interface-type t1
```

For NxDS0 interfaces on which PPP encapsulation needs to be enabled, partition the CE1 and CT1 interfaces by including the **ce1-x/y/z partition *partition-number* timeslots *timeslots* interface-type ds** and **ct1-x/y/z partition *partition-number* timeslots *timeslots* interface-type ds** statements at the **[edit interfaces *interface-name*]** hierarchy level.

The following operational mode commands can be used to view PPP configuration settings and statistical details:

- The **show ppp address-pool** command is used to display PPP address pool information.
- The **show ppp interface** command is used to display PPP session information for an interface.
- The **show ppp statistics** command is used to display PPP session statistics.
- The **show ppp summary** command is used to display summary information about PPP-configured interfaces.
- The **show interfaces e1-fpc/pic/port**, **show interfaces t1-fpc/pic/port**, and **show interfaces ds-fpc/pic/port** commands are used to display the PPP settings of a specific E1, T1, and DS interface, respectively.

RELATED DOCUMENTATION

[Configuring Interface Encapsulation on Physical Interfaces in ACX Series | 66](#)

[encapsulation](#)

[ppp-options](#)

Configuring Interface Encapsulation on Physical Interfaces in ACX Series

IN THIS SECTION

- [Configuring the Encapsulation on a Physical Interface | 66](#)
- [Encapsulation Capabilities | 68](#)

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types. For more information about logical interface encapsulation, see *Configuring Interface Encapsulation on Logical Interfaces*.

This section contains the following topics:

Configuring the Encapsulation on a Physical Interface

By default, PPP is the encapsulation type for physical interfaces. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc | ethernet-ccc |
  ethernet-over-atm | ethernet-tcc | ethernet-vpls | extended-frame-relay-ccc |
  extended-frame-relay-ether-type-tcc | extended-frame-relay-tcc | extended-vlan-ccc | extended-vlan-tcc |
  extended-vlan-vpls | flexible-ethernet-services | flexible-frame-relay | frame-relay | frame-relay-ccc |
  frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-port-ccc | frame-relay-tcc |
  multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc | vlan-vpls);
```

NOTE: ACX Series routers do not support *cisco-hdlc* encapsulation.

The physical interface encapsulation can be one of the following:

- ATM CCC cell relay—Connects two remote virtual circuits or ATM physical interfaces with a label-switched path (LSP). Traffic on the circuit is ATM cells.

For more information, see the *Junos OS Administration Library*.

- ATM PVC—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).
- Ethernet cross-connect—Ethernet interfaces without VLAN tagging can use Ethernet CCC encapsulation. Two related versions are supported:
 - CCC version (**ethernet-ccc**)—Ethernet interfaces with standard Tag Protocol ID (TPID) tagging can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
 - TCC version (**ethernet-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

- VLAN CCC (**vlan-ccc**)—Ethernet interfaces with VLAN tagging enabled can use VLAN CCC encapsulation. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended VLAN cross-connect—Gigabit Ethernet interfaces with VLAN 802.1Q tagging enabled can use extended VLAN cross-connect encapsulation. (Ethernet interfaces with standard TPID tagging can use VLAN CCC encapsulation.) Two related versions of extended VLAN cross-connect are supported:
 - CCC version (**extended-vlan-ccc**)—Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only.
 - TCC version (**extended-vlan-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

NOTE: In ACX Series routers, VPLS is supported only on ACX5048 and ACX5096 routers.

- Ethernet VPLS (**ethernet-vpls**)—Ethernet interfaces with VPLS enabled can use Ethernet VPLS encapsulation. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Ethernet VLAN VPLS (**vlan-vpls**)—Ethernet interfaces with VLAN tagging and VPLS enabled can use Ethernet VLAN VPLS encapsulation. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Extended VLAN VPLS (**extended-vlan-vpls**)—Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled can use Ethernet Extended VLAN VPLS encapsulation. (Ethernet interfaces with standard TPID tagging can use Ethernet VLAN VPLS encapsulation.) Extended Ethernet VLAN VPLS encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Flexible Ethernet services (**flexible-ethernet-services**)—Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) can use flexible Ethernet services encapsulation. Aggregated Ethernet bundles can use this encapsulation type. You use this encapsulation type when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.
- PPP—Defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.

NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

In ACX Series routers, VPLS is supported only on ACX5048 and ACX5096 routers.

Encapsulation Capabilities

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one **unit** statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

Ethernet CCC encapsulation for Ethernet interfaces with standard TPID tagging requires that the physical interface have only a single logical interface. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For Ethernet interfaces in VLAN mode, VLAN IDs are applicable as follows:

- VLAN ID 0 is reserved for tagging the priority of frames.
- For encapsulation type **vlan-ccc**, VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN CCCs.

When you configure Ethernet virtual LAN (VLAN) encapsulation on CCC circuits (by using the **encapsulation vlan-ccc** statement at the **[edit interfaces *interface-name*]** hierarchy level), you can bind a list of VLAN IDs to the interface by using the **vlan-id-list [*vlan-id-numbers*]** statement to configure a CCC for multiple VLANs. Configuring this statement creates a CCC for:

- Each VLAN listed—for example, **vlan-id-list [100 200 300]**
- Each VLAN in a range—for example, **vlan-id-list [100-200]**
- Each VLAN in a list and range combination—for example, **vlan-id-list [50, 100-200, 300]**
- For encapsulation type **vlan-vpls**, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLANs.
- For Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. For interfaces with **flexible-ethernet-services** encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are not reserved.
- For encapsulation types **extended-vlan-ccc** and **extended-vlan-vpls**, all VLAN IDs are valid.

The upper limits for configurable VLAN IDs vary by interface type.

When you configure a TCC encapsulation, some modifications are needed to handle VPN connections over unlike Layer 2 and Layer 2.5 links and terminate the Layer 2 and Layer 2.5 protocol locally.

The router performs the following media-specific change:

- ATM—Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) processing is terminated at the router. Cell relay is not supported. The Junos OS strips all ATM encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to ATM encapsulation.

Example: Configuring the Encapsulation on a Physical Interface

Configure PPP encapsulation on a SONET/SDH interface. The second and third **family** statements allow Intermediate System-to-Intermediate System (IS-IS) and MPLS to run on the interface.

```
[edit interfaces]
so-7/0/0 {
  encapsulation ppp;
  unit 0 {
```

```
point-to-point;  
family inet {  
    address 192.168.1.113/32 {  
        destination 192.168.1.114;  
    }  
}  
family iso;  
family mpls;  
}  
}
```

RELATED DOCUMENTATION

| *Configuring Interface Encapsulation on Logical Interfaces*

Configuring RADIUS Authentication for L2TP

IN THIS CHAPTER

- [Configuring RADIUS Authentication for L2TP | 71](#)
- [RADIUS Attributes for L2TP | 73](#)
- [RADIUS Local Loopback Interface Attribute for L2TP Overview | 77](#)
- [Example: Configuring RADIUS Authentication for L2TP | 78](#)
- [Configuring the RADIUS Disconnect Server for L2TP | 80](#)
- [Configuring RADIUS Authentication for an L2TP Client and Profile | 81](#)
- [Example: Configuring RADIUS Authentication for an L2TP Profile | 82](#)
- [Example: Configuring RADIUS-Based Subscriber Authentication and Accounting | 83](#)
- [Configuring Subscriber Session Timeout Options | 86](#)

Configuring RADIUS Authentication for L2TP

The L2TP network server (LNS) sends RADIUS authentication requests or accounting requests. Authentication requests are sent out to the authentication server port. Accounting requests are sent to the accounting port. To configure RADIUS authentication for L2TP on an M10i or M7i router, include the following statements at the **[edit access]** hierarchy level:

```
[edit access]
radius-server server-address {
  accounting-port port-number;
  port port-number;
  retry attempts;
  routing-instance routing-instance-name;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

NOTE: The RADIUS servers at the **[edit access]** hierarchy level are not used by the network access server process (NASD).

You can specify an accounting port number on which to contact the accounting server (in the **accounting-port** statement). Most RADIUS servers use port number 1813 (as specified in RFC 2866, *Radius Accounting*).

NOTE: If you enable RADIUS accounting at the **[edit access profile profile-name accounting-order]** hierarchy level, accounting is triggered on the default port of 1813 even if you do not specify a value for the **accounting-port** statement.

server-address specifies the address of the RADIUS authentication server (in the **radius-server** statement).

You can specify a port number on which to contact the RADIUS authentication server (in the **port** statement). Most RADIUS servers use port number 1812 (as specified in RFC 2865, *Remote Authentication Dial In User Service [RADIUS]*).

You must specify a password in the **secret** statement. If a password includes spaces, enclose the password in quotation marks. The secret used by the local router must match that used by the RADIUS authentication server.

Optionally, you can specify the amount of time that the local router waits to receive a response from a RADIUS server (in the **timeout** statement) and the number of times that the router attempts to contact a RADIUS authentication server (in the **retry** statement). By default, the router waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds. By default, the router retries connecting to the server three times. You can configure this to be a value in the range from 1 through 30 times. If the maximum number of retries is reached, the radius server is considered dead for 5 minutes (300 seconds).

In the **source-address** statement, specify a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address configured on one of the router interfaces.

To configure multiple RADIUS servers, include multiple **radius-server** statements. For information about how to configure the RADIUS disconnect server for L2TP, see [“Configuring the RADIUS Disconnect Server for L2TP” on page 80](#).

NOTE: When the L2TP network server (LNS) is configured with RADIUS authentication, the default behavior is to accept the preferred RADIUS-assigned IP address. Previously, the default behavior was to accept and install the nonzero peer IP address received by the Internet Protocol Control Protocol (IPCP) configuration request packet.

RELATED DOCUMENTATION

RADIUS Attributes for L2TP 73
Configuring L2TP for Enabling PPP Tunneling Within a Network 18
Configuring the RADIUS Disconnect Server for L2TP 80

RADIUS Attributes for L2TP

Junos OS supports the following types of RADIUS attributes for L2TP:

- Juniper Networks vendor-specific attributes (VSAs)
- Attribute-value pairs (AVPs) defined by the Internet Engineering Task Force (IETF)
- RADIUS accounting stop and start AVPs

Juniper Networks vendor-specific RADIUS attributes are described in RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*. These attributes are encapsulated with the vendor ID set to the Juniper Networks ID number 2636. [Table 3 on page 73](#) lists the Juniper Networks VSAs you can configure for L2TP.

Table 3: Juniper Networks Vendor-Specific RADIUS Attributes for L2TP

Attribute Name	Standard Number	Value
Juniper-Primary-DNS	31	IP address
Juniper-Primary-WINS	32	IP address
Juniper-Secondary-DNS	33	IP address
Juniper-Secondary-WINS	34	IP address
Juniper-Interface-ID	35	String

Table 3: Juniper Networks Vendor-Specific RADIUS Attributes for L2TP (continued)

Attribute Name	Standard Number	Value
Juniper-IP-Pool-Name	36	String
Juniper-Keep-Alive	37	Integer

[Table 4 on page 74](#) lists the IETF RADIUS AVPs supported for L2TP.

Table 4: Supported IETF RADIUS Attributes for L2TP

Attribute Name	Standard Number	Value
User-Name	1	String
User-Password	2	String
CHAP-Password	3	String
NAS-IP-Address	4	IP address
NAS-Port	5	Integer
Service-Type	6	Integer
Framed-Protocol	7	Integer
Framed-IP-Address	8	IP address
Framed-IP-Netmask	9	IP address
Framed-MTU	12	Integer
Framed-Route	22	String
Session-Timeout	27	Integer
Idle-Timeout	28	Integer
Called-Station-ID	30	String
Calling-Station-ID	31	String
CHAP-Challenge	60	String

Table 4: Supported IETF RADIUS Attributes for L2TP (continued)

Attribute Name	Standard Number	Value
NAS-Port-Type	61	Integer
Framed-Pool	88	Integer

[Table 5 on page 75](#) lists the supported RADIUS accounting start AVPs for L2TP.

Table 5: Supported RADIUS Accounting Start Attributes for L2TP

Attribute Name	Standard Number	Value
User-Name	1	String
NAS-IP-Address	4	IP address
NAS-Port	5	Integer
Service-Type	6	Integer
Framed-Protocol	7	Integer
Framed-IP-Address	8	IP address
Called-Station-ID	30	String
Calling-Station-ID	31	String
Acct-Status-Type	40	Integer
Acct-Delay-Time	41	Integer
Acct-Session-ID	44	String
Acct-Authentic	45	Integer
NAS-Port-Type	61	Integer
Tunnel-Client-Endpoint	66	String
Tunnel-Server-Endpoint	67	String
Acct-Tunnel-Connection	68	String

Table 5: Supported RADIUS Accounting Start Attributes for L2TP (continued)

Attribute Name	Standard Number	Value
Tunnel-Client-Auth-ID	90	String
Tunnel-Server-Auth-ID	91	String

[Table 6 on page 76](#) lists the supported RADIUS accounting stop AVPs for L2TP.

Table 6: Supported RADIUS Accounting Stop Attributes for L2TP

Attribute Name	Standard Number	Value
User-Name	1	String
Local-Loopback-Interface	3	String
NAS-IP-Address	4	IP address
NAS-Port	5	Integer
Service-Type	6	Integer
Framed-Protocol	7	Integer
Framed-IP-Address	8	IP address
Called-Station-ID	30	String
Calling-Station-ID	31	String
Acct-Status-Type	40	Integer
Acct-Delay-Time	41	Integer
Acct-Input-Octets	42	Integer
Acct-Output-Octets	43	Integer
Acct-Session-ID	44	String
Acct-Authentic	45	Integer
Acct-Session-Time	46	Integer

Table 6: Supported RADIUS Accounting Stop Attributes for L2TP (*continued*)

Attribute Name	Standard Number	Value
Acct-Input-Packets	47	Integer
Acct-Output-Packets	48	Integer
Acct-Terminate-Cause	49	Integer
Acct-Multi-Session-ID	50	String
Acct-Link-Count	51	Integer
NAS-Port-Type	61	Integer
Tunnel-Client-Endpoint	66	String
Tunnel-Server-Endpoint	67	String
Acct-Tunnel-Connection	68	String
Tunnel-Client-Auth-ID	90	String
Tunnel-Server-Auth-ID	91	String

RELATED DOCUMENTATION

[Example: Configuring RADIUS Authentication for L2TP](#) | 78

RADIUS Local Loopback Interface Attribute for L2TP Overview

You can configure the Local-Loopback-Interface attribute on a RADIUS server to manage multiple LAC devices. This attribute is used as the LAC source address on an LNS tunnel for PPPoE subscribers tunneled over L2TP.

When you use the Tunnel-Client-Endpoint attribute as the LAC source address, you must configure the Tunnel-Client-Endpoint attribute for each MX Series router that uses the same RADIUS server. Starting with this release you can use the Local-Loopback-Interface attribute, which needs to be configured only once. When the LAC initiates an Access-Request message to RADIUS for authentication, RADIUS returns the Local-Loopback-Interface attribute in the Access-Accept message. This attribute contains the name

of the loopback interface, either as a generic interface name such as “lo0” or as a specific name like “lo0.0”. The MX Series router then uses the configured loopback interface IP address as the source address during tunnel negotiation with the LNS.

NOTE: An MX Series router can act as the LAC and use any interface address on it as an L2TP tunnel source address. The source address can be dynamically assigned by RADIUS through the Tunnel-Client-Endpoint or Local-Loopback-Interface attribute. The tunnel source address can be statically configured on the MX Series router by using the L2TP tunnel profile. If RADIUS does not return the Tunnel-Client-Endpoint or Local-Loopback-Interface attribute, and if there is no corresponding L2TP tunnel profile configured on the MX Series router, then the L2TP tunnel fails to initiate because the router does not have a proper tunnel source address. In this case, the router can use the locally configured loopback address as the source address to successfully establish the L2TP tunnel.

RELATED DOCUMENTATION

| [RADIUS Attributes for L2TP](#) | 73

Example: Configuring RADIUS Authentication for L2TP

IN THIS SECTION

- [\[xref target has no title\]](#)
- [\[xref target has no title\]](#)
- [Configuration](#) | 78

Configuration

CLI Quick Configuration

The following example shows how to configure RADIUS authentication for L2TP:

[\[edit access\]](#)

```

profile example_bldg {
  client client_1 {
    chap-secret "$ABC123";
    ppp {
      interface-id west;
    }
    group-profile example_users;
  }
  client client_2 {
    chap-secret "$ABC123";
    group-profile example_users;
  }
  authentication-order radius;
}
radius-server {
  192.168.65.213 {
    port 1812;
    accounting-port 1813;
    secret "$ABC123"; # SECRET-DATA
  }
  192.168.65.223 {
    port 1812;
    accounting-port 1813;
    secret "$ABC123"; # SECRET-DATA
  }
}
radius-disconnect-port 2500;
radius-disconnect {
  192.168.65.152 secret "$ABC123";
  # SECRET-DATA
  192.168.64.153 secret "$ABC123";
  # SECRET-DATA
  192.168.64.157 secret "$ABC123";
  # SECRET-DATA
  192.168.64.173 secret "$ABC123";
  # SECRET-DATA
}

```

RELATED DOCUMENTATION

| [Configuring RADIUS Authentication for L2TP](#) | 71

Configuring the RADIUS Disconnect Server for L2TP

To configure the RADIUS disconnect server to listen for disconnect requests from an administrator and process them, include the following statements at the **[edit access]** hierarchy level:

```
[edit access]
radius-disconnect-port port-number;
radius-disconnect {
  client-address {
    secret password;
  }
}
```

port-number is the server port to which the RADIUS client sends disconnect requests. The L2TP network server, which accepts these disconnect requests, is the server. You can specify a port number on which to contact the RADIUS disconnect server. Most RADIUS servers use port number 1700.

NOTE: The Junos OS accepts only disconnect requests from the client address configured at the **[edit access radius-disconnect client-address]** hierarchy level.

client-address is the host sending disconnect requests to the RADIUS server. The client address is a valid IP address configured on one of the router or switch interfaces.

password authenticates the RADIUS client. Passwords can contain spaces. The secret used by the local router must match that used by the server.

For information about how to configure RADIUS authentication for L2TP, see [“Configuring RADIUS Authentication for L2TP” on page 71](#).

The following example shows the statements to be included at the **[edit access]** hierarchy level to configure the RADIUS disconnect server:

```
[edit access]
radius-disconnect-port 1700;
radius-disconnect {
  192.168.64.153 secret "$9$rtkl87ws4ZDkgokPT3tpEcyIWL7-VY4a";
  # SECRET-DATA
  192.168.64.162 secret "$9$rtkl87ws4ZDkgokPT3tpEcyIWL7-VY4a";
  # SECRET-DATA
}
```

RELATED DOCUMENTATION

| [Configuring RADIUS Authentication for L2TP](#) | 71

Configuring RADIUS Authentication for an L2TP Client and Profile

On an M10i or M7i router, L2TP supports RADIUS authentication and accounting for users with one set of RADIUS servers under the **[edit access]** hierarchy. You can also configure RADIUS authentication for each tunnel client or user profile.

To configure the RADIUS authentication for L2TP tunnel clients on an M10i or M7i router, include the **ppp-profile** statement with the **l2tp** attributes for tunnel clients:

```
[edit access profile profile-name client client-name l2tp]
ppp-profile profile-name;
```

ppp-profile *profile-name* specifies the profile used to validate PPP session requests through L2TP tunnels. Clients of the referenced profile must have only PPP attributes. The referenced group profile must be defined.

To configure the RADIUS authentication for a profile, include following statements at the **[edit access profile *profile-name*]** hierarchy level:

```
[edit access profile profile-name]
radius-server server-address {
  accounting-port port-number;
  port port-number;
  retry attempts;
  routing-instance routing-instance-name;
  secret password;
  source-address source-address;
  timeout seconds;
}
```

When a PPP user initiates a session and RADIUS authentication is configured for the user profile on the tunnel group, the following priority sequence is used to determine which RADIUS server is used for authentication and accounting:

- If the **ppp-profile** statement is configured under the tunnel client (LAC), the RADIUS servers configured under the specified **ppp-profile** are used.
- If RADIUS servers are configured under the user profile for the tunnel group, those servers will be used.

- If no RADIUS server is configured for the tunnel client (LAC) or user profile, then the RADIUS servers configured at the **[edit access]** hierarchy level are used.

RELATED DOCUMENTATION

[Example: Configuring RADIUS Authentication for an L2TP Profile](#) | 82

Example: Configuring RADIUS Authentication for an L2TP Profile

Configuration

The following example shows statements to be included at the **[edit access]** hierarchy level to configure RADIUS authentication for an L2TP profile:

CLI Quick Configuration

```
[edit access]
profile t {
  client LAC_A {
    l2tp {
      ppp-profile u;
    }
  }
}
profile u {
  client client_1 {
    ppp {
    }
  }
  5.5.5.5 {
    port 3333;
    secret $ABC123;
    source-address 1.1.1.1;
    retry 3;
    timeout 3;
  }
  6.6.6.6 secret $ABC123;
  7.7.7.7 secret $ABC123;
}
```

RELATED DOCUMENTATION

[Configuring RADIUS Authentication for an L2TP Client and Profile](#) | 81

Example: Configuring RADIUS-Based Subscriber Authentication and Accounting

Configuration

This example shows a RADIUS-based authentication and accounting configuration.

CLI Quick Configuration

```
[edit access]
radius-server {
  192.168.1.250 {
    port 1812;
    accounting-port 1813;
    accounting-retry 6;
    accounting-timeout 20;
    retry 3;
    secret $ABC123$ABC123;
    source-address 192.168.1.100;
    timeout 45;
  }
  192.168.1.251 {
    port 1812;
    accounting-port 1813;
    accounting-retry 6;
    accounting-timeout 20;
    retry 3;
    secret $ABC123;
    source-address 192.168.1.100;
    timeout 30;
  }
  2001:DB8:0f101::2{
    port 1812;
    accounting-port 1813;
    accounting-retry 6;
    accounting-timeout 20;
    retry 4;
    secret $ABC123$ABC123$ABC123-;
```

```

        source-address 2001:DB8:0f101::1;
        timeout 20;
    }
}
profile isp-bos-metro-fiber-basic {
    authentication-order radius;
    accounting {
        order radius;
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        immediate-update;
        statistics time;
        update-interval 12;
        wait-for-acct-on-ack;
        send-acct-status-on-config-change;
    }
    radius {
        authentication-server 192.168.1.251 192.168.1.252;
        accounting-server 192.168.1.250 192.168.1.251;
        options {
            accounting-session-id-format decimal;
            client-accounting-algorithm round-robin;
            client-authentication-algorithm round-robin;
            nas-identifier 56;
            nas-port-id-delimiter %;
            nas-port-id-format {
                nas-identifier;
                interface-description;
            }
            nas-port-type {
                ethernet {
                    wireless-80211;
                }
            }
        }
    }
    attributes {
        ignore {
            framed-ip-netmask;
        }
        exclude {
            accounting-delay-time [accounting-start accounting-stop];
            accounting-session-id [access-request accounting-on accounting-off
            accounting-start accounting-stop];
            dhcp-gi-address [access-request accounting-start accounting-stop];

```

```

        dhcp-mac-address [access-request accounting-start accounting-stop];
        nas-identifier [access-request accounting-start accounting-stop];
        nas-port [accounting-start accounting-stop];
        nas-port-id [accounting-start accounting-stop];
        nas-port-type [access-request accounting-start accounting-stop];
    }
}
}
[edit logical-systems isp-bos-metro-12 routing-instances isp-cmbrg-12-32]
interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.100/24;
            }
        }
    }
    ge-0/0/0 {
        vlan-tagging;
        unit 0 {
            vlan-id 200;
            family inet {
                unnumbered-address lo0.0;
            }
        }
    }
}

```

Configuring Subscriber Session Timeout Options

Subscriber session timeout options enable you to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both. The subscriber session options apply to both L2TP-tunneled and PPP-terminated subscriber sessions. For DHCP subscribers, the session timeout limits the DHCP lease time.

NOTE: To configure the timeout attributes in RADIUS, refer to the documentation for your RADIUS server.

To configure limitations on subscriber sessions, configure the session options in the client profile that applies to the subscriber:

- Terminate the subscriber when the configured session timeout expires, regardless of activity.

```
[edit access profile profile-name session-options]
user@host# set client-session-timeout minutes
```

- Terminate the subscriber when there is no ingress or egress data traffic for the duration of the configured idle timeout.

```
[edit access profile profile-name session-options]
user@host# set client-idle-timeout minutes
```

- Terminate the subscriber when there is no ingress data traffic for the duration of the configured idle timeout; ignore egress traffic.

```
[edit access profile profile-name session-options]
user@host# set client-idle-timeout minutes
user@host# set client-idle-timeout-ingress-only
```

For example, to configure session timeout options in the **acc-prof** client profile, specifying an idle timeout of 15 minutes, that only ingress traffic is monitored, and that the session times out after 120 minutes:

```
[edit]
access {
  profile {
    acc-prof {
      session-options {
        client-idle-timeout 15;
```

```
        client-idle-timeout-ingress-only;  
        client-session-timeout 120;  
    }  
}  
}
```

Configuring MLPPP

IN THIS CHAPTER

- [Understanding MLPPP Bundles on ACX Series Routers | 88](#)
- [Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers | 90](#)
- [Configuring Encapsulation for Multilink and Link Services Logical Interfaces | 96](#)
- [Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces | 97](#)
- [Configuring MRRU on Multilink and Link Services Logical Interfaces | 98](#)
- [Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces | 100](#)
- [Configuring Multiclass MLPPP on LSQ Interfaces | 100](#)
- [Configuring LSQ Interfaces as NxT1 or Nx E1 Bundles Using MLPPP on ACX Series | 102](#)
- [Example: Configuring an MLPPP Bundle on ACX Series | 109](#)

Understanding MLPPP Bundles on ACX Series Routers

ACX Series routers support MLPPP encapsulations. MLPPP enables you to bundle multiple PPP links into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1 and E1 links.

You configure multilink bundles as logical units or channels on the link services interface. With MLPPP, multilink bundles are configured as logical units on the link service interface—for example, **lsq-0/0/0.0,lsq-0/0/0.1**. MLPPP is supported on ACX1000, ACX2000, ACX2100 routers, and with Channelized OC3/STM1 (Multi-Rate) MICs with SFP and 16-port Channelized E1/T1 Circuit Emulation MIC on ACX4000 routers.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. The following table shows the maximum number of multilink bundles you can create on ACX Series routers:

Table 7: Multilink Bundles Supported by ACX Series Routers

ACX Platform	Maximum Bundles	Maximum Links	Maximum Links Per Bundle
ACX2000 ACX2100	16	16	16
ACX4000 ACX-MIC-16CHE1-T1-CE	16	16	16
ACX4000 ACX-MIC-4COC3-1COC12CE	50	336	16
ACX1000	8	8	8

The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.
- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

With multilink PPP bundles, you can use PPP Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP) for secure transmission over the PPP interfaces. For link services IQ (Isq) interfaces only, the maximum number of multilink classes to be negotiated when a link joins the bundle that you can specify by using the `multilink-max-classes` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level is limited to 4. Fragmentation size is not specified under fragmentation map; instead, fragmentation size configured on the bundle is used. Compressed Real-Time Transport Protocol (RTP) is not supported. HDLC address and control field compression (ACFC) and PPP protocol field compression (PFC) are not supported.

RELATED DOCUMENTATION

Link and Multilink Services Overview

Multilink Interfaces on Channelized MICs Overview

Guidelines for Configuring MLPPP With LSQ Interfaces on ACX Series Routers

You can configure MLPPP bundle interfaces with T1/E1 member links. The traffic that is transmitted over the MLPPP bundle interface is spread over the member links in a round-robin manner. If the packet size is higher than the fragmentation size configured on the MLPPP interface, the packet are fragmented. The fragments are also sent over member links in a round-robin pattern. The PPP control packets received on the interface are terminated on the router. The fragmentation size is configured at the MLPPP bundle-level. This fragmentation size is applied to all the packets on the bundle, regardless of the multilink class.

Multiclass MLPPP segregates the multilink protocol packets in to multiple classes. ACX routers support up to a maximum of four classes. One queue is associated with each of the four classes of multiclass MLPPP (MCML). The packets can be classified to be part of one of the classes. These packets take the queue associated with the class. The packets inside a queue are served in first-in first-out (FIFO) sequence.

Multiclass MLPPP is required to provide preferential treatment to high-priority, delay-sensitive traffic. The delay-sensitive smaller real-time frames are classified such that they end up in higher priority queue. While a lower priority packet is being fragmented, if a higher priority packet is enqueued, the lower priority fragmentation is suspended, the higher priority packet is fragmented and enqueued for transmission, and then the lower priority packet fragmentation is resumed.

Traditional LSQ interfaces (anchored on PICs) are supported to combine T1/E1 interfaces in an MLPPP bundle interface. Inline services (si-) interfaces and inline LSQ interfaces are not supported in MLPPP bundles. On ACX routers, MLPPP bundling is performed on the TDM MICs and traditional LSQ model is most effective mechanism. You can configure channelized OC interfaces (**t1-x/y/z:n:m**, **e1-x/y/z:n**) as members of an MLPPP bundle interface. A maximum of 16 member links per bundle is supported. The MPLS, ISO, and inet address families are supported. The ISO address family is supported only for IS-IS. You can configure MLPPP bundles on network-to-network interface (NNI) direction of an Ethernet pseudowire. Interleaving using multiclass MLPPP is supported.

Keep the following points in mind when you configure MLPPP bundles on ACX routers:

- The physical links must be of the same type and bandwidth.
- Round-robin packet distribution is performed over the member links.
- To add a T1 or E1 member link to the MLPPP bundle as link services LSQ interfaces, include the **bundle** statement at the **[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]** hierarchy level:

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```

- To configure the link services LSQ interface properties, include the following statements at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
encapsulation multilink-ppp;
fragment-threshold bytes;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

You can configure the address family as MPLS for the LSQ interfaces in an MLPPP bundle.

- PPP control protocol support depends on the processing of the PPP application for MLPPP bundle interfaces IPv4, Internet Protocol Control Protocol (IPCP), PPP Challenge Handshake Authentication Protocol (CHAP), and Password Authentication Protocol (PAP) applications are supported for PPP.
- Drop timeout configuration is not applicable to ACX routers
- The member links across MICs cannot be bundled. Only physical interfaces on the same MIC can be bundled.
- Fractional T1 and E1 interfaces are not supported. CoS is supported only for full T1 and E1 interfaces. Selective time slots of T1/E1 cannot be used and full T1/E1 interfaces must be used.
- Detailed statistics displayed depend on the parameters supported by the hardware. The counters that are supported by the hardware are displayed with appropriate values in the output of the **show interfaces lsq-fpc/pic/port detail** command. In the following sample output, the fields that are displayed with a value of 0 denote the fields that are not supported for computation by ACX routers. In the lsq- interface statistics, non- fragment statistics of the bundle are not accounted. Non-fragments are typically treated as single-fragment frames and counted in the fragment statistics.

```
user@host# show interfaces lsq-1/1/0 detail
Physical interface: lsq-1/1/0, Enabled, Physical link is Up
Interface index: 162, SNMP ifIndex: 550, Generation: 165
Description: LSQ-interface
Link-level type: LinkService, MTU: 1504
Device flags    : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x0
Last flapped   : 2015-06-22 19:01:47 PDT (2d 04:56 ago)
Statistics last cleared: 2015-06-23 05:01:49 PDT (1d 18:56 ago)
Traffic statistics:
Input  bytes :           108824           208896 bps
Output bytes :           90185           174080 bps
Input  packets:           1075             256 pps
Output packets:          1061             256 pps
```

```

IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Frame exceptions:
  Oversized frames 0
  Errored input frames 0
  Input on disabled link/bundle 0
  Output for disabled link/bundle 0
  Queuing drops 0
Buffering exceptions:
  Packet data buffer overflow 0
  Fragment data buffer overflow 0
Assembly exceptions:
  Fragment timeout 0
  Missing sequence number 0
  Out-of-order sequence number 0
  Out-of-range sequence number 0
Hardware errors (sticky):
  Data memory error 0
  Control memory error 0

Logical interface lsq-1/1/0.0 (Index 326) (SNMP ifIndex 599) (Generation 177)
  Flags: Up Point-To-Point SNMP-Traps 0x0 Encapsulation: Multilink-PPP
  Last flapped: 2015-06-24 23:57:34 PDT (00:00:51 ago)
  Bandwidth: 6144kbps
  Bundle links information:
    Active bundle links 4
    Removed bundle links 0
    Disabled bundle links 0
  Bundle options:
    MRRU 2000
    Remote MRRU 2000
    Drop timer period 0
    Inner PPP Protocol field compression enabled
    Sequence number format short (12 bits)
    Fragmentation threshold 450
    Links needed to sustain bundle 3
    Multilink classes 4
    Link layer overhead 4.0 %
  Bundle status:
    Received sequence number 0x0
    Transmit sequence number 0x0

```

```

    Packet drops                0 (0 bytes)
    Fragment drops              0 (0 bytes)
    MRRU exceeded               0
    Fragment timeout            0
    Missing sequence number     0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
    Packet data buffer overflow  0
    Fragment data buffer overflow 0
Statistics      Frames      fps      Bytes      bps
Bundle:
  Multilink:
    Input :      1076      256      484200      921600
    Output:      1061      256      477450      921600
  Network:
    Input :      2182      256      201812      208896
    Output:      2168      256      192029      174080
IPv6 Transit Statistics      Packets      Bytes
  Network:
    Input :              0              0
    Output:              0              0
Multilink class 0:
  Multilink:
    Input :      1075      256      483750      921600
    Output:      1061      256      477450      921600
  Network:
    Input :      1061      256      477450      921600
    Output:      1075      256      483750      921600
Multilink class 1:
  Multilink:
    Input :          0          0          0          0
    Output:          0          0          0          0
  Network:
    Input :          0          0          0          0
    Output:          0          0          0          0
Multilink class 2:
  Multilink:
    Input :          0          0          0          0
    Output:          0          0          0          0
  Network:
    Input :          0          0          0          0
    Output:          0          0          0          0
Multilink class 3:
  Multilink:

```

```

      Input :           0           0           0           0
      Output:           0           0           0           0
Network:
      Input :           0           0           0           0
      Output:           0           0           0           0
Link:
  t1-1/1/1.0
    Up time: 00:00:51
      Input :          280          64         126000         230400
      Output:          266          64         119700         230400
  t1-1/1/2.0
    Up time: 00:00:51
      Input :          266          64         119700         230400
      Output:          265          64         119250         230400
  t1-1/1/3.0
    Up time: 00:00:51
      Input :          265          64         119250         230400
      Output:          265          64         119250         230400
  t1-1/1/4.0
    Up time: 00:00:51
      Input :          265          64         119250         230400
      Output:          265          64         119250         230400
Multilink detail statistics:
Bundle:
  Fragments:
    Input :          1076          256         484200         921600
    Output:          1061          256         477450         921600
  Non-fragments:
    Input :           0           0           0           0
    Output:           0           0           0           0
  LFI:
    Input :           0           0           0           0
    Output:           0           0           0           0
Multilink class 0:
  Fragments:
    Input :          1076          256         484200         921600
    Output:          1061          256         477450         921600
  Non-fragments:
    Input :           0           0           0           0
    Output:           0           0           0           0
Multilink class 1:
  Fragments:
    Input :           0           0           0           0
    Output:           0           0           0           0

```

```

    Non-fragments:
      Input :          0          0          0          0
      Output:          0          0          0          0
Multilink class 2:
  Fragments:
    Input :          0          0          0          0
    Output:          0          0          0          0
  Non-fragments:
    Input :          0          0          0          0
    Output:          0          0          0          0
Multilink class 3:
  Fragments:
    Input :          0          0          0          0
    Output:          0          0          0          0
  Non-fragments:
    Input :          0          0          0          0
    Output:          0          0          0          0
NCP state: inet: Opened, inet6: Not-configured, iso: Opened, mpls: Opened
Protocol inet, MTU: 1500, Generation: 232, Route table: 0
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 9.1.9/24, Local: 9.1.9.18, Broadcast: Unspecified, Generation:
212
Protocol iso, MTU: 1500, Generation: 233, Route table: 0
  Flags: Is-Primary
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 234, Route table:
0
  Flags: Is-Primary

```

- For modifying the frame checksum (FCS) in the set of T1 options or E1 options on a MLPPP bundle member link, you must remove the member link out of the bundle by deactivating the link or unconfiguring it as a bundle member, and add the link back to the bundle after FCS modification. You must first remove the link from the bundle and modify FCS. If you are configuring FCS for the first time on the member link, specify the value before it is added to the bundle.

The following MLPPP functionalities are not supported:

- Member links across MICs.
- Fragmentation per class (only configurable at bundle level)
- IPv6 address family header compression (no address and control field compression [ACFC] or protocol field compression [PFC])
- Prefix elision as defined in *RFC 2686, The Multi-Class Extension to Multi-Link PPP*

- A functionality that resembles link fragmentation and interleaving (LFI) can be achieved using multiclass MLPPP (RFC 2686), which interleaves the high priority packets between lower priority packets. This methodology ensures that the delay desitive packets are sent as soon as they arrive. While LFI-classified packets are sent to a specific member link as PPP packets, the ACX implementation of interleaving contains multilink PPP (also referred to as PPP Multilink, MLP, and MP) headers and fragments that are sent on all member links in a round-robin manner.
- PPP over MLPPP bundle interfaces

RELATED DOCUMENTATION

Configuring Encapsulation for Multilink and Link Services Logical Interfaces

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

Multilink and link services interfaces support the following logical interface encapsulation types:

- MLPPP
- MLFR end-to-end

By default, the logical interface encapsulation type on multilink interfaces is MLPPP. The default logical interface encapsulation type on link services interfaces is MLFR end-to-end. For general information on encapsulation, see the *Junos OS Network Interfaces Library for Routing Devices*.

You can also configure physical interface encapsulation on link services interfaces. For more information, see *Configuring Encapsulation for Link Services Physical Interfaces*.

To configure multilink or link services encapsulation, include the **encapsulation** statement:

```
encapsulation type;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You must also configure the T1, E1, or DS0 physical interface with the same encapsulation type.

NOTE: ACX Series routers do not support DSO physical interface as member links.



CAUTION: When you configure the first MLFR encapsulated unit or delete the last MLFR encapsulated unit on a port, it triggers an interface encapsulation change on the port, which causes an interface flap on the other units within the port that are configured with generic Frame Relay.

RELATED DOCUMENTATION

Link and Multilink Services Overview

*Example: Configuring a Link Services Interface with MLPPP
encapsulation (Logical Interface)*

Configuring the Minimum Number of Active Links on Multilink and Link Services Logical Interfaces

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

You can set the minimum number of links that must be up for the multilink bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up. A member link is considered up when the PPP Link Control Protocol (LCP) phase transitions to open state.

The **minimum-links** value should be identical on both ends of the bundle.

To set the minimum number, include the **minimum-links** statement:

```
minimum-links number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For link services interfaces, you also can configure the minimum number of links at the physical interface level by including the **minimum-links** statement at the [edit interfaces *ls-fpc/pic/port:channel* **mlfr-uni-nni-bundle-options**] hierarchy level:

```
minimum-links number;
```

The number can be from 1 through 8. The maximum number of links supported in a bundle is 8. When 8 is specified, all configured links of a bundle must be up.

RELATED DOCUMENTATION

| [Link and Multilink Services Overview](#)

Configuring MRRU on Multilink and Link Services Logical Interfaces

The *maximum received reconstructed unit (MRRU)* is similar to a maximum transmission unit (MTU), but applies only to multilink bundles; it is the maximum packet size that the multilink interface can process. By default, the MRRU is set to 1500 bytes; you can configure a different MRRU value if the peer equipment allows this. The MRRU accounts for the original payload, for example the Layer 3 protocol payload, but does not include the 2-byte PPP header or the additional MLPPP or MLFR header applied while the individual multilink packets are traversing separate links in the bundle.

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

To configure a different MRRU value, include the **mrru** statement:

```
mrru bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

NOTE: ACX Series routers do not support logical systems.

For link services interfaces, you also can configure a different MRRU at the physical interface level by including the **mrru** statement at the **[edit interfaces ls-fpc/pic/port:channel mlfr-uni-nni-bundle-options]** hierarchy level:

```
mrru bytes;
```

The MRRU size can range from 1500 through 4500 bytes.

NOTE: If you set the MRRU on a bundle to a value larger than the MTU of the individual links within it, you must enable a fragmentation threshold for that bundle. Set the threshold to a value no larger than the smallest MTU of any link included in the bundle.

Determine the appropriate MTU size for the bundle by ensuring that the MTU size does not exceed the sum of the encapsulation overhead and the MTU sizes for the links in the bundle.

You can configure separate **family mtu** values on the following protocol families under bundle interfaces: **inet**, **inet6**, **iso**, and **mpls**. If not configured, the default value of 1500 is used on all except for **mpls** configurations, in which the value 1488 is used.

NOTE: ACX Series routers do not support **family inet6** on MLPPP interfaces.

NOTE: The effective family MTU might be different from the MTU value specified for MLPPP configurations, because it is adjusted downward by the remote MRRU's constraints. The remote MRRU configuration is not supported on M120 routers.

RELATED DOCUMENTATION

| [Link and Multilink Services Overview](#)

Configuring the Sequence Header Format on Multilink and Link Services Logical Interfaces

NOTE: Only MLPPP is supported on ACX Series routers. MLFR is not supported on ACX Series routers.

For MLPPP, the sequence header format is set to 24 bits by default. You can configure an alternative value of 12 bits, but 24 bits is considered the more robust value for most networks.

To configure a different sequence header value, include the **short-sequence** statement:

```
short-sequence;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For MLFR FRF.15, the sequence header format is set to 24 bits by default. This is the only valid option.

RELATED DOCUMENTATION

[Link and Multilink Services Overview](#)

[Configuring DLCIs on Link Services Logical Interfaces](#)

Configuring Multiclass MLPPP on LSQ Interfaces

For link services LSQ (**lsq-**) interfaces with MLPPP encapsulation, you can configure multiclass MLPPP (MCML). If you do not configure MCML, fragments from different classes cannot be interleaved. All fragments for a single packet must be sent before the fragments from another packet are sent. Nonfragmented packets can be interleaved between fragments of another packet to reduce latency seen by nonfragmented packets. In effect, latency-sensitive traffic is encapsulated as regular PPP traffic, and bulk traffic is encapsulated as multilink traffic. This model works as long as there is a single class of latency-sensitive traffic, and there is no high-priority traffic that takes precedence over latency-sensitive traffic. This approach to LFI, used on the Link Services PIC, supports only two levels of traffic priority, which is not sufficient to carry the four-to-eight forwarding classes that are supported by M Series and T

Series routers. For more information about the Link Services PIC support of LFI, see *Configuring Delay-Sensitive Packet Interleaving on Link Services Logical Interfaces*.

NOTE: ACX Series routers do not support link fragmentation interleaving (LFI).

For link services LSQ interfaces only, you can configure MCML, as defined in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. MCML makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, MCML allows different classes of traffic to have different latency guarantees. With MCML, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees.

NOTE: Configuring both LFI and MCML on the same bundle is not necessary, nor is it supported, because multiclass MLPPP represents a superset of functionality. When you configure multiclass MLPPP, LFI is automatically enabled.

The Junos OS implementation of MCML does not support compression of common header bytes, which is referred to in RFC 2686 as “prefix elision.”

MCML greatly simplifies packet ordering issues that occur when multiple links are used. Without MCML, all voice traffic belonging to a single flow is hashed to a single link to avoid packet ordering issues. With MCML, you can assign voice traffic to a high-priority class, and you can use multiple links. For more information about voice services support on link services IQ interfaces (**lsq**), see *Configuring Services Interfaces for Voice Services*.

To configure MCML on a link services IQ interface, you must specify how many multilink classes should be negotiated when a link joins the bundle, and you must specify the mapping of a forwarding class into an MCML class.

To specify how many multilink classes should be negotiated when a link joins the bundle, include the **multilink-max-classes** statement:

```
multilink-max-classes number;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The number of multilink classes can be 1 through 8. The number of multilink classes for each forwarding class must not exceed the number of multilink classes to be negotiated.

NOTE: In ACX Series routers, the multilink classes can be 1 through 4.

To specify the mapping of a forwarding class into a MCML class, include the **multilink-class** statement at the **[edit class-of-service fragmentation-maps map-name forwarding-class class-name]** hierarchy level:

```
[edit class-of-service fragmentation-maps map-name forwarding-class class-name]
multilink-class number;
```

The multilink class index number can be 0 through 7. The **multilink-class** statement and **no-fragmentation** statements are mutually exclusive.

NOTE: In ACX Series routers, the multilink class index number can be 0 through 3. ACX Series routers do not support the **no-fragmentation** statement for fragmentation map.

To view the number of multilink classes negotiated, issue the **show interfaces lsq-fpc/pic/port.logical-unit-number detail** command.

RELATED DOCUMENTATION

Layer 2 Service Package Capabilities and Interfaces

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP

Configuring LSQ Interfaces for ATM2 IQ Interfaces Using MLPPP

Configuring LSQ Interfaces for T3 Links Configured for Compressed RTP over MLPPP

Link Services Configuration for Junos Interfaces

Configuring LSQ Interfaces as NxT1 or NxE1 Bundles Using MLPPP on ACX Series

To configure an NxT1 bundle using MLPPP, you aggregate *N* different T1 links into a bundle. The NxT1 bundle is called a logical interface, because it can represent, for example, a routing adjacency. To aggregate T1 links into a an MLPPP bundle, include the **bundle** statement at the **[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]** hierarchy level:

```
[edit interfaces t1-fpc/pic/port unit logical-unit-number family mlppp]
bundle lsq-fpc/pic/port.logical-unit-number;
```

NOTE: LSQ interfaces support both T1 and E1 physical interfaces. These instructions apply to T1 interfaces, but the configuration for E1 interfaces is similar.

To configure the LSQ interface properties, include the following statements at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port unit logical-unit-number]
drop-timeout milliseconds;
encapsulation multilink-ppp;
fragment-threshold bytes;
link-layer-overhead percent;
minimum-links number;
mrru bytes;
short-sequence;
family inet {
    address address;
}
```

NOTE: ACX Series routers do not support drop-timeout and link-layer-overhead properties.

The logical link services IQ interface represents the MLPPP bundle. For the MLPPP bundle, there are four associated queues on M Series routers and eight associated queues on M320 and T Series routers. A scheduler removes packets from the queues according to a scheduling policy. Typically, you designate one queue to have strict priority, and the remaining queues are serviced in proportion to weights you configure.

For MLPPP, assign a single scheduler map to the link services IQ interface (**lsq**) and to each constituent link. The default schedulers for M Series and T Series routers, which assign 95, 0, 0, and 5 percent bandwidth for the transmission rate and buffer size of queues 0, 1, 2, and 3, are not adequate when you configure LFI or multiclass traffic. Therefore, for MLPPP, you should configure a single scheduler with nonzero percent transmission rates and buffer sizes for queues 0 through 3, and assign this scheduler to the link services IQ interface (**lsq**) and to each constituent link, as shown in [“Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP” on page 106](#).

NOTE: For M320 and T Series routers, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent.

If the bundle has more than one link, you must include the **per-unit-scheduler** statement at the **[edit interfaces lsq-fpc/pic/port]** hierarchy level:

```
[edit interfaces lsq-fpc/pic/port]
per-unit-scheduler;
```

To configure and apply the scheduling policy, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  t1-fpc/pic/port unit logical-unit-number {
    scheduler-map map-name;
  }
}
forwarding-classes {
  queue queue-number class-name;
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    priority priority-level;
    transmit-rate (rate | remainder) <exact>;
  }
}
```

For link services IQ interfaces, a strict-high-priority queue might starve the other three queues because traffic in a strict-high priority queue is transmitted before any other queue is serviced. This implementation is unlike the standard Junos CoS implementation in which a strict-high-priority queue does round-robin with high-priority queues, as described in the *Class of Service User Guide (Routers and EX9200 Switches)*.

After the scheduler removes a packet from a queue, a certain action is taken. The action depends on whether the packet came from a multilink encapsulated queue (fragmented and sequenced) or a nonencapsulated queue (hashed with no fragmentation). Each queue can be designated as either multilink

encapsulated or nonencapsulated, independently of the other. By default, traffic in all forwarding classes is multilink encapsulated. To configure packet fragmentation handling on a queue, include the **fragmentation-maps** statement at the **[edit class-of-service]** hierarchy level:

```
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      multilink-class number;
    }
  }
}
```

For NxT1 bundles using MLPPP, the byte-wise load balancing used in multilink-encapsulated queues is superior to the flow-wise load balancing used in nonencapsulated queues. All other considerations are equal. Therefore, we recommend that you configure all queues to be multilink encapsulated. You do this by including the **fragment-threshold** statement in the configuration. You use the **multilink-class** statement to map a forwarding class into a multiclass MLPPP (MCML). For more information about MCML, see [“Configuring Multiclass MLPPP on LSQ Interfaces” on page 100](#). For more information about fragmentation maps, see *Configuring CoS Fragmentation by Forwarding Class on LSQ Interfaces*.

When a packet is removed from a multilink-encapsulated queue, the software gives the packet an MLPPP header. The MLPPP header contains a sequence number field, which is filled with the next available sequence number from a counter. The software then places the packet on one of the *N* different T1 links. The link is chosen on a packet-by-packet basis to balance the load across the various T1 links.

If the packet exceeds the minimum link MTU, or if a queue has a fragment threshold configured at the **[edit class-of-service fragmentation-maps map-name forwarding-class class-name]** hierarchy level, the software splits the packet into two or more fragments, which are assigned consecutive multilink sequence numbers. The outgoing link for each fragment is selected independently of all other fragments.

If you do not include the **fragment-threshold** statement in the fragmentation map, the fragmentation threshold you set at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level is the default for all forwarding classes. If you do not set a maximum fragment size anywhere in the configuration, packets are fragmented if they exceed the smallest MTU of all the links in the bundle.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the maximum received reconstructed unit (MRRU) by including the **mrru** statement at the **[edit interfaces lsq-fpc/pic/port unit logical-unit-number]** hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1500 bytes, and you can configure it to be from 1500 through 4500 bytes. For more information, see *Configuring MRRU on Multilink and Link Services Logical Interfaces*.

When a packet is removed from a nonencapsulated queue, it is transmitted with a plain PPP header. Because there is no MLPPP header, there is no sequence number information. Therefore, the software must take special measures to avoid packet reordering. To avoid packet reordering, the software places

the packet on one of the N different T1 links. The link is determined by hashing the values in the header. For IP, the software computes the hash based on source address, destination address, and IP protocol. For MPLS, the software computes the hash based on up to five MPLS labels, or four MPLS labels and the IP header.

For UDP and TCP the software computes the hash based on the source and destination ports, as well as source and destination IP addresses. This guarantees that all packets belonging to the same TCP/UDP flow always pass through the same T1 link, and therefore cannot be reordered. However, it does not guarantee that the load on the various T1 links is balanced. If there are many flows, the load is usually balanced.

The N different T1 interfaces link to another router, which can be from Juniper Networks or another vendor. The router at the far end gathers packets from all the T1 links. If a packet has an MLPPP header, the sequence number field is used to put the packet back into sequence number order. If the packet has a plain PPP header, the software accepts the packet in the order in which it arrives and makes no attempt to reassemble or reorder the packet.

Example: Configuring an LSQ Interface as an NxT1 Bundle Using MLPPP

```
[edit interfaces]
lsq-1/1/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-ppp;
    mrru 2000;
    multilink-max-classes 4;
    family inet {
      address 20.1.1.1/24;
    }
    family mpls;
  }
}
ct1-1/1/4 {
  enable;
  no-partition interface-type t1;
}
t1-1/1/4 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/1/0.0;
    }
  }
}
```

```

ct1-1/1/5 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/5 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
}
}
class-of-service {
    classifiers {
        inet-precedence myIPv4 {
            forwarding-class best-effort {
                loss-priority low code-points 000;
            }
            forwarding-class expedited-forwarding {
                loss-priority low code-points 001;
            }
            forwarding-class assured-forwarding {
                loss-priority low code-points 011;
            }
            forwarding-class network-control {
                loss-priority low code-points 111;
            }
        }
    }
}
drop-profiles {
    dp1 {
        fill-level 50 drop-probability 0;
        fill-level 100 drop-probability 100;
    }
    dp2 {
        fill-level 50 drop-probability 0;
        fill-level 100 drop-probability 100;
    }
}
interfaces {
    lsq-1/1/0 {
        unit 0 {
            scheduler-map sm;

```

```

        fragmentation-map frag;
        rewrite-rules {
            inet-precedence myRRIPv4;
        }
    }
}

rewrite-rules {
    inet-precedence myRRIPv4 {
        forwarding-class best-effort {
            loss-priority low code-point 111;
        }
        forwarding-class expedited-forwarding {
            loss-priority low code-point 011;
        }
        forwarding-class assured-forwarding {
            loss-priority low code-point 001;
        }
        forwarding-class network-control {
            loss-priority low code-point 000;
        }
    }
}

scheduler-maps {
    sm {
        forwarding-class best-effort scheduler new;
        forwarding-class network-control scheduler new_nc;
        forwarding-class assured-forwarding scheduler new_af;
        forwarding-class expedited-forwarding scheduler new_ef;
    }
}

fragmentation-maps {
    frag {
        forwarding-class {
            best-effort {
                multilink-class 3;
            }
            network-control {
                multilink-class 0;
            }
            assured-forwarding {
                multilink-class 2;
            }
            expedited-forwarding {

```

```

        multilink-class 1;
    }
}
}
}
}
schedulers {
    new {
        transmit-rate 32k;
        shaping-rate 3m;
        priority low;
        drop-profile-map loss-priority low protocol any drop-profile dp1;
        drop-profile-map loss-priority high protocol any drop-profile dp2;
    }
    new_nc {
        transmit-rate 32k;
        shaping-rate 3m;
        priority strict-high;
    }
    new_af {
        transmit-rate 32k;
        shaping-rate 3m;
        priority medium-low;
    }
    new_ef {
        transmit-rate 32k;
        shaping-rate 3m;
        priority medium-high;
    }
}
}
}

```

Example: Configuring an MLPPP Bundle on ACX Series

IN THIS SECTION

- [\[xref target has no title\]](#)
- [\[xref target has no title\]](#)
- [Configuration | 110](#)

Configuration

The following is a sample for configuring an MLPPP bundle on ACX Series routers:

CLI Quick Configuration

```
[edit]
user@host# show interfaces
lsq-1/1/0 {
  description LSQ-interface;
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-ppp;
    mrru 2000;
    short-sequence;
    fragment-threshold 450;
    minimum-links 3;
    multilink-max-classes 4;
    family inet {
      address 9.1.9.18/24
    }
    family iso;
    family mpls;
  }
}
ct1-1/1/1 {
  enable;
  no-partition interface-type t1;
}
t1-1/1/1 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
      bundle lsq-1/1/0.0;
    }
  }
}
ct1-1/1/2 {
  enable;
  no-partition interface-type t1;
}
t1-1/1/2 {
  encapsulation ppp;
  unit 0 {
    family mlppp {
```

```
        bundle lsq-1/1/0.0;
    }
}
}
ct1-1/1/3 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/3 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
ct1-1/1/4 {
    enable;
    no-partition interface-type t1;
}
t1-1/1/4 {
    encapsulation ppp;
    unit 0 {
        family mlppp {
            bundle lsq-1/1/0.0;
        }
    }
}
```

Configuration Statements

IN THIS CHAPTER

- Access Configuration Statements | 114
- accounting (Access Profile) | 119
- accounting-order | 120
- accounting-stop-on-access-deny | 121
- accounting-stop-on-failure | 122
- address-assignment (Address-Assignment Pools) | 123
- address-pool | 126
- attributes (RADIUS Attributes) | 128
- authentication-order | 131
- cell-overhead | 133
- circuit-type (DHCP Local Server) | 134
- client | 136
- compression (PPP Properties) | 140
- dead-peer-detection | 141
- default-chap-secret | 143
- default-pap-password | 144
- dhcp-attributes (Address-Assignment Pools) | 145
- encapsulation-overhead | 152
- exclude (RADIUS Attributes) | 153
- framed-pool | 161
- group-profile (Group Profile) | 162
- host (Address-Assignment Pools) | 164
- idle-timeout (Access) | 165
- ike (Access Profile) | 167
- immediate-update | 169
- interface-description-format | 170
- interface-id | 172
- keepalive | 173

- [keepalive-retries | 174](#)
- [l2tp \(Group Profile\) | 175](#)
- [l2tp \(Profile\) | 176](#)
- [lcp-renegotiation | 181](#)
- [local-chap | 182](#)
- [maximum-sessions-per-tunnel | 183](#)
- [multilink | 184](#)
- [nas-port-extended-format | 186](#)
- [network | 188](#)
- [option-82 \(Address-Assignment Pools\) | 189](#)
- [option-match | 190](#)
- [options \(Access Profile\) | 191](#)
- [order | 201](#)
- [pool \(Address-Assignment Pools\) | 202](#)
- [ppp \(Group Profile\) | 204](#)
- [ppp \(Profile\) | 207](#)
- [primary-dns | 208](#)
- [primary-wins | 209](#)
- [profile \(Access\) | 210](#)
- [radius \(Access Profile\) | 217](#)
- [radius-disconnect | 221](#)
- [radius-disconnect-port | 222](#)
- [radius-server | 223](#)
- [range \(Address-Assignment Pools\) | 229](#)
- [revert-interval \(Access\) | 230](#)
- [secondary-dns | 231](#)
- [secondary-wins | 232](#)
- [secret \(RADIUS\) | 233](#)
- [session-options | 234](#)
- [statistics \(Access Profile\) | 237](#)
- [update-interval | 238](#)

Access Configuration Statements

To configure access, include the following statements at the **[edit access]** hierarchy level:

```
[edit access]
address-assignment {
  neighbor-discovery-router-advertisement;
  pool pool-name {
    family inet {
      dhcp-attributes {
        [protocol-specific-attributes];
      }
      host hostname {
        hardware-address mac-address;
        ip-address ip-address;
      }
      network address-or-prefix </subnet-mask>;
      range name {
        high upper-limit;
        low lower-limit;
        prefix-length prefix-length;
      }
    }
  }
}
address-pool pool-name {
  address address-or-prefix;
  address-range low <lower-limit> high <upper-limit>;
}
domain {
  delimiter;
  map;
  parse-direction;
};
group-profile profile-name {
  l2tp {
    interface-id interface-id;
    lcp-renegotiation;
    local-chap;
    maximum-sessions-per-tunnel number;
    multilink {
      drop-timeout milliseconds;
      fragment-threshold bytes;
    }
  }
}
```

```

}
ppp {
    cell-overhead;
    encapsulation-overhead bytes;
    framed-pool pool-id;
    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;
}
}
profile profile-name {
    accounting {
        accounting-stop-on-access-deny;
        accounting-stop-on-failure;
        coa-immediate-update;
        duplication;
        immediate-update;
        order [ accounting-method ];
        statistics (time | volume-time);
        update-interval minutes;
    }
    accounting-order radius;
    authentication-order [ authentication-methods ];
    authorization-order jsr;
    client client-name {
        chap-secret chap-secret;
        client-group [ group-names ];
        firewall-user {
            password password;
        }
        group-profile profile-name;
        ike {
            allowed-proxy-pair {
                local local-proxy-address remote remote-proxy-address;
            }
            ike-policy policy-name;
            initiate-dead-peer-detection;
            interface-id interface-id;
            ipsec-policy policy-name;
            pre-shared-key (ascii-text character-key-string | hexadecimal hexadecimal-digits-key-string);

```

```

}
l2tp {
    interface-id interface-identifier;
    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;
}
provisioning-order jsr;
user-group-profile profile-name;
}
radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
        accounting-authentic [ accounting-on | accounting-off ];
        accounting-delay-time [ accounting-on | accounting-off ];
        accounting-session-id [ access-request | accounting-on | accounting-off | accounting-stop ];
        accounting-terminate-cause [ accounting-off ];
        called-station-id [ access-request | accounting-start | accounting-stop ];
        calling-station-id [ access-request | accounting-start | accounting-stop ];
        class [ accounting-start | accounting-stop ];
        dhcp-options [ access-request | accounting-start | accounting-stop ];
        dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    }
}

```

```

    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    output-filter [ accounting-start | accounting-stop ];
    event-timestamp [ accounting-on | accounting-off | accounting-start | accounting-stop ];
    framed-ip-address [ accounting-start | accounting-stop ];
    framed-ip-netmask [ accounting-start | accounting-stop ];
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets [ accounting-stop ];
    input-gigawords [ accounting-stop ];
    interface-description [ access-request | accounting-start | accounting-stop ];
    nas-identifier [ access-request | accounting-on | accounting-off | accounting-start | accounting-stop ];
    nas-port [ access-request | accounting-start | accounting-stop ];
    nas-port-id [ access-request | accounting-start | accounting-stop ];
    nas-port-type [ access-request | accounting-start | accounting-stop ];
    output-gigapackets [ accounting-stop ];
    output-gigawords [ accounting-stop ];
}
ignore {
    framed-ip-netmask;
    input-filter;
    logical-system-routing-instance;
    output-filter;
}
}
authentication-server [ ip-address ];
options {
    accounting-session-id-format (decimal | description);
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm (direct | round-robin);
    ethernet-port-type-virtual;
    interface-description-format [sub-interface | adapter];
    nas-identifier identifier-value;
    nas-port-extended-format {
        adapter-width width;
        port-width width;
        slot-width width;
        stacked-vlan-width width;
        vlan-width width;
    }
    revert-interval interval;
    vlan-nas-port-stacked-format;
}
}
radius-options {
    revert-interval interval;

```

```

}
radius-disconnect {
    client-address {
        secret password;
    }
}
radius-disconnect-port port-number;
radius-options {
    revert-interval interval;
}
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address address;
    timeout seconds;
}
session-options {
    client-group [ group-names ];
    client-idle-timeout minutes;
    client-session-timeout minutes;
}

```

RELATED DOCUMENTATION

[Configuring the PPP Authentication Protocol | 5](#)

[Example: Configuring PPP CHAP | 13](#)

[Configuring the PPP Authentication Protocol | 5](#)

[Example: Configuring CHAP Authentication with RADIUS | 14](#)

[Configuring L2TP for Enabling PPP Tunneling Within a Network | 18](#)

[Defining the Minimum L2TP Configuration | 20](#)

[Configuring the Address Pool for L2TP Network Server IP Address Allocation | 21](#)

[Configuring the Group Profile for Defining L2TP Attributes | 24](#)

accounting (Access Profile)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

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

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Understanding RADIUS Accounting Duplicate Reporting

accounting-order

Syntax

```
accounting-order (radius | [accounting-order-data-list]);
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.0.

Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description

Specify the order in which accounting methods are used.

Options

radius—Use the RADIUS accounting method.

[*accounting-order-data-list*]—Set of data listing the accounting order to be used, enclosed in brackets. This can be any combination of accounting methods, up to and including a list of the entire accounting order.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the Accounting Order | 31](#)

accounting-stop-on-access-deny

Syntax

```
accounting-stop-on-access-deny;
```

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.

Description

Configure RADIUS accounting to send an Acct-Stop message when the AAA server refuses a client request for access.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Authentication and Accounting Parameters for Subscriber Access*

accounting-stop-on-failure

Syntax

```
accounting-stop-on-failure;
```

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.

Description

Configure RADIUS accounting to send an Acct-Stop message when a subscriber session has been successfully authenticated and authorized, but then fails before an Acct-Start message is sent. By default, an Acct-Stop message is sent only if an Acct-Start message has been exchanged with the accounting server.

Consider a situation where RADIUS address pools are used to assign IP/IPv6 addresses. After a subscriber session is successfully authenticated, the RADIUS server authorizes the session by assigning an IP address from the RADIUS address pool and conveying that address in the Framed-IP-Address attribute. If a negotiation failure occurs at this point, the session is terminated before activating. The Acct-Start message is never sent because it is initiated by session activation. By default, an Acct-Stop message cannot be sent because the Acct-Start is never sent. However, if the **acct-stop-on-failure** statement is configured, the negotiation failure causes the Acct-Stop message to be sent, which explicitly notifies the RADIUS server that the session is disconnected and that it can free the allocated IP address back to the pool.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Authentication and Accounting Parameters for Subscriber Access*

address-assignment (Address-Assignment Pools)

Syntax

```

address-assignment {
  abated-utilization percentage;
  abated-utilization-v6 percentage;
  high-utilization percentage;
  high-utilization-v6 percentage;
  neighbor-discovery-router-advertisement ndra-pool-name;
  pool pool-name {
    active-drain;
    family family {
      dhcp-attributes {
        protocol-specific attributes;
      }
      excluded-address ip-address;
      excluded-range name low minimum-value high maximum-value;
      host hostname {
        hardware-address mac-address;
        ip-address ip-address;
      }
      network ip-prefix/<prefix-length>;
      prefix ipv6-prefix;
      range range-name {
        high upper-limit;
        low lower-limit;
        prefix-length prefix-length;
      }
    }
    hold-down;
    link pool-name;
  }
}

```

Hierarchy Level

[edit access]

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

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

NOTE: Support for subordinate statements is platform-specific. See individual statement topics for support information.

Options

abated-utilization—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate SNMP traps for DHCP address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

Values: *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

abated-utilization-v6—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate SNMP traps for DHCPv6 address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

Values: *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

high-utilization—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate an SNMP trap when the DHCP address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

Values: *percentage*—Percentage used to generate a trap. Range: 2 through 99.

high-utilization-v6—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate an SNMP trap when the DHCPv6 address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

Values: *percentage*—Percentage used to generate a trap. Range: 2 through 99.

neighbor-discovery-router-advertisement—(M Series, MX Series, SRX Series, T Series only) Configure the name of the address-assignment pool used to assign the router advertisement prefix.

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

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Address-Assignment Pools Overview](#)

[Address-Assignment Pool Configuration Overview](#)

[Configuring an Address-Assignment Pool for L2TP LNS with Inline Services](#)

[Configuring Address-Assignment Pool Usage Threshold Traps](#)

[Configuring an Address-Assignment Pool Used for Router Advertisements](#)

address-pool

Syntax

```
address-pool pool-name {
  address address-or-prefix;
  address-range <low lower-limit> <high upper-limit>;
}
```

Hierarchy Level

[edit access]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Allocate IP addresses for clients.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

pool-name—Name assigned to an address pool.

address—(EX Series, M Series, PTX Series, T Series only) Configure the IP address or prefix value for clients.

Values: *address-or-prefix*—An address or prefix value.

address-range—Configure the address range.

Values:

- high *upper-limit*—Upper limit of an address range.
- low *lower-limit*—Lower limit of an address range.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

attributes (RADIUS Attributes)

Syntax

```

attributes {
  exclude {
    attribute-name packet-type;
    standard-attribute number {
      packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
    vendor-id id-number {
      vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system-routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}

```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Specify how the router or switch processes RADIUS attributes.

Options

exclude—Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the **ignore** statement.

The options for this statement are explained separately. Click the linked statement for details.

ignore—Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. Standard attributes and VSAs received in RADIUS messages take precedence over internally provisioned attribute values. Ignoring the attributes enables your internally provisioned values to be used instead. Contrast this behavior with that provided by the **exclude** statement.

Starting in Junos OS Release 18.1R1, you can specify RADIUS standard attributes with the attribute number. You can specify vendor-specific attributes (VSAs) with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be ignored. The configuration has no effect if you can configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to ignore only a subset of all attributes that can be received in Access-Accept messages.

Values:

- **dynamic-iflset-name**—Ignore Juniper Networks VSA 26-130, Qos-Set-Name.
- **framed-ip-netmask**—Ignore RADIUS attribute 9, Framed-IP-Netmask.
- **idle-timeout**—Ignore RADIUS attribute 28, Idle-Timeout.
- **input-filter**—Ignore Juniper Networks VSA 26-10, Ingress-Policy-Name.
- **logical-system-routing-instance**—Ignore Juniper Networks VSA 26-1.
- **output-filter**—Ignore Juniper Networks VSA 26-11, Egress-Policy-Name.
- **session-timeout**—Ignore RADIUS attribute 27, Session-Timeout.
- **standard-attribute *number***—RADIUS standard attribute number supported by your platform. You can enclose multiple values in square brackets to specify a list of attributes. If you configure an unsupported attribute, that configuration has no effect. Range: 1 through 255.
- **vendor-attribute *vsa-number***—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. You can enclose multiple values in square brackets to specify a list of VSAs. If you configure an unsupported VSA, that configuration has no effect. Range: 1 through 255.
- **vendor-id *id-number***—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect. Range: 1 through 16777215.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[RADIUS Servers and Parameters for Subscriber Access](#)

[Standard and Vendor-Specific RADIUS Attributes](#)

[AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS](#)

[AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS](#)

authentication-order

Syntax

```
authentication-order [ authentication-methods ];
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

none option added in Junos OS Release 11.2.

nasreq option added in Junos OS Release 16.1.

s6a option added in Junos OS Release 19.3R1.

Description

Set the order in which AAA tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, AAA tries the authentication methods in order, from first to last.

A given subscriber does not undergo both authentication and authorization as separate steps. When both **authentication-order** and **authorization-order** are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Starting in Junos OS Release 18.2R1, the **password** option can also be used to specify that local authentication and local authorization is attempted for individual subscribers that are configured with the **subscriber** statement at the `[edit access profile profile-name]` hierarchy level.

Options

authentication-methods—Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination:

- **nasreq**—Verify subscribers using the Diameter-based Network Access Server Requirements (NASREQ) protocol.
- **none**—No authentication is performed. Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.

NOTE: Subscriber access management does not support the **none** option; authentication fails when this option is specified.

- **password**—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

Subscriber access management does not support the **password** option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, this option is used to enable local authentication and optionally local authorization for individual subscribers. Local authentication is typically used when you do not have external authentication and authorization servers. The password itself must be configured with the **subscriber** statement in the same access profile. Local authentication is performed when a subscriber logs in with a matching username; it succeeds if the subscribers login password matches the password in the profile.

If you have external authentication and authorization servers, you can use local authentication as a backup authentication method. In this case, configure **password** other than first in the list of methods.

- **radius**—Verify the client using RADIUS authentication services.
- **s6a**—Verify subscribers using the Diameter-based s6a protocol.

Default: password

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Example: Configuring CHAP Authentication with RADIUS | 14](#)

Specifying the Authentication and Accounting Methods for Subscriber Access

[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

Configuring Local Authentication and Authorization for Subscribers

Example: Configure S6a Application

cell-overhead

Syntax

```
cell-overhead;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],  
[edit access profile profile-name client client-name ppp]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

circuit-type (DHCP Local Server)

Syntax

```
circuit-type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  group group-name authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems logical-system-name system services dhcp-local-server dhcpv6 group group-name authentication
  username-include],
[edit logical-systems logical-system-name system services dhcp-local-server group group-name authentication
  username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  dhcpv6 authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  dhcpv6 group group-name authentication username-include],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
  group group-name authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 authentication
  username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server dhcpv6 group group-name
  authentication username-include],
[edit routing-instances routing-instance-name system services dhcp-local-server group group-name authentication
  username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],
[edit system services dhcp-local-server group group-name authentication username-include]
```

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Specifying Authentication Support*

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 number;
        maximum-sessions-per-tunnel number;
        multilink {
            drop-timeout milliseconds;
            fragment-threshold bytes;
        }
        override-result-code session-out-of-resource;
        ppp-authentication (chap | pap);
        ppp-profile profile-name;
        sessions-limit-group;
        service-profile profile-name(parameter)&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.

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**.

chap-secret—(M Series, MX Series, SRX Series, T Series only) For interfaces with PPP encapsulation on which the PPP Challenge Handshake Authentication Protocol (CHAP) is configured, configure the shared secret (the CHAP secret key associated with a peer), as defined in RFC 1994. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *chap-secret*—The secret key associated with a peer.

group-profile—Associate a group profile with a client. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *profile-name*—Name assigned to the group profile.

pap-password—(M Series, MX Series, T Series only) Configure the Password Authentication Protocol (PAP) password. This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *password*—PAP password.

user-group-profile—Apply a configured PPP group profile to PPP users. If **user-group-profile** is modified or deleted, the existing LNS subscribers, which were using this Layer 2 Tunneling Protocol client configuration, go down.

Values:

- *profile-name*—Name of a PPP group profile configured at the [edit access group-profile *profile-name*] hierarchy level.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the CHAP Secret for an L2TP Profile | 37](#)

[Configuring PPP CHAP Authentication | 55](#)

[Referencing the Group Profile from the L2TP Profile | 38](#)

[Configuring the L2TP Client | 35](#)

[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

[Configuring the PAP Password for an L2TP Profile | 43](#)

[Applying a Configured PPP Group Profile to a Tunnel | 46](#)

[*Configuring an L2TP Access Profile on the LNS*](#)

compression (PPP Properties)

Syntax

```
compression {  
    acfc;  
    pfc;  
}
```

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

For interfaces with PPP encapsulation, set Link Control Protocol (LCP) compression options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring PPP Address and Control Field Compression](#) | 52

[Configuring the PPP Protocol Field Compression](#) | 9

dead-peer-detection

Syntax

```
dead-peer-detection {  
    (always-send | optimized | probe-idle-tunnel);  
    interval seconds;  
    threshold number;  
}
```

Hierarchy Level

```
[edit security ike gateway gateway-name]
```

Release Information

Statement introduced in Junos OS Release 8.5. Support for the **optimized** and **probe-idle-tunnel** options added in Junos OS Release 12.1X46-D10.

Description

Enable the device to use dead peer detection (DPD). DPD is a method used by devices to verify the current existence and availability of IPsec peers. A device performs this verification by sending encrypted IKE Phase 1 notification payloads (R-U-THERE messages) to a peer and waiting for DPD acknowledgements (R-U-THERE-ACK messages) from the peer.

Options

interval—Specify the amount of time that the peer waits for traffic from its destination peer before sending a dead-peer-detection (DPD) request packet.

Default: 10 seconds

Range: 2 through 60 seconds

always-send—Instructs the device to send dead peer detection (DPD) requests regardless of whether there is outgoing IPsec traffic to the peer.

optimized—Send dead peer detection (DPD) messages if there is no incoming IKE or IPsec traffic within the configured interval after outgoing packets are sent to the peer. This is the default DPD mode.

probe-idle-tunnel—Send dead peer detection (DPD) messages during idle traffic time between peers.

threshold—Specify the maximum number of unsuccessful dead peer detection (DPD) requests to be sent before the peer is considered unavailable.

Default: 5

Range: 1 through 5

Required Privilege Level

security—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding AutoVPN

IPsec VPN Overview

default-chap-secret

Syntax

```
default-chap-secret name;
```

Hierarchy Level

```
[edit interfaces interface-name ppp-options chap],
[edit interfaces interface-name unit logical-unit-number ppp-options chap],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-options chap]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Description

Define the default CHAP secret to be used when no matching CHAP access profile exists.

For ATM2 IQ interfaces only, you can configure a default CHAP secret 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.

Default

If you do not include the **default-chap-secret** statement in the configuration, and an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a Default CHAP Secret

access-profile

default-pap-password

Syntax

```
default-pap-password password;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number ppp-options pap],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number ppp-options pap]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Description

For PAP authentication, the default PAP password.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a Default PAP Password

access-profile

dhcp-attributes (Address-Assignment Pools)

Syntax

```
dhcp-attributes {
  boot-file filename;
  boot-server (address | hostname);
  dns-server [ ipv6-address ];
  domain-name domain-name;
  exclude-prefix-len exclude-prefix-length;
  grace-period seconds;
  maximum-lease-time seconds;
  name-server [ server-list ];
  netbios-node-type node-type;
  option {
    [ (id-number option-type option-value)
      (id-number array option-type option-value) ];
  }
  option-match {
    option-82 {
      circuit-id value range named-range;
      remote-id value range named-range;
    }
  }
  preferred-lifetime seconds;
  router [ router-address ];
  server-identifier ip4-address;
  sip-server-address [ ipv6-address ];
  sip-server-domain-name domain-name;
  t1-percentage percentage;
  t1-renewal-time;
  t2-percentage percentage;
  t2-rebinding-time;
  tftp-server address;
  valid-lifetime seconds;
  wins-server [ servers ];
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family family]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.3 for EX Series switches.

exclude-prefix-len statement introduced in Junos OS Release 17.3 for MX Series.

Description

Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

Options

boot-file—(EX Series, M Series, MX Series, SRX Series, T Series only) Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup. This configuration is equivalent to DHCP option 67.

Values: *filename*—Location of the boot file on the boot server. The filename can include a pathname.

boot-server—(EX Series, M Series, MX Series, SRX Series, T Series only) Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup. This configuration is equivalent to DHCP option 66.

Values:

- *address*—IPv4 address of a boot server.
- *hostname*—Fully qualified hostname of a boot server.

dns-server—(MX Series only) Specify a DNS server to which clients can send DNS queries. This is equivalent to DHCPv6 option 23. To specify multiple DNS servers, add multiple **dns-server** statements in order of preference.

Values: *ipv6-address*—IPv6 address of a DNS server.

domain-name—(EX Series, MX Series only) Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.

Values: *domain-name*—Name of the domain.

exclude-prefix-len *exclude-prefix-length*—Specify the length of the IPv6 prefix to be excluded from the delegated prefix. Range: 1 through 128.

grace-period—(M Series, MX Series, SRX Series, T Series only) Configure the amount of time that the client retains the address lease after the lease expires. The address cannot be reassigned to another client during the grace period.

Values: *seconds*—Number of seconds the lease is retained.

Range: 0 through 4,294,967,295 seconds.

Default: 0 (no grace period).

maximum-lease-time—(EX Series, M120, MX Series, SRX240, SRX3400, T640, T1600 only) Specify the maximum length of time, in seconds, that the lease is held for a client if the client does not renew the lease. This is equivalent to DHCP option 51. The **maximum-lease-time** is mutually exclusive with both the **preferred-lifetime** and the **valid-lifetime**, and cannot be configured with either timer.

Values: *seconds*—Maximum number of seconds the lease can be held.

Range: 30 through 4,294,967,295 seconds.

Default: 86,400 (24 hours).

name-server—(M120, MX Series, SRX Series, T640, T1600 only) Configure one or more Domain Name System (DNS) name servers available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

Values: *server-names*—IP addresses of the domain name servers, listed in order of preference.

netbios-node-type—(M Series, MX Series, SRX Series, T Series only) Specify the NetBIOS node type. This is equivalent to DHCP option 46.

Values: *node-type*—One of the following node types:

- *b-node*—Broadcast node.
- *h-node*—Hybrid node.
- *m-node*—Mixed node.
- *p-node*—Peer-to-peer node.

option—(EX Series, M Series, MX Series, SRX Series, T Series only) Specify user-defined options that are added to client packets. Starting in Junos OS Release 13.3, the **hex-string** option type was introduced.

Values:

- *array*—An option can include an array of option types.
- *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
- *option-type*—Any of the following types: byte, byte-stream, flag, hex-string, integer, ip-address, short, string, unsigned-integer, or unsigned-short.
- *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an On or Off value for a flag type).

preferred-lifetime—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix active. When the lifetime expires, the address is deprecated. If the **valid-lifetime** is also configured, the **preferred-lifetime** must be less than the **valid-lifetime**. The **preferred-lifetime** and the **maximum-lease-time** are mutually exclusive and cannot both be configured.

Values: *seconds*—Number of seconds that the IPv6 prefix is active.

Range: 30 through 4,294,967,295 seconds.

Default: 86,400 (24 hours).

router—(EX Series, MX Series only) Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

Values: *router-address*—IP address of one or more routers.

server-identifier—(EX Series, MX Series, SRX Series only) Specify the IP address that is used as the source address the DHCP server includes in IP packets when communicating with clients. The address is included in the DHCP packet in option 54.

Values: *ipv4-address*—IP address.

sip-server-address—(EX Series, M Series, MX Series, SRX Series, T Series only) Specify a SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 22. To specify multiple servers, add multiple **sip-server-address** statements in order of preference.

Values: *ipv6-address*—IPv6 address of a SIP outbound proxy server.

sip-server-domain-name—(M Series, MX Series, SRX Series, T Series only) Configure the domain name of the SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 21.

Values: *domain-name*—Name of the domain.

t1-percentage—(EX Series, M Series, MX Series only) Specify a percentage of the **preferred-lifetime** value. After this percentage of the **preferred-lifetime** value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from the originating DHCP local server. The **t1-percentage** is also referred to as the renewal time. The **t1-percentage** value must be less than the **t2-percentage** value. DHCPv4 server support was added in Junos OS Release 17.2.

Values: *percentage*—Percentage of the **preferred-lifetime** value.

Range: 0 through 100.

Default: If the **t1-percentage** value is not configured, the default is based on the **preferred-lifetime** value:

- If the **preferred-lifetime** value is finite, the default is 50 percent of the **preferred-lifetime** value.
- If the **preferred-lifetime** value is infinite, the default is also infinite.

t1-renewal-time—(MX Series only) Starting in Junos OS Release 17.2, specify the time (T1) at which the DHCPv4 or DHCPv6 client requests an extension (renewal) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the **t1-percentage** statement.

Values: *seconds*—Number of seconds.

Range: 30 through 4,294,967,295 seconds.

Default: 50 percent of the lease duration (**preferred-lifetime**).

t2-percentage—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify a percentage of the **preferred-lifetime** value. After this percentage of the **preferred-lifetime** value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from any available DHCPv4 or DHCPv6 server. The **t2-percentage** is also referred to as the rebinding time. The **t2-percentage** value must be greater than the **t1-percentage** value. DHCPv4 server support was added in Junos OS Release 17.2.

Values: *percentage*—Percentage of the **preferred-lifetime** value.

Range: 0 through 100.

Default: Default: If the **t2-percentage** value is not configured, the default is based on the **preferred-lifetime** value:

- If the **preferred-lifetime** value is finite, the default is 80 percent of the **preferred-lifetime** value.
- When the **preferred-lifetime** value is infinite, the default is also infinite.

t2-rebinding-time—(MX Series only) Starting in Junos OS Release 17.2, specify the time (T2) at which the DHCPv4 or DHCPv6 client attempts to contact any DHCP server to request an extension (rebinding) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the **t2-percentage** statement.

Values: *seconds*—Number of seconds.

Range: 30 through 4,294,967,295 seconds.

Default: The default value depends on the client:

- (DHCPv4 clients) 87.5 percent of the lease duration (**preferred-lifetime**).
- (DHCPv6 clients) 80 percent of the lease duration (**preferred-lifetime**).

tftp-server—(ACX Series, SRX Series only) Specify the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file. This is equivalent to DHCP option 150.

Values: *ip-address*—IP address of the TFTP server.

valid-lifetime—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix valid. When the lifetime expires, the address becomes invalid. If the **preferred-lifetime** is also configured, the **valid-lifetime** must be greater than the **preferred-lifetime**. The **valid-lifetime** and the **maximum-lease-time** are mutually exclusive and cannot both be configured.

Values: *seconds*—Number of seconds that the IPv6 prefix is valid.

Range: 30 through 4,294,967,295 seconds.

Default: 86,400 (24 hours).

wins-server—(M Series, MX Series, SRX Series, T Series only) Specify one or more NetBIOS name servers (NBNS) that the client uses to resolve NetBIOS names. This is equivalent to DHCP option 44.

Values: *ipv4-address*—IP address of each NetBIOS name server. Add them to the configuration in order of preference.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>Address-Assignment Pools Overview</i>
<i>Attributes That Can Be Applied to DHCP Clients</i>
<i>Address-Assignment Pool Configuration Overview</i>
<i>Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address</i>
<i>DHCP Lease Timers</i>
<i>Configuring DHCP Attributes for All Clients or a Group of Clients</i>

encapsulation-overhead

Syntax

```
encapsulation-overhead bytes;
```

Hierarchy Level

```
[edit access group-profile profile-name ppp],  
[edit access profile profile-name client client-name ppp]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Configure the encapsulation overhead for class-of-service calculations.

Options

bytes—The number of bytes used as encapsulation overhead for the session.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

exclude (RADIUS Attributes)

Syntax

```
exclude {
  acc-aggr-cir-id-asc [ access-request | accounting-start | accounting-stop ];
  acc-aggr-cir-id-bin [ access-request | accounting-start | accounting-stop ];
  acc-loop-cir-id [ access-request | accounting-start | accounting-stop ];
  acc-loop-encap [ access-request | accounting-start | accounting-stop ];
  acc-loop-remote-id [ access-request | accounting-start | accounting-stop ];
  accounting-authentic [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  accounting-delay-time [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  accounting-session-id access-request;
  accounting-terminate-cause accounting-off;
  acct-request-reason [ accounting-start | accounting-stop ];
  acct-tunnel-connection [ access-request | accounting-start | accounting-stop ];
  act-data-rate-dn [ access-request | accounting-start | accounting-stop ];
  act-data-rate-up [ access-request | accounting-start | accounting-stop ];
  act-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
  act-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
  att-data-rate-dn [ access-request | accounting-start | accounting-stop ];
  att-data-rate-up [ access-request | accounting-start | accounting-stop ];
  called-station-id [ access-request | accounting-start | accounting-stop ];
  calling-station-id [ access-request | accounting-start | accounting-stop ];
  chargeable-user-identity access-request;
  class [ accounting-start | accounting-stop ];
  cos-shaping-rate [ accounting-start | accounting-stop ];
  delegated-ipv6-prefix [ accounting-start | accounting-stop ];
  dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
  dhcp-header access-request;
  dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
  dhcp-options [ access-request | accounting-start | accounting-stop ];
  dhcpv6-header access-request;
  dhcpv6-options [ access-request | accounting-start | accounting-stop ];
  downstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
  dsl-forum-attributes [ access-request | accounting-start | accounting-stop ];
  dsl-line-state [ access-request | accounting-start | accounting-stop ];
  dsl-type [ access-request | accounting-start | accounting-stop ];
  dynamic-iflset-name [ accounting-start | accounting-stop ];
  event-timestamp [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  filter-id [ accounting-start | accounting-stop ];
  first-relay-ipv4-address [ access-request | accounting-start | accounting-stop ];
  first-relay-ipv6-address [ access-request | accounting-start | accounting-stop ];
  framed-interface-id [ access-request | accounting-start | accounting-stop ];
  framed-ip-address [ access-request | accounting-start | accounting-stop ];
```



```

framed-ip-netmask [ access-request | accounting-start | accounting-stop ];
framed-ip-route [ accounting-start | accounting-stop ];
framed-ipv6-address [ access-request | accounting-start | accounting-stop ];
framed-ipv6-pool [ accounting-start | accounting-stop ];
framed-ipv6-prefix [ accounting-start | accounting-stop ];
framed-ipv6-route [ accounting-start | accounting-stop ];
framed-pool [ accounting-start | accounting-stop ]; input-ipv6-gigawords accounting-stop;
input-filter [ accounting-start | accounting-stop ];
input-gigapackets accounting-stop;
input-gigawords accounting-stop;
input-ipv6-octets accounting-stop;
input-ipv6-packets accounting-stop;
interface-description [ access-request | accounting-start | accounting-stop ];
l2c-downstream-data [ access-request | accounting-start | accounting-stop ];
l2c-upstream-data [ access-request | accounting-start | accounting-stop ];
l2tp-rx-connect-speed [ access-request | accounting-start | accounting-stop ];
l2tp-tx-connect-speed [ access-request | accounting-start | accounting-stop ];
max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
max-data-rate-up [ access-request | accounting-start | accounting-stop ];
max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
min-data-rate-up [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
nas-identifier [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
nas-port [ access-request | accounting-start | accounting-stop ];
nas-port-id [ access-request | accounting-start | accounting-stop ];
nas-port-type [ access-request | accounting-start | accounting-stop ];
output-filter [ accounting-start | accounting-stop ];
output-gigapackets accounting-stop;
output-gigawords accounting-stop;
output-ipv6-gigawords accounting-stop;
output-ipv6-octets accounting-stop;
output-ipv6-packets accounting-stop;
pppoe-description [ access-request | accounting-start | accounting-stop ];
standard-attribute number {
    packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
}
tunnel-assignment-id [ access-request | accounting-start | accounting-stop ];
tunnel-client-auth-id [ access-request | accounting-start | accounting-stop ];
tunnel-client-endpoint [ access-request | accounting-start | accounting-stop ];
tunnel-medium-type [ access-request | accounting-start | accounting-stop ];
tunnel-server-auth-id [ access-request | accounting-start | accounting-stop ];

```

```

tunnel-server-endpoint [ access-request | accounting-start | accounting-stop ];
tunnel-type [ access-request | accounting-start | accounting-stop ];
upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
vendor-id id-number {
    vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
}
virtual-router [ access-request | accounting-start | accounting-stop ];
}

```

Hierarchy Level

[edit access profile *profile-name* radius [attributes](#)]

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

downstream-calculated-qos-rate, **dsl-forum-attributes**, and **upstream-calculated-qos-rate** options added in Junos OS Release 11.4.

cos-shaping-rate and **filter-id** options added in Junos OS Release 13.2.

pppoe-description option added in Junos OS Release 14.2.

virtual-router option added in Junos OS Release 15.1.

first-relay-ipv4-address and **first-relay-ipv6-address** options added in Junos OS Release 16.1.

acc-loop-encap and **acc-loop-remote-id** options added in Junos OS Release 16.1R4.

access-request option support for all tunnel attributes added in Junos OS Release 15.1R7, 16.1R5, 16.2R2, 17.1R2, 17.2R2, and 17.3R1 for MX Series.

packet-type, **standard-attribute**, **vendor-attribute**, and **vendor-id** options added in Junos OS Release 18.1R1.

Description

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the **ignore** statement.

You can specify attribute exclusion for multiple RADIUS message types by enclosing the message types, separated by spaces, within brackets ([]). You do not need brackets when specifying a single message type.

Starting in Junos OS Release 18.1R1, you can specify standard RADIUS attributes with the attribute number. You can specify VSAs with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be excluded. The configuration has no effect if you configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to exclude only a subset of all attributes that can be received in Access-Accept messages.

Not all attributes are available in all types of RADIUS messages.

NOTE: If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

NOTE: VSAs with dedicated option names include Juniper Networks (IANA vendor ID 4874) and DSL Forum (vendor ID 3561) VSAs.

Options

RADIUS attribute—RADIUS standard attribute or VSA:

- **acc-aggr-cir-id-asc**—Exclude Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- **acc-aggr-cir-id-bin**—Exclude Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- **acc-loop-cir-id**—Exclude Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- **acc-loop-encap**—Exclude Juniper Networks VSA 26-183, Acc-Loop-Encap.
- **acc-loop-remote-id**—Exclude Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- **accounting-authentic**—Exclude RADIUS attribute 45, Acct-Authentic.
- **accounting-delay-time**—Exclude RADIUS attribute 41, Acct-Delay-Time.
- **accounting-session-id**—Exclude RADIUS attribute 44, Acct-Session-Id.
- **accounting-terminate-cause**—Exclude RADIUS attribute 49, Acct-Terminate-Cause.
- **acct-request-reason**—Exclude Juniper Networks VSA 26-210, Acct-Request-Reason.
- **acct-tunnel-connection**—Exclude RADIUS attribute 68, Acct-Tunnel-Connection.
- **act-data-rate-dn**—Exclude Juniper Networks VSA 26-114, Act-Data-Rate-Dn.
- **act-data-rate-up**—Exclude Juniper Networks VSA 26-113, Act-Data-Rate-Up.
- **act-interlv-delay-dn**—Exclude Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn.
- **act-interlv-delay-up**—Exclude Juniper Networks VSA 26-124, Act-Interlv-Delay-Up.
- **att-data-rate-dn**—Exclude Juniper Networks VSA 26-118, Att-Data-Rate-Dn.
- **att-data-rate-up**—Exclude Juniper Networks VSA 26-117, Att-Data-Rate-Up.
- **called-station-id**—Exclude RADIUS attribute 30, Called-Station-Id.
- **calling-station-id**—Exclude RADIUS attribute 31, Calling-Station-Id.
- **chargeable-user-identity**—Exclude RADIUS attribute 89, Chargeable-User-Identity.
- **class**—Exclude RADIUS attribute 25, Class.
- **cos-shaping-rate**—Exclude Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- **delegated-ipv6-prefix**—Exclude RADIUS attribute 123, Delegated-IPv6-Prefix.
- **dhcp-gi-address**—Exclude Juniper Networks VSA 26-57, DHCP-GI-Address.
- **dhcp-header**—Exclude Juniper Networks VSA 26-208, DHCP-Header.
- **dhcp-mac-address**—Exclude Juniper Networks VSA 26-56, DHCP-MAC-Address.
- **dhcp-options**—Exclude Juniper Networks VSA 26-55, DHCP-Options.
- **dhcpv6-header**—Exclude Juniper Networks VSA 26-209, DHCPv6-Header.
- **dhcpv6-options**—Exclude Juniper Networks VSA 26-207, DHCPv6-Options.

- **dynamic-iflset-name**—Exclude Juniper Networks VSA 26-130, Qos-Set-Name.
- **downstream-calculated-qos-rate**—Exclude Juniper Networks VSA 26-141.
- **dsl-forum-attributes**—Exclude DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*.
- **dsl-line-state**—Exclude Juniper Networks VSA 26-127, DSL-Line-State.
- **dsl-type**—Exclude Juniper Networks VSA 26-128, DSL-Type.
- **event-timestamp**—Exclude RADIUS attribute 55, Event-Timestamp.
- **filter-id**—Exclude RADIUS attribute 11, Filter-Id.
- **first-relay-ipv4-address** —Exclude Juniper Networks VSA 26-189, DHCP-First-Relay-IPv4-Address.
- **first-relay-ipv6-address** —Exclude Juniper Networks VSA 26-190, DHCP-First-Relay-IPv6-Address.
- **framed-interface-id**—Exclude RADIUS attribute 96, Framed-Interface-ID.
- **framed-ip-address**—Exclude RADIUS attribute 8, Framed-IP-Address.
- **framed-ip-netmask**—Exclude RADIUS attribute 9, Framed-IP-Netmask.
- **framed-ip-route**—Exclude RADIUS attribute 22, Framed-Route.
- **framed-ipv6-address**—Exclude RADIUS attribute 168, Framed-IPv6-Address.
- **framed-ipv6-pool**—Exclude RADIUS attribute 100, Framed-IPv6-Pool.
- **framed-ipv6-prefix**—Exclude RADIUS attribute 97, Framed-IPv6-Prefix.
- **framed-ipv6-route**—Exclude RADIUS attribute 99, Framed-IPv6-Route.
- **framed-pool**—Exclude RADIUS attribute 88, Framed-Pool.
- **input-filter**—Exclude Juniper Networks VSA 26-10, Ingress-Policy-Name.
- **input-gigapackets**—Exclude Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- **input-gigawords**—Exclude RADIUS attribute 52, Acct-Input-Gigawords.
- **input-ipv6-gigawords**—Exclude Juniper Networks VSA 26-155, Acct-Input-IPv6-Gigawords.
- **input-ipv6-octets**—Exclude Juniper Networks VSA 26-151, Acct-Input-IPv6-Octets.
- **input-ipv6-packets**—Exclude Juniper Networks VSA 26-153, Acct-Input-IPv6-Packets.
- **interface-description**—Exclude Juniper Networks VSA 26-53, Interface-Desc.
- **l2c-downstream-data**—Exclude Juniper Networks VSA 26-93, L2C-Down-Stream-Data.
- **l2c-upstream-data**—Exclude Juniper Networks VSA 26-92, L2C-Up-Stream-Data.
- **l2tp-rx-connect-speed**—Exclude Juniper Networks VSA 26-163, Rx-Connect-Speed.
- **l2tp-tx-connect-speed**—Exclude Juniper Networks VSA 26-162, Tx-Connect-Speed.
- **max-data-rate-dn**—Exclude Juniper Networks VSA 26-120, Max-Data-Rate-Dn.

- **max-data-rate-up**—Exclude Juniper Networks VSA 26-119, Max-Data-Rate-Up.
- **max-interlv-delay-dn**—Exclude Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn.
- **max-interlv-delay-up**—Exclude Juniper Networks VSA 26-123, Max-Interlv-Delay-Up.
- **min-data-rate-dn**—Exclude Juniper Networks VSA 26-116, Min-Data-Rate-Dn.
- **min-data-rate-up**—Exclude Juniper Networks VSA 26-115, Min-Data-Rate-Up.
- **min-lp-data-rate-dn**—Exclude Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn.
- **min-lp-data-rate-up**—Exclude Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up.
- **nas-identifier**—Exclude RADIUS attribute 32, NAS-Identifier.
- **nas-port**—Exclude RADIUS attribute 5, NAS-Port.
- **nas-port-id**—Exclude RADIUS attribute 87, NAS-Port-Id.
- **nas-port-type**—Exclude RADIUS attribute 61, NAS-Port-Type.
- **output-filter**—Exclude Juniper Networks VSA 26-11, Egress-Policy-Name.
- **output-gigapackets**—Exclude Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- **output-gigawords**—Exclude RADIUS attribute 53, Acct-Output-Gigawords.
- **output-ipv6-gigawords**—Exclude Juniper Networks VSA 26-156, Acct-Output-IPv6-Gigawords.
- **output-ipv6-octets**—Exclude Juniper Networks VSA 26-152, Acct-Output-IPv6-Octets.
- **output-ipv6-packets**—Exclude Juniper Networks VSA 26-154, Acct-Output-IPv6-Packets.
- **packet-type**—Specify the RADIUS message type to exclude; term required when excluding a standard attribute or VSA by number rather than name. You can enclose multiple values in square brackets to specify a list of message types. Message types include Access-Request, Accounting-Off, Accounting-Off, Accounting-Start, and Accounting-Stop.
- **pppoe-description**—Exclude Juniper Networks VSA 26-24, PPPoE-Description.
- **standard-attribute *number***—RADIUS standard attribute number supported by your platform. If you configure an unsupported attribute, that configuration has no effect. When you use this option, you must use the **packet-type** term to specify the message from which the attribute is excluded.
- **tunnel-assignment-id**—Exclude RADIUS attribute 82, Tunnel-Assignment-ID.
- **tunnel-client-auth-id**—Exclude RADIUS attribute 90, Tunnel-Client-Auth-ID.
- **tunnel-client-endpoint**—Exclude RADIUS attribute 66, Tunnel-Client-Endpoint.
- **tunnel-medium-type**—Exclude RADIUS attribute 65, Tunnel-Medium-Type.
- **tunnel-server-auth-id**—Exclude RADIUS attribute 91, Tunnel-Server-Auth-ID.
- **tunnel-server-endpoint**—Exclude RADIUS attribute 67, Tunnel-Server-Endpoint.
- **tunnel-type**—Exclude RADIUS attribute 64, Tunnel-Type.

- **upstream-calculated-qos-rate**—Exclude Juniper Networks VSA 26-142
- **vendor-attribute** *vsa-number*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. If you configure an unsupported VSA, that configuration has no effect. When you use this option, you must use the **packet-type** term to specify the message from which the attribute is excluded.
- **vendor-id** *id-number*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect.
- **virtual-router**—Exclude Juniper Networks VSA 26-1.

RADIUS message type:

- **access-request**—RADIUS Access-Request messages.
- **accounting-off**—RADIUS Accounting-Off messages.
- **accounting-on**—RADIUS Accounting-On messages.
- **accounting-start**—RADIUS Accounting-Start messages.
- **accounting-stop**—RADIUS Accounting-Stop messages.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Filtering RADIUS Attributes and VSAs from RADIUS Messages

RADIUS Servers and Parameters for Subscriber Access

Standard and Vendor-Specific RADIUS Attributes

framed-pool

Syntax

```
framed-pool framed-pool;
```

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 address pool.

Options

framed-pool—References a configured address pool.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

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 {
      aaa-options aaa-options-name;
      chap;
      ignore-magic-number-mismatch;
      initiate-ncp (ip | ipv6 | dual-stack-passive)
      ipcp-suggest-dns-option;
      mru;
      mtu;
      pap;
      peer-ip-address-optional;
    }
    primary-dns primary-dns;
    primary-wins primary-wins;
    secondary-dns secondary-dns;
    secondary-wins secondary-wins;
  }
}
```

Hierarchy Level

[edit access]

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the Group Profile for Defining L2TP Attributes | 24](#)

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

host (Address-Assignment Pools)

Syntax

```
host hostname {
    hardware-address mac-address;
    ip-address ip-address;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Configure a static binding for the specified client.

Options

hostname—Name of the client.

hardware-address *mac-address*—(M Series, MX Series, SRX Series, T Series only) Specify the MAC address of the client. This is the hardware address that identifies the client on the network.

- **mac-address**—MAC address of the client.

ip-address *ip-address*—(SRX Series, T Series only) Specify the reserved IP address assigned to the client.

- **ip-address**—IP version 4 (IPv4) address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Address-Assignment Pools Overview](#)

[Address-Assignment Pool Configuration Overview](#)

[Configuring Static Address Assignment](#)

idle-timeout (Access)

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit access group-profile profile-nameppp ppp],  
[edit access profile profile-name client client-nameppp ]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

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 PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

ike (Access Profile)

Syntax

```
ike {
  allowed-proxy-pair {
    remote remote-proxy-address local local-proxy-address;
  }
  dead-peer-detection {
    interval seconds
    threshold number
  }
  ike-policy policy-name;
  initiate-dead-peer-detection;
  interface-id string-value;
  ipsec-policy ipsec-policy;
  pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
  reverse-route
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 7.4.

ike-policy statement introduced in Junos OS Release 8.2.

Description

Configure an IKE access profile.

The remaining statements are explained below.

NOTE: This statement is not supported on MX Series routers.

Options

allowed-proxy-pair—(M Series, MX Series, PTX Series, T Series only) Specify the network address of the local and remote peer associated with an IKE access profile.

Values:

- local *local-proxy-address*—Network address of the local peer. Default: 0.0.0.0

- remote *remote-proxy-address*—Network address of the remote peer. Default: 0.0.0.0

ike-policy *policy-name*—(MX Series, SRX Series only) Specify the IKE policy used to authenticate dynamic peers during IKE negotiation.

Values:

- *policy-name*—The name of an IKE policy configured at the **[edit services ipsec-vpn ike policy *policy-name*]** hierarchy level. The IKE policy defines either the local digital certificate or the pre-shared key used for IKE authentication with dynamic peers. For more information about how to configure the IKE policy, see the *Junos OS Services Interfaces Library for Routing Devices*.

initiate-dead-peer-detection—(MX Series, T Series only) Detect inactive peers on dynamic IPsec tunnels.

pre-shared-key (*ascii-text character-string* | hexadecimal *hexadecimal-digits*)—(M Series, MX Series, T Series only) Configure the key used to authenticate a dynamic peer during IKE phase 1 negotiation. Specify the key in either ASCII or hexadecimal format.

Values:

- *ascii-text character-string*—Authentication key in ASCII format.
- hexadecimal *hexadecimal-digits*—Authentication key in hexadecimal format.

reverse-route—(M Series and MX Series routers with an AS or MultiServices PIC only) Starting in Junos OS Release 10.4, configure a reverse route for dynamic endpoint IPsec tunnels.

Values: preference *metric-value*

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring an IKE Access Profile | 33](#)

Junos OS Services Interfaces Library for Routing Devices

Services Interface Naming Overview

immediate-update

Syntax

```
immediate-update;
```

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.

Description

Configure the router or switch to send an Acct-Update message to the RADIUS accounting server on receipt of a response (for example, an ACK or timeout) to the Acct-Start message.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Per-Subscriber Session Accounting

RADIUS Servers and Parameters for Subscriber Access

interface-description-format

Syntax

```
interface-description-format {
  exclude-adapter;
  exclude-channel;
  exclude-sub-interface;
}
```

Hierarchy Level

[edit access profile *profile-name* radius [options](#)]

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

exclude-adapter and **exclude-sub-interface** options added in Junos OS Release 10.4.

exclude-channel option added in Junos OS Release 17.3R1.

Description

Specify the information that is excluded from the interface description that the device passes to RADIUS for inclusion in the RADIUS attributes such as NAS-Port-ID (87) or Calling-Station-ID (31).

The default format for nonchannelized interfaces is as follows:

interface-type-slot/adapter/port.subinterface[:svlan-vlan]

For example, consider physical interface ge-1/2/0, with a subinterface of 100 and SVLAN identifier of 100. The interface description used in the NAS-Port-ID is ge-1/2/0.100:100. If you exclude the subinterface, the description becomes ge-1/2/0:100.

The default format for channelized interfaces is as follows:

interface-type-slot/adapter/channel.subinterface[:svlan-vlan]

The channel information (logical port number) is determined by this formula:

Logical port number = 100 + (*actual-port-number* × 20) + *channel-number*.

For example, consider a channelized interface 3 on port 2 where the:

- Physical interface is xe-0/1/2:3.
- Subinterface is 4.
- SVLAN is 5.

- VLAN is 6.

Using the formula, the logical port number = $100 + (2 \times 20) + 3 = 143$. Consequently, the default interface description is xe-0/1/143.4-5.6. If you exclude the channel information, the description becomes xe-0/1/2.4-5.6.

Options

exclude-adapter—(Optional) Exclude the adapter from the interface description.

exclude-channel—(Optional) Exclude the channel information from the interface description.

exclude-sub-interface—(Optional) Exclude the subinterface from the interface description.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Interface Text Descriptions for Inclusion in RADIUS Attributes

RADIUS Servers and Parameters for Subscriber Access

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 L2TP for a Group Profile | 25](#)

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

[Configuring an IKE Access Profile | 33](#)

Configuring an L2TP Access Profile on the LNS

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 PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

keepalive-retries

Syntax

```
keepalive-retries number-of-retries;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Configure the number of retry attempts for checking the keepalive status of a Point-to-Point (PPP) protocol session. Configure this setting to reduce the detection time for PPP client session timeouts or failures if you have configured the keepalive timeout interval (using the **keepalive** statement).

Options

number-of-retries—The maximum number of retries the L2TP network server (LNS) attempts by sending LCP echo requests to the peer to check the keepalive status of the PPP session. If there is no response from the PPP client within the specified number of retries, the PPP session is considered to have timed out.

Range: 3 through 32,767 times

Default: 10 times

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

[keepalive | 173](#)

I2tp (Group Profile)

Syntax

```
I2tp {  
  interface-id interface-id;  
  lcp-renegotiation;  
  local-chap;  
  maximum-sessions-per-tunnel number;  
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the Layer 2 Tunneling Protocol for a group profile.

NOTE: This statement is not supported 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 L2TP for a Group Profile | 25](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Configuring an L2TP Access Profile on the LNS

I2tp (Profile)

Syntax

```

l2tp {
  interface-id interface-id;
  lcp-renegotiation;
  local-chap;
  maximum-sessions number;
  maximum-sessions-per-tunnel number;
  multilink {
    drop-timeout milliseconds;
    fragment-threshold bytes;
  }
  override-result-code session-out-of-resource;
  ppp-authentication (chap | pap);
  ppp-profile profile-name;
  sessions-limit-group;
  service-profile profile-name(parameter)&profile-name;
  shared-secret shared-secret;
}

```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the L2TP properties for a profile.

NOTE: Only the **interface-id**, **lcp-renegotiation**, **maximum-sessions**, **maximum-sessions-per-tunnel**, **sessions-limit-group** and **shared-secret** statements are supported for L2TP LNS on MX Series routers.

Options

interface-id—Configure the interface identifier.

Values:

- **interface-id**—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces interface-name unit local-unit-number dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

lcp-renegotiation—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.

local-chap—Configure the Junos OS so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

maximum-sessions—Specify the maximum number of L2TP sessions for the chassis, all tunnels, a tunnel group, a session limit group, or a client.

Values:

- **number**—Number of sessions allowed.
- Range: (Chassis, tunnel group, session limit group, or client) 1 through the default maximum chassis limit
- Range: (Tunnel) 1 through 65,536

maximum-sessions-per-tunnel—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.

Values:

- *number*—Maximum number of sessions for a Layer 2 tunnel.

multilink—Configure Multilink PPP for Layer 2 Tunneling Protocol (L2TP).

The options for this statement are explained separately. Click the linked statement for details.

override-result-code—Configure the LNS to override result codes in Call-Disconnect-Notify (CDN) messages.

Values:

- *session-out-of-resource*—Override result codes 4 and 5 with result code 2. These result codes indicate that the number of L2TP sessions have reached the configured maximum value and the LNS can support no more sessions. When the LAC receives the code, it fails over to another LNS to establish subsequent sessions. Some third-party LACs respond only to result code 2.

ppp-authentication—(T Series only) Configure PPP authentication.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Values:

- *chap*—Challenge Handshake Authentication Protocol.
- *pap*—Password Authentication Protocol.

ppp-profile—(M Series, T Series only) Specify the profile used to validate PPP session requests through L2TP tunnels.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Values: *profile-name*—Identifier for the PPP profile.

sessions-limit-group—(MX Series only) Starting in Junos OS Release 16.1, specify in an L2TP access profile the session limit group to which a client is assigned by the profile.

Values: *limit-group-name*—Identifier of the session-limit group to which a client is assigned.

service-profile—Configure one or more dynamic service profiles to be applied to subscriber sessions at activation for all subscribers in the specified tunnel group or on the specified LAC. Services are typically applied to L2TP sessions with RADIUS VSAs or CoA requests. In multivendor environments, you might use only standard attributes to simplify management of multiple vendor VSAs. This statement enables you to apply services without using an external authority such as RADIUS. The locally configured list of services (service profiles) serves as local authorization that is applied by authd during client session activation. This list of services is subject to the same validation and processing as services originating from an external authority, such as RADIUS.

You can optionally specify parameters that are passed to the corresponding service when it is activated for the session. The parameter might override values configured in the profile itself, such as a downstream shaping rate for a CoS service. This enables you to use the same service profile for multiple situations with different requirements, or to modify a previously applied value for a service.

You can still use RADIUS VSAs or CoA requests together with the service profiles. If services are sourced from an external authority as authorization during authentication or during subscriber session provisioning (activation), the services from the external authority take strict priority over those in the local configuration. If a service applied with RADIUS is the same as a service applied with a service profile in the CLI, but with different parameters, the RADIUS service is applied with a new session ID and takes precedence over the earlier service profile.

When service profiles are configured on a LAC client and on a tunnel group that uses that LAC client, the LAC configuration overrides the tunnel group configuration. Only the service profile configured on the LAC client is applied to subscribers in the tunnel group.

Values:

- *profile-name*—Name of a dynamic service profile that defines a service to be applied to L2TP subscriber sessions. You can specify one or more service profiles, separated by an ampersand (&).
- *parameter*—(Optional) Value to be passed to the service when it is activated on the subscriber session.

shared-secret—Configure the shared secret.

Values:

- *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 | 39](#)[Configuring an L2TP Access Profile on the LNS](#)[Limiting the Number of L2TP Sessions Allowed by the LAC or LNS](#)[Configuring an L2TP LAC](#)[Configuring an L2TP LNS with Inline Service Interfaces](#)[Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces](#)[L2TP for Subscriber Access Overview](#)

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 L2TP for a Group Profile | 25](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Configuring an L2TP Access Profile on the LNS

local-chap

Syntax

```
local-chap;
```

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 Junos OS so that the LNS ignores proxy authentication attribute-value pairs (AVPs) from the L2TP access concentrator (LAC) and reauthenticates the PPP client using a Challenge Handshake Authentication Protocol (CHAP) challenge. When you do this, the LNS directly authenticates the PPP client.

NOTE: This statement is not supported 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 L2TP for a Group Profile | 25](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

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 L2TP for a Group Profile | 25](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Configuring an L2TP Access Profile on the LNS

multilink

Syntax

```
multilink {  
    drop-timeout milliseconds;  
    fragment-threshold bytes;  
}
```

Hierarchy Level

[edit access [profile](#) *profile-name* [client](#) *client-name* [l2tp](#)]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure Multilink PPP for Layer 2 Tunneling Protocol (L2TP).

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options

drop-timeout—(M Series, MX Series, PTX Series, T Series only) Configure the drop timeout for a multilink bundle.

Values:

- *milliseconds*—Number of milliseconds for the timeout that is associated with the first fragment on the reassembly queue. If the timeout expires before all the fragments have been collected, the fragments at the beginning of the reassembly queue are dropped. If the drop timeout is not specified, the Junos OS holds on to the fragments. (Fragments may still be dropped if the multilink reassembly algorithm determines that another fragment belonging to the packet on a reassembly queue has been lost.)

fragment-threshold—(M Series, MX Series, PTX Series, T Series only) Configure the fragmentation threshold for a multilink bundle.

Values:

- *bytes*—The maximum number of bytes in a packet. If a packet exceeds the fragmentation threshold, the Junos OS fragments it into two or more multilink fragments.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring L2TP Properties for a Client-Specific Profile](#) | 39

nas-port-extended-format

Syntax

```
nas-port-extended-format {
  adapter-width bits;
  ae-width bits;
  atm {
    adapter-width bits;
    port-width bits;
    slot-width bits;
    vci-width bits;
    vpi-width bits;
  }
  port-width bits;
  pw-width bits;
  slot-width bits;
  stacked-vlan-width bits;
  vlan-width bits;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

ae-width option added in Junos OS Release 12.1.

atm option added in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.

atm option supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

pw-width option added in Junos OS Release 15.1.

Description

Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width in bits of the fields in the NAS-Port attribute.

The NAS-Port attribute specifies the physical port number of the NAS that is authenticating the user, and is formed by a combination of the physical port's slot number, port number, adapter number, VLAN ID, and S-VLAN ID. The NAS-Port extended format specifies the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, aggregated, Ethernet, VLAN, and S-VLAN.

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

Options

adapter-width *width*—Number of bits in the adapter field.

ae-width *width*—(Ethernet subscribers only) Number of bits in the aggregated Ethernet identifier field.

atm—Specify width for fields for ATM subscribers.

port-width *width*—Number of bits in the port field.

pw-width *width*—(Ethernet subscribers only) Number of bits in the pseudowire field. Appears in the Cisco NAS-Port-Info AVP (100).

slot-width *width*—Number of bits in the slot field.

stacked-vlan-width *width*—Number of bits in the SVLAN ID field.

vci-width *width*—(ATM subscribers only) Number of bits in the ATM virtual circuit identifier (VCI) field.

vlan-width *width*—Number of bits in the VLAN ID field.

vpi-width *width*—(ATM subscribers only) Number of bits in the ATM virtual path identifier (VPI) field.

NOTE: The total of the widths must not exceed 32 bits, or the configuration will fail.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers

RADIUS Servers and Parameters for Subscriber Access

network

Syntax

```
network ip-prefix</prefix-length>;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Configure subnet information for an IPv4 address-assignment pool.

Options

ip-prefix—IP version 4 address or prefix value.

prefix-length—(Optional) Subnet mask.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Address-Assignment Pool Configuration Overview*

option-82 (Address-Assignment Pools)

Syntax

```
option-82 {
  circuit-id value range named-range;
  remote-id value range named-range;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes option-match],
[edit access protocol-attributes attribute-set-name option-match]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Specify the list of option 82 suboption match criteria used to select the named address range used for the client. The server matches the option 82 value in the user PDU to the specified option 82 match criteria and uses the named address range associated with the string.

Options

circuit-id—(EX Series, MX Series only) Configure the address-assignment pool named-range to use for a particular option 82 Agent Circuit ID value.

Values:

- *value*—String for the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) in DHCP packets.
- *range named-range*—Name of the address-assignment pool range to use.

remote-id—(SRX Series only) Specify the address-assignment pool named range to use based on the particular option 82 Agent Remote ID value.

Values:

- *range named-range*—Name of the address-assignment pool range to use.
- *value*—String for Agent Remote ID suboption (suboption 2) of the DHCP relay agent information option (option 82) in DHCP packets.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Address-Assignment Pool Configuration Overview](#)

option-match

Syntax

```
option-match {  
  option-82 {  
    circuit-id value range named-range;  
    remote-id value range named-range;  
  }  
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet dhcp-attributes],  
[edit access protocol-attributes attribute-set-name]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Specify a list of match criteria used to determine which named address range in the address-assignment pool to use. The extended DHCP local server matches this information to the match criteria specified in the client PDUs. For example, for option 82 match criteria, the server matches the option 82 value in the user PDU to the specified option 82 string and uses the named range associated with the string.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Address-Assignment Pool Configuration Overview](#)

options (Access Profile)

Syntax

```
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    nas-identifier;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
  client-authentication-algorithm (direct | round-robin);
  coa-dynamic-variable-validation;
  ethernet-port-type-virtual;
  interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
  }
  ip-address-change-notify message;
  juniper-access-line-attributes;
  nas-identifier identifier-value;
  nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
      adapter-width width;
      port-width width;
      pw-width width;
      slot-width width;
      vci-width width;
      vpi-width width;
    }
  }
  nas-port-id-delimiter delimiter-character;
  nas-port-id-format {
    agent-circuit-id;
```

```

agent-remote-id;
interface-description;
interface-text-description;
nas-identifier;
order {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    postpend-vlan-tags;
}
postpend-vlan-tags;
}
nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback;
remote-circuit-id-format {
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}

```

Hierarchy Level

[edit access profile *profile-name* [radius](#)]

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

juniper-dsl-attributes introduced in Junos OS Release 11.4.

nas-port-id-delimiter introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

calling-station-id-delimiter introduced in Junos OS Release 13.1.

ip-address-change-notify introduced in Junos OS Release 13.1.

coa-dynamic-variable-validation, **client-authentication-algorithm**, and **client-accounting-algorithm** introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

remote-circuit-id-delimiter, **remote-circuit-id-fallback**, and **remote-circuit-id-format** introduced in Junos OS Release 13.3R1 on MX Series.

chap-challenge-in-request-authenticator introduced in Junos OS Release 15.1.

nas-identifier introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

service-activation introduced in Junos OS Release 16.2.

juniper-access-line-attributes introduced in Junos OS Release 19.2R1

Description

Configure the options used by RADIUS authentication and accounting servers.

Options

accounting-session-id-format—(EX Series, MX Series only) Configure the format the router or switch uses to identify the accounting session. The default is **decimal**.

Values:

- decimal—Use the decimal format.
- description—Use the generic format, in the form: **jnpr interface-specifier:subscriber-session-id**.

calling-station-id-delimiter—(MX Series, T Series only) Starting in Junos OS Release 13.1, specify the character that the router uses as a separator between the concatenated values in the Calling-Station-ID (RADIUS IETF attribute 31) string. The router uses the delimiter when you configure more than one value in the **calling-station-id-format** statement. The default is the hash (#) character.

Values:

- *delimiter-character*—Character to use for the delimiter. You must enclose the delimiter character in quotation marks (" ").

chap-challenge-in-request-authenticator—(MX Series only) Starting in Junos OS Release 15.1, configure the **authd** process to insert the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets, if the challenge value is 16 bytes long. If you enable the **chap-challenge-in -request-authenticator** statement and the random challenge is not 16 bytes long, **authd** ignores the statement and uses the default behavior, which inserts the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets.

client-accounting-algorithm—(EX Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the access method the router uses to access RADIUS accounting servers. The default is the **direct** option.

Values:

- direct—Use the direct method.
- round-robin—Use the round-robin method.

client-authentication-algorithm—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, configure the method that the authenticator uses to access RADIUS authentication servers when there are multiple servers configured. Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The router or switch, acting as the authenticator, waits for a response from the server before sending another request.

When there are multiple RADIUS server connections configured for a client, the authenticator attempts to reach the different servers in the order that they are configured. If there is no response from the first RADIUS server, the authenticator attempts to reach the next RADIUS server. This process repeats until the client is either granted access or there are no more configured servers.

If the **direct** method is configured, the authenticator always treats the first server in the list as the primary server. The authenticator moves on to the second server only if the attempt to reach the first

server fails. If the **round-robin** method is configured, the server chosen first will be rotated based on which server was used last. The first server in the list is treated as a primary for the first authentication request, but for the second request, the second server configured is treated as primary, and so on. With this method, all of the configured servers receive roughly the same number of requests on average so that no single server has to handle all of the requests.

NOTE: The **round-robin** access method is not recommended for use with EX Series switches.

Default: The default is the **direct** option.

Values:

- **direct**—Use the direct access method. The authenticator contacts the first RADIUS server on the list for each request, the second server if the first one fails, and so on.
- **round-robin**—Use the round-robin method. The authenticator contacts the first RADIUS server for the first request, the second server for the second request, and so on.

coa-dynamic-variable-validation—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.2X50-D10 for EX Series switches, specify that when a CoA operation includes a change to a client profile dynamic variable that cannot be applied (such as an update to a non-existent filter), the router does not apply any changes to client profile dynamic variables in the request, and responds with a NACK message.

Default: If you do not configure this statement, the router does not apply any incorrect variable updates, but does make any other changes to the client profile dynamic variables, and responds with an ACK message.

ethernet-port-type-virtual—(EX Series, M Series, MX Series only) Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of **ethernet** in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of **virtual**.

NOTE: This statement takes precedence over the **nas-port-type** statement if you include both statements in the same access profile.

access-loop-id-local—Specify that the Agent-Remote-Id and Agent-Circuit-Id are generated locally when these values are not present in the client database.

ip-address-change-notify—(MX Series only) Starting in Junos OS Release 13.1, for on-demand address allocation for dual-stack PPP subscribers, specify that the BNG includes the IPv4-Release-Control VSA (26-164) in the Access-Request that is sent during on-demand IP address allocation and in the Interim-Accounting messages that are sent to report an address change. The configuration of this statement has no effect when on-demand IP address allocation or deallocation is not configured.

Optionally, configure a message that is included in the VSA when it is sent to the RADIUS server.

Default: This functionality is disabled by default.

Values: *message*—VSA message.

Range: Up to 32 characters.

juniper-access-line-attributes—Configure AAA to add Juniper Networks access line VSAs to the RADIUS authentication and accounting request messages for subscribers. If the router has not received and processed the corresponding ANCP attributes from the access node, then AAA provides only the following in these RADIUS messages:

- Downstream-Calculated-QoS-Rate (IANA 4874, 26-141)—Default configured advisory transmit speed.
- Upstream-Calculated-QoS-Rate (IANA 4874, 26-142)—Default configured advisory receive speed.

NOTE: Starting in Junos OS Release 19.2R1, the **juniper-access-line-attributes** option replaces the **juniper-dsl-attributes** option. The difference between these options is that **juniper-dsl-attributes** supported only DSL TLVs received in the ANCP Port Status message. The **juniper-access-line-attributes** option supports PON TLVs in addition to DSL TLVs, and will be extensible to future access technologies.

For backward compatibility with existing scripts, the **juniper-dsl-attributes** option redirects to the new **juniper-access-line-attributes** option. We recommend that you use **juniper-access-line-attributes**.

NOTE: The **juniper-access-line-attributes** option is not backward compatible with Junos OS Release 19.1 or earlier releases. This means that if you have configured **juniper-access-line-attributes** option in Junos OS Release 19.2 or higher releases, you must perform the following steps to downgrade to Junos OS Release 19.1 or earlier releases:

1. Delete the **juniper-access-line-attributes** option from all access profiles that include it.
2. Perform the software downgrade.
3. Add the **juniper-dsl-attributes** option to the affected access profiles.

Default: The Juniper Networks access line VSAs are not added to the RADIUS authentication and accounting request messages. However, the DSL Forum VSA—if available—is added to RADIUS messages by default.

nas-identifier—(EX Series, MX Series, SRX Series only) Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests. This statement was introduced in Junos OS Release 15.1X49-D110 for SRX300, SRX320, SRX340, SRX345, and SRX550M Series devices.

Values: *identifier-value*—String to use for authentication and accounting requests.

Range: 1 through 64 characters.

nas-port-id-delimiter—(MX Series only) Starting in Junos OS Release 11.4, specify the character that the router uses as a separator between the concatenated values in the NAS-Port-ID string. The router uses the delimiter when you configure more than one value in the **nas-port-id-format** statement. The default is the hash (#) character. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Values: *delimiter-character*—Character used for the delimiter.

remote-circuit-id-delimiter—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure a delimiter character for the remote circuit ID string when you use the **remote-circuit-id-format** statement to configure the string to use instead of the Calling-Station ID in L2TP Calling Number AVP 22. If more than one value is configured for the remote circuit ID format, the delimiter character is used as a separator between the concatenated values in the resulting remote circuit ID string. The default is the hash (#) character.

Values: *delimiter*—Delimiter character to be used between components of the remote circuit ID string.

remote-circuit-id-fallback—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the fallback value for the LAC to send in L2TP Calling Number AVP 22, either the configured Calling-Station-ID or the default underlying interface. Use of the fallback value is triggered when the components of the override string you configured with the **remote-circuit-id-format** statement—the ACI, the ARI, or both ACI and ARI—are not received by the LAC in the PPPoE Active Discovery Request (PADR) packet.

Values:

- **configured-calling-station-id**—Send the configured Calling-Station-ID in the Calling Number AVP.
- **default**—Send the underlying interface value in the Calling Number AVP.

remote-circuit-id-format—(MX Series only) Starting in Junos OS Release 13.3R1 on MX Series, configure the format of the string that overrides the Calling-Station-ID format in the Calling Number AVP 22 sent by the LAC to the LNS in the ICRQ packet when an L2TP session is being established. You can specify the ACI, the ARI, or both the ACI and ARI. This statement enables you to decouple the AVP 22 value from the RADIUS Calling-Station-ID attribute (31); the values for AVP 22 and the Calling-Station-ID attribute are the same when you use the **calling-station-id-format** statement to configure AVP 22.

NOTE: You must configure the **override calling-circuit-id remote-circuit-id** statement for the remote circuit ID format to be used in the calling number AVP.

Values:

- **agent-circuit-id**—Specifies use of the ACI string that uniquely identifies the subscriber's access node and the digital subscriber line (DSL) on the access node. For PPPoE traffic, the ACI string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.
- **agent-remote-id**—Specifies use of the ARI string that identifies the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request. The agent remote identifier (ARI) string is stored in the DSL Forum Agent-Remote-ID VSA [26-2] for PPPoE traffic.

service-activation—(MX Series only) Starting in Junos OS Release 16.2, specify whether subscribers are allowed to log in even when service activation failures related to configuration errors occur during family activation request processing by authd for a newly authenticated subscriber. Configuration errors include missing or incorrect syntax, missing or incomplete references to dynamic profiles, and missing or incomplete variables.

NOTE: This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policies.

You can enable separate configurations for subscriber login services for two **service-activation** types: **dynamic-profile** and **extensible-service**. You configure the **dynamic-profile** type services in the dynamic profile at the **[edit dynamic-profiles]** hierarchy level; the profile is used to provide dynamic subscriber access and services for broadband applications. The **extensible-service** type is for business services configured in an operation script and provisioned by the Extensible Subscriber Services Manager daemon (essmd).

Default:

Default behavior depends on the service type:

- For **extensible-service** services: **optional-at-login**.
- For **dynamic-profile** services: **required-at-login**.

Values:

- **optional-at-login**—Service activation is optional. Failure due to configuration errors does not prevent activation of the address family; it allows subscriber access. Failure for any other reason causes network family activation to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.
- **required-at-login**—Service activation is required. Failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.

vlan-nas-port-stacked-format—(MX Series only) Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Access Profile Options for Interactions with RADIUS Servers

RADIUS Servers and Parameters for Subscriber Access

Configuring Authentication and Accounting Parameters for Subscriber Access

Configuring a Calling-Station-ID with Additional Options

order

Syntax

```
order [ accounting-method ];
```

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.

Description

Set the order in which the Junos OS tries different accounting methods for client activity. When a client logs in, the software tries the accounting methods in the specified order.

Options

accounting-method—One or more accounting methods. When a client logs in, the software tries the accounting methods in the following order, from first to last. The only valid value is **radius** for RADIUS accounting.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

pool (Address-Assignment Pools)

Syntax

```
pool pool-name {
  active-drain;
  family family {
    dhcp-attributes {
      [ protocol-specific attributes ]
    }
    excluded-address ip-address;
    excluded-range name low minimum-value high maximum-value;
    host hostname {
      hardware-address mac-address;
      ip-address ip-address;
    }
    network ip-prefixprefix-length>;
    prefix ipv6-prefix;
    range range-name {
      high upper-limit;
      low lower-limit;
      prefix-length prefix-length;
    }
  }
  hold-down;
  link pool-name;
}
```

Hierarchy Level

[edit access [address-assignment](#)]

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure the name of an address-assignment pool.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

pool-name—Name assigned to the address-assignment pool.

active-drain—(MX Series only) Starting in Junos OS Release 17.2, configure the DHCP local server to stop allocating addresses from this pool. When this is configured, the DHCP local server gracefully shifts clients from this address pool to an alternative pool for which active drain is not configured. When existing clients with an address from this pool submit a DHCPv4 request or DHCPv6 renew, they receive a NAK, forcing them to renegotiate. The server responds with a DHCPv4 offer or DHCPv6 advertise message with an address from a different pool.

family —Configure the protocol family for the address-assignment pool.

The options for this statement are explained separately. Click the linked statement for details.

hold-down—(MX Series only) Starting in Junos OS Release 16.1, configure an address-assignment pool that is currently in use to be unavailable for further address allocation. When a pool is in the hold-down state, the pool is no longer used to allocate IP addresses for subscribers. Current subscribers who previously obtained an address from the pool are not affected; they can continue to renew their leases. As each of these users disconnects, their address is not reallocated. The pool becomes inactive when all subscribers have disconnected and their addresses are returned to the pool.

link—(M Series, MX Series, SRX Series, T Series only) Designate a secondary address-assignment pool that is linked to the pool being configured. When the pool being configured has no addresses available for allocation, the secondary pool can be searched for a free address. You can configure a chain of linked pools, but you cannot directly link more than one pool to or from any other pool. Each linked pool in the chain serves as a backup pool for the pool immediately before it in the chain.

Values: *pool-name*—Name assigned to the secondary address-assignment pool.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Address-Assignment Pools Overview](#)

[Address-Assignment Pool Configuration Overview](#)

[Configuring Address-Assignment Pool Linking](#)

[Address Allocation from Linked Address Pools](#)

[Configuring DHCP Local Address Pool Rapid Drain](#)

[Attributes That Can Be Applied to DHCP Clients](#)

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 {
    aaa-options aaa-options-name;
    chap;
    ignore-magic-number-mismatch;
    initiate-ncp (ip | ipv6 | dual-stack-passive)
    ipcp-suggest-dns-option;
    mru;
    mtu;
    pap;
    peer-ip-address-optional;
  }
  primary-dns primary-dns;
  primary-wins primary-wins;
  secondary-dns secondary-dns;
  secondary-wins secondary-wins;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement **idle-timeout** introduced in Junos OS Release 11.1 for the QFX Series.

Statement **idle-timeout** introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description

Configure PPP properties for a group profile.

Options

cell-overhead—(M Series, MX Series, PTX Series, T Series only) Configure the session to use Asynchronous Transfer Mode (ATM)-aware egress shaping on the IQ2 PIC.

encapsulation-overhead—(MX Series, T Series only) Configure the encapsulation overhead for class-of-service calculations.

Values: *bytes*—The number of bytes used as encapsulation overhead for the session.

framed-pool—(M Series, MX Series, PTX Series, T Series only) Configure the address pool.

Values: *framed-pool*—References a configured address pool.

idle-timeout—(EX4600, M Series, MX Series, OCX1100, PTX Series, QFX Series, T Series only) Configure the idle timeout for a user. Starting in Junos OS Release 11.1, this statement is available on the QFX Series. Starting in Junos OS Release 14.1X53-D20, this statement is available on OCX Series switches. The router might consider a PPP session to be idle because of the following reasons:

- There is no ingress traffic on the PPP session.
- There is no egress traffic.
- There is neither ingress or egress traffic on the PPP session.
- There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.

Values: *seconds*—Number of seconds a user can remain idle before the session is terminated.

Range: 0 through 4,294,967,295 seconds

Default: 0

interface-id *interface-id*—(M Series, MX Series, PTX Series, T Series only) Configure the interface identifier.

Values: *interface-id*—Identifier for the interface representing a Layer 2 Tunneling Protocol (L2TP) session configured at the **[edit interfaces *interface-name* unit *local-unit-number* dial-options]** hierarchy level. For more information about the interface ID, see *Services Interface Naming Overview*.

keepalive—(M Series, MX Series, PTX Series, T Series only) Configure the keepalive interval for an L2TP tunnel.

Values: *seconds*—Time period that must elapse before the Junos OS checks the status of the Point-to-Point Protocol (PPP) session by sending an echo request to the peer. 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

primary-dns—(EX Series, SRX Series only) Configure the primary Domain Name System (DNS) server.

Values: *primary-dns*—An IPv4 address.

primary-wins—(M Series, MX Series, PTX Series, T Series only) Configure the primary Windows Internet name server.

Values: *primary-wins*—An IPv4 address.

secondary-dns—(SRX Series only) Configure the secondary DNS server.

Values: *secondary-dns*—An IPv4 address.

secondary-wins—(M Series, MX Series, PTX Series, SRX Series, T Series only) Configure the secondary Windows Internet name server.

Values: *secondary-wins*—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

ppp (Profile)

Syntax

```
ppp {
  cell-overhead;
  encapsulation-overhead bytes;
  framed-ip-address address;
  framed-pool framed-pool;
  idle-timeout seconds;
  interface-id interface-id;
  keepalive seconds;
  primary-dns primary-dns;
  primary-wins primary-wins;
  secondary-dns secondary-dns;
  secondary-wins secondary-wins;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement **idle-timeout** introduced in Junos OS Release 11.1 for the QFX Series.

Statement **idle-timeout** introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description

Configure PPP properties for a client profile.

Options

framed-ip-address—(SRX Series, T Series only) Specify a framed IP address.

Values: *address*—The IP version 4 (IPv4) prefix.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

primary-dns

Syntax

```
primary-dns primary-dns;
```

Hierarchy Level

```
[edit access group-profile profile-name client client-name ppp],  
[edit access profile profile-name ppp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the primary Domain Name System (DNS) server.

Options

primary-dns—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

primary-wins

Syntax

```
primary-wins primary-wins;
```

Hierarchy Level

```
[edit access group-profile profile-name client client-name ppp],  
[edit access profile profile-name ppp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the primary Windows Internet name server.

Options

primary-wins—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

profile (Access)

Syntax

```

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

```

```

    multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
    }
    override-result-code session-out-of-resource;
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    service-profile profile-name(parameter)&profile-name;
    sessions-limit-group limit-group-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;
}
domain-name-server;
domain-name-server-inet;
domain-name-server-inet6;
local {
    flat-file-profile profile-name;
}
preauthentication-order preauthentication-method;
provisioning-order (gx-plus | jsr | pcrf);

```

```

radius {
  accounting-server [ ip-address ];
  attributes {
    exclude {
      attribute-name packet-type;
      standard-attribute number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
      vendor-id id-number {
        vendor-attribute vsa-number {
          packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
        }
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system:routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
authentication-server [ ip-address ];
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
}

```

```

client-authentication-algorithm (direct | round-robin);
coa-dynamic-variable-validation;
ethernet-port-type-virtual;
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    pw-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
        adapter-width width;
        port-width width;
        slot-width width;
        vci-width width;
        vpi-width width;
    }
}
nas-port-id-delimiter delimiter-character;
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback {
    remote-circuit-id-format;
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    max-outstanding-requests value;
    source-address source-address;
    timeout seconds;
}

```

```

service {
  accounting {
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order (activation-protocol | local | radius);
}
session-limit-per-username number;
session-options {
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeoutminutes;
  pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
  }
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
subscriber username {
  delegated-pool delegated-pool-name;
  framed-ip-address ipv4-address;
  framed-ipv6-pool ipv6-pool-name;
  framed-pool ipv4-pool-name;
  password password;
  target-logical-system logical-system-name <target-routing-instance (default | routing-instance-name)>;
  target-routing-instance (default | routing-instance-name);
}
}

```

Hierarchy Level

[edit access]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a subscriber access profile that includes subscriber access, L2TP, or PPP properties.

Options

profile-name—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Authentication Protocol | 5](#)

[Configuring Access Profiles for L2TP or PPP Parameters | 28](#)

[Configuring L2TP Properties for a Client-Specific Profile | 39](#)

Configuring an L2TP Access Profile on the LNS

Configuring an L2TP LNS with Inline Service Interfaces

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

Configuring Service Accounting with JSRC

Configuring Service Accounting in Local Flat Files

AAA Service Framework Overview

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

radius (Access Profile)

Syntax

```
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude
      attribute-name packet-type;
    standard-attribute number {
      packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
    vendor-id id-number {
      vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system-routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
authentication-server [ ip-address ];
options {
  accounting-session-id-format (decimal | description);
  calling-station-id-delimiter delimiter-character;
  calling-station-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    nas-identifier;
  }
  chap-challenge-in-request-authenticator;
  client-accounting-algorithm (direct | round-robin);
  client-authentication-algorithm (direct | round-robin);
```



```

coa-dynamic-variable-validation;
ethernet-port-type-virtual;
interface-description-format {
    exclude-adapter;
    exclude-channel;
    exclude-sub-interface;
}
ip-address-change-notify message;
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
    atm {
        adapter-width width;
        port-width width;
        slot-width width;
        vci-width width;
        vpi-width width;
    }
}
nas-port-id-delimiter delimiter-character;
nas-port-id-format {
    agent-circuit-id;
    agent-remote-id;
    interface-description;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback;
remote-circuit-id-format {
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}

```

Hierarchy Level

[edit access [profile](#) *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

Options

accounting-server—(MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

Values: *ip-address*—IP version 4 (IPv4) address.

authentication-server—(SRX Series only) Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

Values: *ip-address*—IPv4 address.

preauthentication-server—(MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

NOTE: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the **exclude** statement.

Values: *ip-address*—IPv4 address.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Authentication and Accounting Parameters for Subscriber Access](#)

[RADIUS Logical Line Identifier \(LLID\) Overview](#)

[Configuring Logical Line Identification \(LLID\) Preauthentication](#)

radius-disconnect

Syntax

```
radius-disconnect {  
  client-address {  
    secret (RADIUS) password;  
  }  
}
```

Hierarchy Level

[edit access]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a disconnect server that listens on a configured User Datagram Protocol (UDP) port for disconnect messages from a configured client and processes these disconnect messages.

Options

client-address—A valid IP address configured on one of the router interfaces.

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

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

[Example: Configuring CHAP Authentication with RADIUS | 14](#)

[Configuring RADIUS Authentication for L2TP | 71](#)

[Configuring the RADIUS Disconnect Server for L2TP | 80](#)

radius-disconnect-port

Syntax

```
radius-disconnect-port port-number;
```

Hierarchy Level

```
[edit access]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify a port number on which to contact the RADIUS disconnect server. Most RADIUS servers use port number 1700.

Options

port-number—The server port to which disconnect requests from the RADIUS client are sent. The L2TP network server, which accepts these disconnect requests, is the server.

NOTE: The Junos OS accepts disconnect requests only from the client address configured at the `[edit access radius-disconnect client client-address]` hierarchy level.

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 RADIUS Disconnect Server for L2TP](#) | 80

radius-server

Syntax

```
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    max-outstanding-requests value;
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

max-outstanding-requests introduced in Junos OS Release 11.4.

accounting-retry and **accounting-timeout** introduced in Junos OS Release 14.1.

dynamic-request-port option added in Junos OS Release 14.2R1 for MX Series routers.

preauthentication-port and **preauthentication-secret** options added in Junos OS Release 15.1 for MX Series routers.

accounting-port introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

Support for IPv6 **server-address** introduced in Junos OS Release 16.1.

Description

Configure RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—IPv4 or IPv6 address of the RADIUS server.

accounting-port—(EX Series, M Series, MX Series, PTX Series, T Series only) Configure the port number on which to contact the RADIUS accounting server. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

NOTE: Specifying the accounting port is optional, and port 1813 is the default. However, we recommend that you configure it in order to avoid confusion, as some RADIUS servers might refer to an older default.

Values: *port-number*—Port number on which to contact the RADIUS accounting server. Most RADIUS servers use port 1813, as specified in RFC 2866.

Default: 1813

accounting-retry—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure the number of times the device retransmits RADIUS accounting messages when no response is received from the server. When you do not configure this statement, the number of retry attempts is determined by the **retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *number*—Number of retry attempts.

Range: 0 through 100

Default: 0 (disabled)

accounting-timeout—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure how long the local device waits to receive a response from a RADIUS accounting server before retransmitting the message. When you do not configure this statement, the length of the timeout is determined by the **timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Duration of timeout period.

Range: 0 through 1000 seconds

Default: 0 (disabled)

dynamic-request-port—(MX Series only) Starting in Junos OS Release 14.2R1, specify the port that the router monitors for dynamic (CoA) requests from the specified RADIUS servers. You can configure a port globally or for a specific access profile.

You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

NOTE: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

Values: *port-number*—Number of the monitored port.

Default: 3799 (as specified in RFC 5176)

max-outstanding-requests—(MX Series only) Starting in Junos OS Release 11.4, configure the maximum number of outstanding requests for this RADIUS server. An increase in this value is immediate while a decrease is more gradual if the current number of outstanding requests exceeds the new value.

Values: *requests*—Maximum number of outstanding requests for this RADIUS server.

Range: 0 through 2000 outstanding requests per server

Default: 1000 outstanding requests per server

port—(EX Series, M Series, MX Series, SRX Series, T Series only) Configure the port number on which to contact the RADIUS server.

Values: *port-number*—Port number on which to contact the RADIUS server.

Default: 1812 (as specified in RFC 2865)

preauthentication-port—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the port number on which to contact the RADIUS server for logical line identification (LLID) preauthentication requests. If you do not configure a separate UDP port for preauthentication purposes, the same UDP port that you configure for authentication messages by including the **port port-number** statement is used.

Values: *port-number*—Port number used for preauthentication requests to contact the RADIUS server.

preauthentication-secret—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the password to use with the RADIUS server for LLID preauthentication requests. If you do not configure a separate UDP password for preauthentication purposes, the same password that you configure for authentication messages by including the **secret password** statement is used. The secret password used by the local router must match that used by the server.

Values: *password*—Password to use. To include spaces enclose the character string in quotation marks.

retry—(EX Series, M Series, MX Series, PTX Series, T Series only) Specify the number of times that the device is allowed to attempt to contact a RADIUS authentication or accounting server. You can override the retry limit for accounting servers with the **accounting-retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *attempts*—Number of times that the router is allowed to attempt to contact a RADIUS server.

Range: 1 through 100

Default: 3

routing-instance—(SRX Series, vSRX only) Configure the routing instance used to send RADIUS packets to the RADIUS server.

Values: *routing-instance-name*—Routing instance name.

source-address—(SRX Series, vSRX only) Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. Support for IPv6 **source-address** was introduced in Junos OS Release 16.1.

Values: *source-address*—Valid IPv4 or IPv6 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

timeout—(SRX Series, vSRX only) Configure the amount of time that the local device waits to receive a response from RADIUS authentication and accounting servers. You can override the timeout value for accounting servers with the **accounting-timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Amount of time to wait.

Range: 1 through 1000 seconds

Default: 3 seconds

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

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

[Configuring the PPP Authentication Protocol | 5](#)

[Configuring RADIUS Authentication for L2TP | 71](#)

Configuring RADIUS System Accounting

Configuring RADIUS-Initiated Dynamic Request Support

RADIUS Logical Line Identifier (LLID) Overview

RADIUS Attributes for LLID Preauthentication Requests

show network-access aaa statistics

[clear network-access aaa statistics | 242](#)

range (Address-Assignment Pools)

Syntax

```
range range-name {
    high upper-limit;
    low lower-limit;
    prefix-length prefix-length;
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family (inet | inet6)]
```

Release Information

Statement introduced in Junos OS Release 9.0.

IPv6 support introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Configure a named range of IPv4 addresses or IPv6 prefixes, used within an address-assignment pool.

Options

high upper-limit—Upper limit of an address range or IPv6 prefix range.

low lower-limit—Lower limit of an address range or IPv6 prefix range.

prefix-length prefix-length—Assigned length of the IPv6 prefix.

range-name—Name assigned to the range of IPv4 addresses or IPv6 prefixes.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a Named Address Range for Dynamic Address Assignment

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

revert-interval (Access)

Syntax

```
revert-interval interval;
```

Hierarchy Level

```
[edit access profile profile-name radius options],  
[edit access radius-options]
```

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure the amount of time the router or switch waits after a server has become unreachable. The router or switch rechecks the connection to the server when the specified interval expires. If the server is then reachable, it is used in accordance with the order of the server list.

Options

interval—Amount of time to wait.

Range: 0 through 604,800 seconds

Default: 60 seconds

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Configuring Authentication and Accounting Parameters for Subscriber Access

secondary-dns

Syntax

```
secondary-dns secondary-dns;
```

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 secondary DNS server.

Options

secondary-dns—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

secondary-wins

Syntax

```
secondary-wins secondary-wins;
```

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 secondary Windows Internet name server.

Options

secondary-wins—An IPv4 address.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the PPP Attributes for a Group Profile | 25](#)

[Configuring PPP Properties for a Client-Specific Profile | 44](#)

secret (RADIUS)

Syntax

```
secret password;
```

Hierarchy Level

```
[edit access profile profile-name radius-server server-address],  
[edit access radius-disconnect client-address],  
[edit access radius-server server-address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.

Options

password—Password to use; it can include spaces if the character string is enclosed in quotation marks.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

RADIUS Authentication and Accounting Server Definition

[Example: Configuring CHAP Authentication with RADIUS | 14](#)

[Configuring RADIUS Authentication for L2TP | 71](#)

[Configuring the RADIUS Disconnect Server for L2TP | 80](#)

session-options

Syntax

```
session-options {
  client-group [ group-names ];
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeout minutes;
  pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
  }
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
```

Hierarchy Level

[edit access [profile](#) *profile-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

(MX Series and SRX Series devices) Define options to place limits on subscriber access based on how long the session has been up, how long the user has been inactive, or both.

(MX Series) Define options to modify a subscriber username at login based on the subscriber's access profile.

(MX Series) Specify characteristics related to policy and charging control (PCC) rules, such as the PCEF profile that contains the rules, service sets to process the rules, and service filters for the service sets.

Options

client-idle-timeout—(MX Series only) Specify the grace period that begins after an authenticated user terminates all sessions and connections. Authentication is not required if a new connection is initiated during the grace period by the same user.

During this period, the router determines whether the subscriber is inactive by monitoring data traffic, both upstream from the user (ingress) and downstream to the user (egress). Control traffic is ignored. The subscriber is not considered idle as long as data traffic is detected in either direction. When no traffic is detected for the duration of the idle time out, non-DHCP subscribers (such as L2TP or PPP) are gracefully logged out, similarly to a RADIUS-initiated disconnect or a CLI-initiated logout; DHCP subscribers are disconnected.

When you additionally configure the related **client-idle-timeout-ingress-only** statement (MX Series only), the router monitors only ingress traffic to determine whether the subscriber is inactive; it does not monitor any egress traffic. The related **client-session-timeout** statement terminates the subscriber session when the session timeout expires regardless of user activity.

Client idle timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model. It is not practical for DHCP or DHCPv6 subscribers.

Although you can use the **client-idle-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the idle timeout for VLANs, the timeout period starts when the VLAN is instantiated. It resets when a client session is created or an existing session is reactivated. When no traffic is detected on an authenticated VLAN for the duration of the timeout, the VLAN is considered inactive and is deleted. If no client sessions are ever created on the VLAN, then the VLAN is removed when the timeout expires.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes of idle time that elapse before the session is terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 10 through 1440 minutes

client-idle-timeout-ingress-only—(MX Series only) Starting in Junos OS Release 16.2, specify that only ingress traffic is monitored for subscriber idle timeout processing for the duration of the idle timeout period that you specify with the **client-idle-timeout** statement. If no ingress traffic is received for the duration of the timeout, then the subscriber is gracefully logged out (non-DHCP subscribers) or disconnected (DHCP subscribers).

If you configure **client-idle-timeout** alone, then both ingress and egress traffic are monitored during the idle timeout. Monitoring only ingress traffic is useful in cases where the LNS sends traffic to the remote peer even when the peer is not up, such as when the LNS does not have PPP keepalives enabled and therefore does not detect that the peer is not up. Because the LAC monitors both ingress and egress traffic by default, in this situation it receives the egress traffic from the LNS and either

does not log out the subscriber or delays detection of inactivity until the egress traffic ceases. When you specify that only ingress traffic is monitored in this case, the LAC can detect that the peer is inactive and then initiate logout.

client-session-timeout—(SRX Series, vSRX only) Specify the amount of time after which user sessions are terminated, regardless of user activity (also known as a forced or hard authentication timeout).

Alternatively, when you want subscribers to be identified as inactive before they are terminated, use the related statements, **client-idle-timeout** and **client-idle-timeout-ingress-only**. Use **client-idle-timeout** alone to specify a period of time during which both ingress and egress subscriber data traffic is monitored; if no traffic is detected for the duration of the period, the subscriber is considered inactive and is terminated. Add the **client-idle-timeout-ingress-only** statement to monitor only ingress traffic for the duration of the timeout set with the **client-idle-timeout** statement.

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

Client session timeouts are most often used for residential services rather than business services. The most practical use case for this timeout is in a PPP access model when no voice services are offered. For DHCP or DHCPv6 subscribers, the session timeout is used as the DHCP lease timer if no other lease time configuration is present.

Although you can use the **client-session-timeout** statement for dynamically configured subscriber VLANs, this configuration is useful only in limited circumstances (such as IP over Ethernet without DHCP and with fixed addresses) and is not typically used. If you do use the session timeout for VLANs, the timeout period starts when the VLAN is instantiated.

Default: The timeout is not configured.

Values: *minutes*—Number of minutes after which user sessions are terminated. The value that you specify must be determined locally with consideration of the services and policies that you offer.

Range: 1 through 527040 minutes

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

Required Privilege Level

access—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding Session Options for Subscriber Access

Configuring Subscriber Session Timeout Options

Configuring Username Modification for Subscriber Sessions

Removing Inactive Dynamic Subscriber VLANs

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

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.

volume-time option added 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.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

update-interval

Syntax

```
update-interval minutes;
```

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.

Description

Enable interim accounting updates and configure the amount of time that the router or switch waits before sending a new accounting update.

Interim accounting updates are included in the exchange of messages between the client and the accounting server. In RADIUS accounting, the client is the network access server (NAS), which can be the router or switch. The NAS sends Accounting-Request messages to the server, which acknowledges receipt of the requests with Accounting-Response messages. Interim accounting updates are sent in Accounting-Request packets with the Acct-Status-Type attribute set to Interim-Update.

When a user is authenticated, the authentication server issues an Access-Accept message in response to a successful Access-Request message. The interval between interim updates can be configured directly on the server using the Acct-Interim-Interval attribute of the Access-Accept message. However, if the update interval is configured on the NAS using **update-interval**, then the locally configured value overrides the value found in an Access-Accept message from the server.

NOTE: All information in an interim update message is cumulative from the beginning of the session, not from the last interim update message.

Default

No interim updates are sent from the client to the accounting server.

Options

minutes—Amount of time between updates, in minutes. All values are rounded to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

Range: 10 through 1440 minutes

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications

3

PART

Administration

Administrative Commands | **241**

Monitoring Commands | **259**

Administrative Commands

IN THIS CHAPTER

- [clear network-access aaa statistics | 242](#)
- [clear network-access aaa subscriber | 247](#)
- [clear services l2tp session | 249](#)
- [clear services l2tp tunnel statistics | 252](#)
- [show services l2tp radius | 254](#)

clear network-access aaa statistics

Syntax

```
clear network-access aaa statistics
<accounting>
<address-assignment (client | pool pool-name)>
<authentication>
<dynamic-requests>
<radius>
<re-authentication>
<session-limit-per-username username username access-profile profile-name>
<terminate-code>
```

Release Information

Command introduced in Junos OS Release 10.0.

radius option introduced in Junos OS Release 11.4

terminate-code option introduced in Junos OS Release 11.4.

session-limit-per-username option introduced in Junos OS Release 18.4R1 on MX Series routers.

Description

Clear AAA statistics.

Options

accounting—(Optional) Clear AAA accounting statistics.

address-assignment client—(Optional) Clear AAA address-assignment statistics for the client.

address-assignment pool *pool-name*—(Optional) Clear AAA address-assignment pool statistics.

authentication—(Optional) Clear AAA authentication statistics.

dynamic-requests—(Optional) Clear AAA dynamic-request statistics.

radius—(Optional) Clears the values in the **Peak** and **Exceeded** columns only.

re-authentication—(Optional) Clear AAA reauthentication statistics.

session-limit-per-username—(MX Series routers only) (Optional) Clear all blocked request statistics for all access profiles from the username session-limit table. You can also specify additional options:

- **username *username***—Clear the blocked request statistics for the specified username across all access profiles. A given username can be used in more than one access profile.
- **access-profile *profile-name***—Clear the blocked request statistics for all usernames in the specified access profile.

NOTE: This command does not clear (delete) the entry in the session-limit table. Entries in the table are added or deleted during session login or logout processing.

terminate-code—(Optional) Clear AAA termination code statistics.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

Understanding Session Options for Subscriber Access

Limiting the Number of Active Sessions per Username and Access Profile

show network-access aaa statistics

List of Sample Output

[clear network-access aaa statistics accounting on page 243](#)

[clear network-access aaa statistics address-assignment pool on page 244](#)

[clear network-access aaa statistics radius on page 244](#)

[clear network-access aaa statistics session-limit-per-username \(All Usernames Across All Access Profiles\) on page 244](#)

[clear network-access aaa statistics session-limit-per-username \(Specific Username Across All Access Profiles\) on page 244](#)

[clear network-access aaa statistics session-limit-per-username \(All Usernames for Specific Access Profiles\) on page 245](#)

[clear network-access aaa statistics session-limit-per-username \(Specific Username in Specific Access Profile\) on page 245](#)

Output Fields

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

Sample Output

clear network-access aaa statistics accounting

```
user@host> clear network-access aaa statistics accounting
```

clear network-access aaa statistics address-assignment pool

```
user@host> clear network-access aaa statistics address-assignment pool isp_1
```

clear network-access aaa statistics radius

```
user@host> clear network-access aaa statistics radius
```

clear network-access aaa statistics session-limit-per-username (All Usernames Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail
```

Username	Access-profile	Blocked requests	Session count
rkx@example.net	BNG1	1	4
xyz@example.net	BNG1	3	5
abc@example.net	BNG2	2	5
pqr@example.net	BNG2	3	4

```
user@host> clear network-access aaa statistics session-limit-per-username
```

```
user@host> show network-access aaa statistics on-limit-per-username detail
```

Username	Access-profile	Blocked requests	Session count
rkx@example.net	BNG1	0	4
xyz@example.net	BNG1	0	5
abc@example.net	BNG2	0	5
pqr@example.net	BNG2	0	4

clear network-access aaa statistics session-limit-per-username (Specific Username Across All Access Profiles)

```
user@host> show network-access aaa statistics on-limit-per-username detail
```

Username	Access-profile	Blocked requests	Session count
rkx@example.net	BNG1	1	4
xyz@example.net	BNG1	3	5
rkx@example.net	BNG2	2	5
pqr@example.net	BNG2	3	4

```
user@host> clear network-access aaa statistics session-limit-per-username username rkx@example.net
```

```
user@host> show network-access aaa statistics on-limit-per-username detail
```

Username	Access-profile	Blocked requests	Session count
rkv@example.net	BNG1	0	4
xyz@example.net	BNG1	3	5
rkv@example.net	BNG2	0	5
pqr@example.net	BNG2	3	4

clear network-access aaa statistics session-limit-per-username (All Usernames for Specific Access Profiles)

user@host> **show network-access aaa statistics on-limit-per-username detail**

Username	Access-profile	Blocked requests	Session count
rkv@example.net	BNG1	1	4
xyz@example.net	BNG1	3	5
rkv@example.net	BNG2	2	5
pqr@example.net	BNG2	3	4

user@host> **clear network-access aaa statistics session-limit-per-username access-profile BNG2**

user@host> **show network-access aaa statistics on-limit-per-username detail**

Username	Access-profile	Blocked requests	Session count
rkv@example.net	BNG1	1	4
xyz@example.net	BNG1	3	5
123@example.net	BNG2	0	5
pqr@example.net	BNG2	0	4

clear network-access aaa statistics session-limit-per-username (Specific Username in Specific Access Profile)

user@host> **show network-access aaa statistics on-limit-per-username detail**

Username	Access-profile	Blocked requests	Session count
rkv@example.net	BNG1	1	4
xyz@example.net	BNG1	3	5
rkv@example.net	BNG2	2	5
pqr@example.net	BNG2	3	4

user@host> **clear network-access aaa statistics session-limit-per-username username rk@example.net**
access-profile BNG2

user@host> **show network-access aaa statistics on-limit-per-username detail**

Username	Access-profile	Blocked requests	Session count
rk@example.net	BNG1	3	4
xyz@example.net	BNG1	3	5
rk@example.net	BNG2	0	5
pqr@example.net	BNG2	3	4

clear network-access aaa subscriber

Syntax

```
clear network-access aaa subscriber
<session-id identifier <reconnect>>
<statistics username username>
<username username <reconnect>>
```

Release Information

Command introduced in Junos OS Release 9.1.

reconnect and **session-id** options added in Junos OS Release 16.1R4.

Description

Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.

Options

reconnect—(Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16). You can apply this option to either a Layer 2 wholesale session or a conventionally auto-sensed dynamic VLAN supporting a PPPoE session.

In the latter case, this option triggers a PPPoE session logout and removal of the dynamic VLAN logical interface. This is followed by authorization of the access-line to attempt creation of a dynamic VLAN IFL supporting Layer 2 wholesale session in its place.

session-id identifier—(Optional) Log out the subscriber based on the subscriber session identifier.

statistics username username—(Optional) Clear AAA subscriber statistics and log out the subscriber.

username username—(Optional) Log out the AAA subscriber.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

List of Sample Output

[clear network-access aaa subscriber statistics username on page 248](#)

[clear network-access aaa subscriber username on page 248](#)

[clear network-access aaa subscriber session-id on page 248](#)

Output Fields

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

Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

clear 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.

NOTE: On MX Series routers, you cannot issue the **clear services l2tp session** command in parallel with statistics-related **show services l2tp** commands from separate terminals. If this **clear** command is running, then you must press Ctrl+c to make the command run in the background before issuing any of the **show** commands listed in the following table:

show services l2tp destination extensive	show services l2tp summary statistics
show services l2tp destination statistics	show services l2tp tunnel extensive
show services l2tp session extensive	show services l2tp tunnel statistics
show services l2tp session statistics	

Options

all—Close all L2TP sessions.

BEST PRACTICE: The **all** option is not intended to be used as a means to perform a bulk logout of L2TP subscribers. We recommend that you do not use the **all** option in a production environment. Instead of clearing all subscribers at once, consider clearing subscribers in smaller group, based on interface, tunnel, or destination end point.

interface *interface-name*—Clear only the L2TP sessions using the specified adaptive services or inline services interface. The interface type depends on the line card as follows:

- **si-fpc/pic/port**—MPCs on MX Series routers only. 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](#)

[show services l2tp session](#) | [260](#)

List of Sample Output

[clear services l2tp session on page 251](#)

[clear services l2tp session interface on page 251](#)

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 ID	Remote ID	State	Interface unit	Interface 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 ID	Remote ID	State	Interface unit	Interface Name
34915	2	Established	1073741829	si-2/1/0
46142	4	Established	1073741831	si-2/1/0

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

show services l2tp tunnel

List of Sample Output

[clear services l2tp tunnel statistics all on page 253](#)

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
```

show services l2tp radius

Syntax

```
show services l2tp radius
<accounting (servers | statistics)>
<authentication (servers | statistics)>
<servers>
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0.

Description

(M7i, M10i, and M120 routers only) Display RADIUS servers and statistics information for the RADIUS servers configured on the router.

Options

You must include one of the following keywords to provide a valid completion for the command:

accounting (servers | statistics)—(Optional) Display RADIUS servers or statistical accounting information only.

authentication (servers | statistics)—(Optional) Display RADIUS servers or statistical authentication information only.

servers—(Optional) Display RADIUS authentication and accounting server information only.

statistics—(Optional) Display RADIUS authentication and accounting statistics information only.

Required Privilege Level

view

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

List of Sample Output

[show services l2tp radius servers on page 256](#)

[show services l2tp radius statistics on page 257](#)

Output Fields

Table 8 on page 255 lists the output fields for the **show services l2tp radius** command. Output fields are listed in the approximate order in which they appear.

Table 8: show services l2tp radius Output Fields

Field Name	Field Description
IP Address	IP address of the server.
State	(servers keyword only) Present state of the server.
UDP Port	Number of the UDP port used to send authentication or accounting messages to the server.
Retry Count	(servers keyword only) Number of times the RADIUS client resends a packet if no ACK is received.
Timeout	(servers keyword only) Length of time the client waits for an ACK before retransmission.
Pending Requests	(servers keyword only) Number of client pending authentication or accounting requests.
Maximum Sessions	(servers keyword only) Maximum number of pending requests on each RADIUS client before the server moves to the next RADIUS client, which is 200 times the maximum number of clients that can be created on a server (which is 12).
Dead Time	(servers keyword only) Interval to wait before retrying a server after it fails to send a response to an authentication or accounting request.
Secret Type	(servers keyword only) Secret type configured on the RADIUS server.
Profile	(servers keyword only) Name of profile configured for the RADIUS server.
Access requests	(statistics keyword only) Number of access requests sent to the server.
Rollover requests	(statistics keyword only) Number of requests coming into the server as a result of the previous server timing out.
Retransmissions	(statistics keyword only) Number of retransmissions.
Access accepts	(statistics keyword only) Number of access accept messages received from the server.
Access rejects	(statistics keyword only) Number of access reject messages received from the server.
Access challenges	(statistics keyword only) Number of access challenges received from the server.

Table 8: show services l2tp radius Output Fields (*continued*)

Field Name	Field Description
Malformed responses	(statistics keyword only) Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).
Bad authenticators	(statistics keyword only) Number of responses in which the authenticator is incorrect for the matching request. This can occur if the RADIUS secrets for the client and server do not match.
Requests pending	(statistics keyword only) Number of requests waiting for a response.
Request timeouts	(statistics keyword only) Number of requests that timed out.
Unknown responses	(statistics keyword only) Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.
Packets dropped	(statistics keyword only) Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request. For example, if the router sends a request that times out, the router removes the request from the list and sends a new request. If the server is slow and sends a response to the first request after the router removes the request, the packet is dropped.

Sample Output

show services l2tp radius servers

user@host> show services l2tp radius servers

RADIUS Authentication Servers								
IP Address	State	UDP Port	Retry Count	Timeout	Pending Requests	Maximum Sessions	Dead Time	Secret Type
192.0.2.1	Active	1812	2	25	0	2400	300	radius-key
198.51.100.1	Active	1812	5	35	0	2400	300	radius-key
203.0.113.1	Active	1812	2	25	0	2400	300	radius-key
172.28.30.174	Active	1812	7	75	0	2400	300	radius-key
172.28.30.175	Active	1812	7	75	0	2400	300	radius-key
172.28.30.176	Active	1812	4	55	0	2400	300	radius-key
172.31.30.176	Active	1812	3	3	0	2400	300	none-set

```
172.31.130.174 Active 1812 7 75 0 2400 300 radius-key
```

RADIUS Accounting Servers

IP Address	State	UDP Port	Retry Count	Timeout	Pending Requests	Maximum Sessions	Dead Time	Secret Type
192.0.2.1	Active	1813	2	25	0	2400	300	radius-key
198.51.100.1	Active	1813	5	35	0	2400	300	radius-key
203.0.113.1	Active	1813	2	25	0	2400	300	radius-key
172.28.30.174	Active	1813	7	75	0	2400	300	radius-key
172.28.30.175	Active	1813	7	75	0	2400	300	radius-key
172.28.30.176	Active	1813	4	55	0	2400	300	radius-key
172.31.30.176	Active	1813	3	3	0	2400	300	none-set
172.31.130.174	Active	1813	7	75	0	2400	300	radius-key

RADIUS Accounting Servers

```
Profile: user1
```

show services l2tp radius statistics

```
user@host> show services l2tp radius statistics
```

RADIUS Authentication Statistics

Authentication statistics:

```
Server 192.0.2.1, UDP port: 1812
```

```
Access requests      : 40
Rollover requests    : 5
Retransmissions      : 2
Access accepts       : 39
Access rejects       : 1
Access challenges     : 3
Malformed responses   : 0
Bad authenticators    : 0
Requests pending     : 1
Request timeouts     : 0
Unknown responses     : 0
Packets dropped      : 0
```

RADIUS Accounting Statistics

Accounting statistics:

Server 172.31.130.174, UDP port: 1813

Total requests	: 9
Start requests	: 6
Interim requests	: 1
Stop requests	: 2
Rollover requests	: 0
Retransmissions	: 1
Total response	: 9
Start responses	: 6
Interim responses	: 1
Stop responses	: 2
Malformed responses	: 0
Bad authenticators	: 0
Requests pending	: 1
Request timeouts	: 0
Unknown responses	: 0
Packets dropped	: 0

Monitoring Commands

IN THIS CHAPTER

- [show services l2tp session | 260](#)
- [show services l2tp radius | 271](#)
- [show services l2tp summary | 276](#)

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.

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](#) | 249

List of Sample Output

[show services l2tp session \(LNS on M Series Routers\) on page 266](#)

[show services l2tp session \(LNS on MX Series Routers\) on page 266](#)

[show services l2tp session \(LAC\) on page 267](#)

[show services l2tp session detail \(LAC\) on page 267](#)

[show services l2tp session extensive \(LAC\) on page 267](#)

[show services l2tp session extensive \(LAC on MX Series Routers\) on page 268](#)

[show services l2tp session extensive \(LNS on M Series Routers\) on page 268](#)

[show services l2tp session extensive \(LNS on MX Series Routers\) on page 269](#)

[show services l2tp session statistics \(MX Series Routers\) on page 270](#)

Output Fields

Table 9 on page 262 lists the output fields for the **show services l2tp session** command. Output fields are listed in the approximate order in which they appear.

Table 9: 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
Session remote ID	Identifier of the remote endpoint of the L2TP session, as assigned by the L2TP access concentrator (LAC).	All levels
State	<p>State of the L2TP session:</p> <ul style="list-style-type: none"> • 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. • Ins-ic-accept-new—New session is being accepted. • Ins-ic-idle—Session has been created and is idle. • Ins-ic-reject-new—New session is being rejected. • Ins-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	<p>(LNS) Mode of the interface representing the session: shared or exclusive.</p> <p>(LAC) Mode of the interface representing the session: shared or dedicated. Only dedicated is currently supported for the LAC.</p>	extensive
Local IP	IP address of local endpoint of the Point-to-Point Protocol (PPP) session.	extensive

Table 9: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
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	<p>Transmit speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived.</p> <p>Starting in Junos OS Release 14.1, either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>For Junos OS Release 17.2 and Release 17.3, only the current (update) line speed can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 17.4R1, once again either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	extensive

Table 9: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Rx speed	<p>Receive speed of the session conveyed from the LAC to the LNS, in bits per second (bps) and the source method from which the speed is derived.</p> <p>Starting in Junos OS Release 14.1, either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers:</p> <ul style="list-style-type: none"> • When connection speed updates are not enabled, then only the initial line speed is displayed. • When connection speed updates are enabled, then both the initial and the current speeds are displayed. <p>For Junos OS Release 17.2 and Release 17.3, only the current (update) line speed can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 17.4R1, once again either the initial (initial) line speed or both the initial and current (update) line speeds can be displayed on MX Series routers.</p> <p>Starting in Junos OS Release 15.1, when the Tx connect speed method is set to none, the value of zero (0) is displayed.</p>	extensive
Bearer type	<p>Type of bearer enabled:</p> <ul style="list-style-type: none"> • 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	<p>Type of framing enabled:</p> <ul style="list-style-type: none"> • 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

Table 9: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Idle time	Length of time elapsed since the call became idle, in hours, minutes, and seconds.	extensive

Table 9: show services l2tp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics since	<p>Date and time when collection of the following statistics began:</p> <ul style="list-style-type: none"> • 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 timeout—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 ID	Remote ID	State	Interface unit	Interface Name
31694	1	Established	311	pp0

show services l2tp session detail (LAC)

user@host> show services l2tp session detail

```

Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:      1, Interface unit: 311
    State: Established, Interface: pp0, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns

```

show services l2tp session extensive (LAC)

user@host> show services l2tp session extensive

```

Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:      1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.2:1701, Remote IP: 203.0.113.1:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 0, Rx speed: 0
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A

```

show services l2tp session extensive (LAC on MX Series Routers)

user@host> **show services l2tp session extensive**

```
Tunnel local ID: 31889
  Session local ID: 31694, Session remote ID:      1
    Interface unit: 311
    State: Established, Mode: Dedicated
    Local IP: 203.0.113.102:1701, Remote IP: 203.0.113.101:1701
    Local name: ce-lac, Remote name: ce-lns
    Tx speed: 256000, source service-profile
    Rx speed: 128000, source ancp
    Bearer type: 1, Framing type: 1
    LCP renegotiation: N/A, Authentication: None, Interface ID: N/A
    Interface unit: 311, Call serial number: 0
    Policer bandwidth: 0, Policer burst size: 0
    Policer exclude bandwidth: 0, Firewall filter: 0
    Session encapsulation overhead: 0, Session cell overhead: 0
    Create time: Tue Aug 24 14:38:23 2010, Up time: 01:06:25
    Idle time: N/A
```

show services l2tp session extensive (LNS on 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: 203.0.113.121:1701, Remote IP: 203.0.113.202:1701
    Username: user@example.com, Assigned IP address: 203.0.113.51/32
    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: 203.0.113.121:1701, Remote IP: 203.0.113.222:1701
Username: usrl@company.example.com, Assigned IP address: 203.0.113.3/24
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: 192.0.2.2:1701, Remote IP: 192.0.2.3:1701
Local name: lns-mx960, Remote name: testlac
Tx speed: initial 64000, Update 256000
Rx speed: initial 64000, Update 256000
Bearer type: 2, Framing type: 1
LCP renegotiation: Off, Authentication: None
Call serial number: 1
Create time: Mon Apr 25 20:27:50 2011, Up time: 00:01:48

```

```

Idle time: N/A
Statistics since: Mon Apr 25 20:27:50 2011

      Packets      Bytes
Control Tx         4       219
Control Rx         4       221
Data Tx            0         0
Data Rx          10       228
Errors Tx          0
Errors Rx          0

```

show services l2tp session statistics (MX Series Routers)

user@host>**show services l2tp session statistics local session-id 1**

```

Tunnel local ID: 17185
Session local ID: 1, Session remote ID: 14444, Interface unit: 1073788352
State: Established
Statistics since: Mon Aug 1 13:27:47 2011

      Packets  Bytes
Data Tx    4    51
Data Rx    3    36

```

show services l2tp radius

Syntax

```
show services l2tp radius
<accounting (servers | statistics)>
<authentication (servers | statistics)>
<servers>
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0.

Description

(M7i, M10i, and M120 routers only) Display RADIUS servers and statistics information for the RADIUS servers configured on the router.

Options

You must include one of the following keywords to provide a valid completion for the command:

accounting (servers | statistics)—(Optional) Display RADIUS servers or statistical accounting information only.

authentication (servers | statistics)—(Optional) Display RADIUS servers or statistical authentication information only.

servers—(Optional) Display RADIUS authentication and accounting server information only.

statistics—(Optional) Display RADIUS authentication and accounting statistics information only.

Required Privilege Level

view

RELATED DOCUMENTATION

L2TP Services Configuration Overview

L2TP Minimum Configuration

List of Sample Output

[show services l2tp radius servers on page 273](#)

[show services l2tp radius statistics on page 274](#)

Output Fields

Table 8 on page 255 lists the output fields for the **show services l2tp radius** command. Output fields are listed in the approximate order in which they appear.

Table 10: show services l2tp radius Output Fields

Field Name	Field Description
IP Address	IP address of the server.
State	(servers keyword only) Present state of the server.
UDP Port	Number of the UDP port used to send authentication or accounting messages to the server.
Retry Count	(servers keyword only) Number of times the RADIUS client resends a packet if no ACK is received.
Timeout	(servers keyword only) Length of time the client waits for an ACK before retransmission.
Pending Requests	(servers keyword only) Number of client pending authentication or accounting requests.
Maximum Sessions	(servers keyword only) Maximum number of pending requests on each RADIUS client before the server moves to the next RADIUS client, which is 200 times the maximum number of clients that can be created on a server (which is 12).
Dead Time	(servers keyword only) Interval to wait before retrying a server after it fails to send a response to an authentication or accounting request.
Secret Type	(servers keyword only) Secret type configured on the RADIUS server.
Profile	(servers keyword only) Name of profile configured for the RADIUS server.
Access requests	(statistics keyword only) Number of access requests sent to the server.
Rollover requests	(statistics keyword only) Number of requests coming into the server as a result of the previous server timing out.
Retransmissions	(statistics keyword only) Number of retransmissions.
Access accepts	(statistics keyword only) Number of access accept messages received from the server.
Access rejects	(statistics keyword only) Number of access reject messages received from the server.
Access challenges	(statistics keyword only) Number of access challenges received from the server.

Table 10: show services l2tp radius Output Fields (*continued*)

Field Name	Field Description
Malformed responses	(statistics keyword only) Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one).
Bad authenticators	(statistics keyword only) Number of responses in which the authenticator is incorrect for the matching request. This can occur if the RADIUS secrets for the client and server do not match.
Requests pending	(statistics keyword only) Number of requests waiting for a response.
Request timeouts	(statistics keyword only) Number of requests that timed out.
Unknown responses	(statistics keyword only) Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.
Packets dropped	(statistics keyword only) Number of packets dropped because they are too short or because the router receives a response for which there is no corresponding request. For example, if the router sends a request that times out, the router removes the request from the list and sends a new request. If the server is slow and sends a response to the first request after the router removes the request, the packet is dropped.

Sample Output

show services l2tp radius servers

user@host> show services l2tp radius servers

RADIUS Authentication Servers								
IP Address	State	UDP Port	Retry Count	Timeout	Pending Requests	Maximum Sessions	Dead Time	Secret Type
192.0.2.1	Active	1812	2	25	0	2400	300	radius-key
198.51.100.1	Active	1812	5	35	0	2400	300	radius-key
203.0.113.1	Active	1812	2	25	0	2400	300	radius-key
172.28.30.174	Active	1812	7	75	0	2400	300	radius-key
172.28.30.175	Active	1812	7	75	0	2400	300	radius-key
172.28.30.176	Active	1812	4	55	0	2400	300	radius-key
172.31.30.176	Active	1812	3	3	0	2400	300	none-set


```
172.31.130.174 Active 1812 7 75 0 2400 300 radius-key
```

RADIUS Accounting Servers

IP Address	State	UDP Port	Retry Count	Timeout	Pending Requests	Maximum Sessions	Dead Time	Secret Type
192.0.2.1	Active	1813	2	25	0	2400	300	radius-key
198.51.100.1	Active	1813	5	35	0	2400	300	radius-key
203.0.113.1	Active	1813	2	25	0	2400	300	radius-key
172.28.30.174	Active	1813	7	75	0	2400	300	radius-key
172.28.30.175	Active	1813	7	75	0	2400	300	radius-key
172.28.30.176	Active	1813	4	55	0	2400	300	radius-key
172.31.30.176	Active	1813	3	3	0	2400	300	none-set
172.31.130.174	Active	1813	7	75	0	2400	300	radius-key

RADIUS Accounting Servers

```
Profile: user1
```

show services l2tp radius statistics

```
user@host> show services l2tp radius statistics
```

RADIUS Authentication Statistics

Authentication statistics:

```
Server 192.0.2.1, UDP port: 1812
```

```
Access requests      : 40
Rollover requests    : 5
Retransmissions      : 2
Access accepts       : 39
Access rejects       : 1
Access challenges     : 3
Malformed responses   : 0
Bad authenticators    : 0
Requests pending     : 1
Request timeouts     : 0
Unknown responses     : 0
Packets dropped       : 0
```

RADIUS Accounting Statistics

Accounting statistics:

Server 172.31.130.174, UDP port: 1813

Total requests	: 9
Start requests	: 6
Interim requests	: 1
Stop requests	: 2
Rollover requests	: 0
Retransmissions	: 1
Total response	: 9
Start responses	: 6
Interim responses	: 1
Stop responses	: 2
Malformed responses	: 0
Bad authenticators	: 0
Requests pending	: 1
Request timeouts	: 0
Unknown responses	: 0
Packets dropped	: 0

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 280](#)

[show services l2tp summary \(LAC on MX Series routers\) on page 281](#)

[show services l2tp summary \(LNS on MX Series routers\) on page 281](#)

[show services l2tp summary \(LNS on M Series routers\) on page 282](#)

[show services l2tp summary statistics \(MX Series routers\) on page 282](#)

Output Fields

Table 11 on page 277 lists the output fields for the **show services l2tp summary** command. Output fields are listed in the approximate order in which they appear.

Table 11: show services l2tp summary Output Fields

Field Name	Field Description
Administrative state	Administrative state of the tunnel is drain. In this state you cannot configure new sessions, destinations, or tunnels at the LAC or LNS.
Failover within a preference level	State of this tunnel selection method on the LAC. When enabled, tunnel selection fails over within a preference level. When disabled, tunnel selection drops to the next lower preference level. Not displayed for LNS on M Series routers.
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.
Destination equal load balancing	State of this tunnel selection method on the LAC. When enabled, the LAC selects tunnels based on the session count for destinations and the tunnel session count. 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.

Table 11: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Failover Protocol	<p>When the state is enabled, the LAC operates in the default <i>failover-protocol-fall-back-to-silent-failover</i> 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.</p>
Tx connect speed method	<p>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:</p> <ul style="list-style-type: none"> • actual This is the default value in Junos OS Releases 15.1, 16.1, 16.2, and 17.1. It is deprecated in Junos Releases 17.2 and higher. • ancp • none • pppoe-ia-tag • service-profile • static This is the default value in Junos Releases 13.3, 14.1, 14.2, 17.2 and higher. It is deprecated in Junos OS Releases 15.1, 16.1, 16.2, and 17.1.
Rx speed avp when equal	<p>Indicates if the Rx connect speed when equal configuration is enabled or disabled.</p>
Tunnel assignment id	<p>Format of the tunnel name.</p> <p>Format of the tunnel name, based on RADIUS attributes returned from the AAA server:</p> <ul style="list-style-type: none"> • 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.

Table 11: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Tunnel Tx Address Change	<p>Action taken by LAC when it receives a request from a peer to change the destination IP address, UDP port, or both:</p> <ul style="list-style-type: none"> • accept—Accepts change requests for the IP address or UDP port. This is the default action. • ignore—Ignores all change requests. • ignore-ip-address—Ignores change requests for the IP address but accepts them for the UDP port. • ignore-udp-port—Ignores change requests for the UDP port but accepts them for the IP address.
Min Retransmission Timeout for control packets	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
Min Retransmission Timeout for control packets	Minimum number of seconds that the local peer waits for the initial response after transmitting an L2TP control packet. If no response has been received by the time the period expires, the local peer retransmits the packet.
Max Retransmissions for Established Tunnel	Maximum number of times control messages are retransmitted for established tunnels.
Max Retransmissions for Not Established Tunnel	Maximum number of times control messages are retransmitted for tunnels that are not established.
Tunnel Idle Timeout	Period that a tunnel can be inactive—that is, carrying no traffic—before it times out and is torn down.
Destruct Timeout	Period that the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed.
Reassembly Service Set	Indicates active IP reassembly configured for the interface.
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.

Table 11: show services l2tp summary Output Fields (*continued*)

Field Name	Field Description
Access Line Information	<p>State of LAC global configuration for forwarding subscriber line information to the LNS, Enabled or Disabled.</p> <p>Indicates active IP reassembly configured for the interface.</p> <p>Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for information it receives from the LAC.</p>
IPv6 Services for LAC Sessions	<p>State of LAC IPv6 service configuration for creating the IPv6 (inet6) address family for LAC subscribers, allowing the application of IPv6 firewall filters, Enabled or Disabled.</p>
Speed Updates	<p>State of LAC global configuration for including connection speed updates when it forwards subscriber line information to the LNS, Enabled or Disabled.</p> <p>Starting in Junos OS Release 17.4R1, this information can also be displayed on the LNS for updates it receives from the LAC.</p>
Destinations	<p>Number of L2TP destinations for the LAC. Not displayed for LNS on M Series routers.</p>
Tunnels	<p>Number of L2TP tunnels established on the router.</p>
Sessions	<p>Number of L2TP sessions established on the router.</p>
Switched sessions	<p>Number of L2TP tunnel-switched sessions established on the router.</p>
Control	<p>Count of L2TP control packets and bytes sent and received.</p>
Data	<p>Count of L2TP data packets and bytes sent and received.</p>
Errors	<p>Count of L2TP error packets and bytes sent and received.</p>

Sample Output

show services l2tp summary (LAC on M Series routers)

```
user@host> show services l2tp summary
```

```

Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Enabled
Destination equal load balancing is Disabled
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**

```

Administrative state is Drain
  Failover within a preference level is Disabled
  Weighted load balancing is Disabled
  Destination equal load balancing is Enabled
  Tunnel authentication challenge is Enabled
  Calling number avp is Enabled
  Failover Protocol is Disabled
  Tx Connect speed method is static
  Rx speed avp when equal is enabled
  Tunnel Tx Address Change is Accept
  Min Retransmissions Timeout for control packets is 2 seconds
  Max Retransmissions for Established Tunnel is 7
  Max Retransmissions for Not Established Tunnel is 5
  Tunnel Idle Timeout is 60 seconds
  Destruct Timeout is 300 seconds
  Destination Lockout Timeout is 300 seconds
  Reassembly Service Set is ssnr3
  Access Line Information is Enabled, Speed Updates is Enabled
  IPv6 Services For LAC Sessions is Enabled
  Destinations: 0, Tunnels: 0, Sessions: 0, Switched sessions: 0

```

show services l2tp summary (LNS on MX Series routers)

user@host **show services l2tp summary**


```

Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is static
reassembly Service Set is ssnr3
Destinations: 4, Tunnels: 19, Sessions: 65, Switched sessions: 2
Access Line Information is Enabled, Speed Updates is Enabled

```

show services l2tp summary (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
```

```

Administrative state is Drain
Failover within a preference level is Disabled
Weighted load balancing is Disabled
Destination equal load balancing is Disabled
Tunnel authentication challenge is Enabled
Calling number avp is Enabled
Failover Protocol is Enabled
Tx Connect speed method is advisory
Tunnel assignment id format is assignment-id
Tunnel Tx Address Change is Accept
Min Retransmissions Timeout for control packets is 4 seconds
Max Retransmissions for Established Tunnel is 7
Max Retransmissions for Not Established Tunnel is 5
Tunnel Idle Timeout is 60 seconds
Destruct Timeout is 300 seconds
Destination Lockout Timeout is 300 seconds
Destinations: 1, Tunnels: 1, Sessions: 31815, Switched sessions: 0
      Tx packets  Rx packets  Memory (bytes)

```

Control	90.4k	32.0k	245678080
Data	127.3k	100.8kk	0
Errors	0	0	